Blueprint for the Digital University

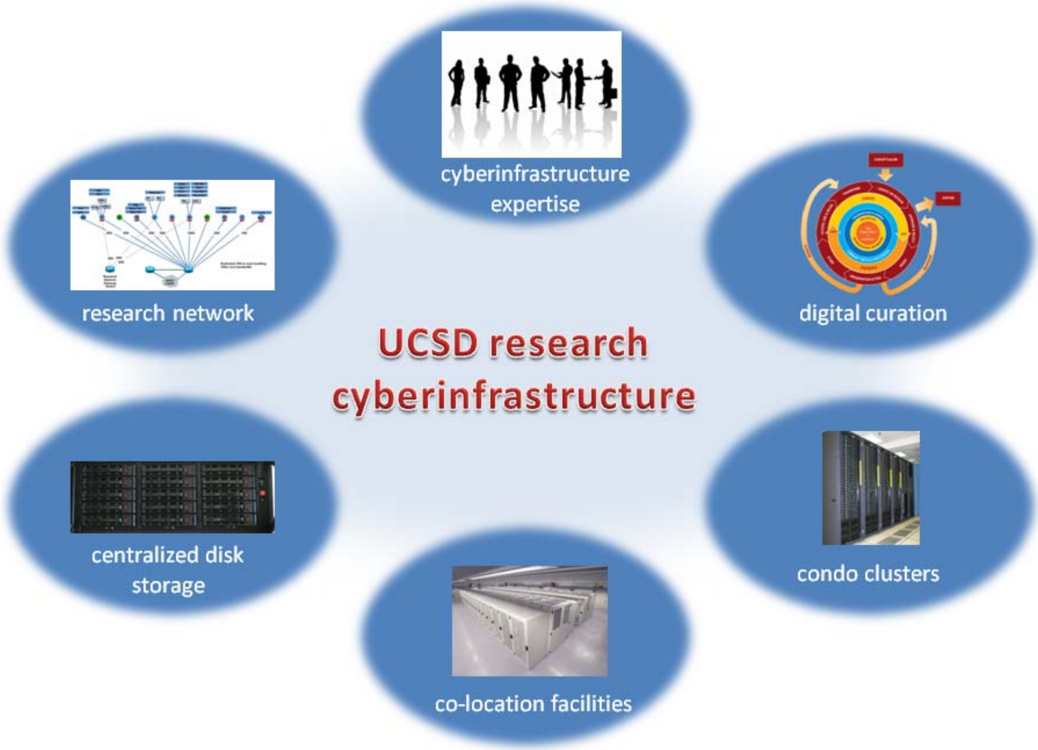

A Report of the UCSD Research Cyberinfrastructure Design Team

(4/24/09)

Contents

Executive Summary

This report provides the rationale and design for a campus-wide research cyberinfrastructure (RCI) that will meet federal mandates for research data preservation, improve UCSD's academic competitiveness, and achieve economies-of-scale savings through centralization of core infrastructure elements, while at the same time recognizing the diverse and distributed nature of UCSD's research enterprise.

Informed by a campus survey conducted in 2008, as well as a number of case studies analyzed by the Research Cyberinfrastructure Design Team (RCIDT), our design focuses primarily (but not exclusively) on *digital data*. Driven by proliferating high throughput instruments, wireless sensor nets, and supercomputer simulations, the amount of digital data at the heart of the modern research enterprise is growing exponentially at UCSD and around the world. Our proposed data-centric RCI will not only allow us to cope with the data deluge, it will position UCSD to prosper from it, and become an international leader and innovator in this area.

The proposed RCI design consists of 5 core elements which will require new campus investments to realize, and a sixth element that will require campus-wide coordination to effect. These design elements and our specific recommendations are as follows:

- <u>Colocation facilities</u>: Energy-efficient, centrally managed datacenter space for hosting computer equipment and related components from individuals, labs, departments, ORUs or other UC campuses, achieving economies of scale through capital and operating costs. SDSC has recently announced such a facility for the UC system (http://www.sdsc.edu/ucsd/colo.php) based on a cost-recovery model.

  The CIDT recommends that UCSD fund the use of *at least* 45 racks in this facility for near-term needs of campus researchers to freely host their equipment, and begin discussions on how to meet long-term needs.

- <u>Centralized disk storage</u>: A centrally administered disk storage farm for UCSD, to be housed in the colocation facility, which features high performance, high accessibility, high reliability, and scalability. Its low cost/Byte will make this an attractive and cost-effective alternative to UCSD researchers using commercial storage clouds or ad hoc local storage solutions that do not scale and do not provide for data replication and backup.

  The CIDT recommends an initial purchase of 2 PB of raw storage capacity to supplement the Data Oasis component of SDSC's Triton Resource, and operating funds to manage and scale up the UCSD storage resource to meet demand.

  Storage would be available to every UCSD researcher over the campus network for reliable data backup and long-term storage to meet federal requirements for digital data preservation.

- Digital curation and data services: Digital data is an important and growing part of the intellectual capital of the University.

  The CIDT recommends the establishment of the Research Data Depot, a suite of three core services designed to meet the needs of modern researchers. The three services are 1) data curation, 2) data discovery and integration, and 3) data analysis and visualization.

  Drawing on the combined expertise of the UCSD Libraries, CalIT2, and SDSC, the Research Data Depot will provide software tools and services to store, analyze, preserve, and disseminate research results in digital form to knowledge consumers in academia and industry.

- Research cyberinfrastructure network: An uncongested, leading-edge network that facilitates research collaborations, high-performance data exchanges, access to co-location facilities, remote mounts of storage, and real-time communications. The CIDT recommends that a Research Cyberinfrastructure Network (RCN) be built using fiber optic cables already in the ground that will be used by every research lab whose requirements go beyond the standard production network. The RCN will complement the standard NGN production network and will be designed for ultra-high performance.

  The CIDT recommends that the current RCN pilot be expanded, and requests funds to connect 25 buildings using 10 Gb/s Ethernet networking within the next several years. Funding and access philosophy would aim to encourage usage of the network.

- Condo clusters: Many UCSD research teams and departments regularly procure and deploy small to medium sized compute clusters to support their research needs. These are typically deployed as standalone resources in "campus computing closets" and are used by a small team of researchers. Many universities have found more efficient and strategic ways to deploy these resources on their campuses. One of the most interesting is the concept of "condo clustering". Condo clusters provide a means of unifying these compute resources by integrating several clusters into one or more larger clusters. Condo clusters allow the owners to leverage a pool of system administrators to manage the systems and other economies of scale (e.g., procurements, software licenses, security, and networking) that come with integrating into single, centrally administered cluster. From a campus perspective condo clusters allow the campus to make use of the idle cycles on the cluster by opening these up to the broader campus community (e.g., students, researchers). Another benefit of condo clusters is that they are housed in a co-location facility and thereby reap the benefits of a "green datacenter". This central location of a condo cluster also provides a strategic hub for Cyberinfrastructure activities that could result in a more competitive posture for educational and research opportunities.

  The CIDT recommends UCSD embrace the concept of condo clusters and exploit the deployment of the Triton Resource to launch the initiative.

- Cyberinfrastructure expertise: While not one of the five core components that form the focus of the CIDT's design recommendations and budget request, we acknowledge a campus-wide need for cyberinfrastructure expertise in areas not directly addressed by the Research Data Depot. As reflected in campus survey responses, there is need for advanced consulting services in software development, porting, and optimization; documentation; portal/interface development; database design; workflow systems; statistical analysis; etc. There is a considerable amount of expertise in these areas at SDSC, as well as other large ORUs on campus.

UC San Diego's pre-eminence in cyberinfrastructure will provide our researchers the unique opportunity to train tomorrow's leaders and stewards of data. Postdoctoral researchers, graduate students, and undergraduate students who participate in research with our faculty will benefit from the availability of the facilities here, and also from their integration into the activities. Future faculty, researchers, and entrepreneurs can be educated and nurtured in this environment.

The CIDT recommends that a coordinating body be established to maintain a labor pool of such experts and work out mechanisms that would allow customers to pay for their services.

The above recommendations would be implemented beginning in FY10 by a set of coordinated activities involving SDSC, UCSD Libraries, CalIT2 and ACT which would be overseen by a steering committee composed of resource/service providers and customers. The realignment of SDSC from primarily a national focus to a primarily UC focus, and the 2009 deployment of Triton and UC3 cluster resources, makes this a propitious time to begin implementing this plan.

The start-up cost and annual operating costs for the RCI are estimated to be $5.5M and $5.7M/yr respectively. After 5 years, the annual energy cost savings from the collocation facility are estimated to be between $1M and $2M.

The report discusses a number of mechanisms for recovering the additional costs, including an increase in ICR, an NGN-like fee, campus subsidies, and recharge. There is no single preferred mechanism given the diverse nature of the resources and services included in our design. Rather, a business model is suggested for each design element based on what has been proven to work elsewhere.

1. Introduction

What is Research Cyberinfrastructure?

The term "cyberinfrastructure" is an imprecisely defined term. For example, the NSF Blue Ribbon Advisory Panel on Cyberinfrastructure [1] defined it to be:

> *"Like the physical infrastructure of roads, bridges, power grids, telephone lines, and water systems that support modern society, "cyberinfrastructure" refers to the distributed computer, information and communication technologies combined with the personnel and integrating components that provide a long-term platform to empower the modern scientific research endeavor."*

On a practical basis, cyberinfrastructure (or CI) encompasses the computing, raw data storage, data analysis and preservation, networks, software technologies, know-how and people to support scientific research that uses any or all of the above technologies.

The CI Design Team (CIDT) wrestled with the definition of cyberinfrastructure, and more importantly, what elements, if any, should be treated in a systematic way by the university as a whole. Throughout the process of hearing from researchers both inside and outside of UCSD, evaluating results of our survey, and examining the practical impediments that computing/data-enabled researchers at UCSD were encountering every day, the target of our design and considerations focused primarily (but not exclusively) on digital data.

Over the past decade the amount of digital data has grown exponentially, driven by the ease with which digital sensors can produce it, and our ability to store it and manipulate it with commodity storage and computing devices. A large number of articles and reports have spotlighted the growth of digital data in research and public domains. So large is the amount of digital data generated—roughly doubling every year—that one speaks of the "Data Deluge" or "Data Tsunami" [5-8]. UCSD is at the forefront of this trend with the School of Medicine's acquisition of high-throughput DNA sequencers, CalIT2's emphasis on sensor nets, SDSC's access to petascale supercomputers via the NSF TeraGrid, and the Department of Physics role as a Tier 2 data center for the Large Hadron Collider (LHC), to name just a few examples.

Why Invest in Research Cyberinfrastructure?

> *"If [physical] infrastructure is required for an industrial economy, then we could say that cyberinfrastructure is required by a knowledge economy"*
>
> *Atkins Report, "Revolutionizing Science and Engineering through Cyberinfrastructure", [1]*

A major research university such as UCSD is a microcosm of the knowledge economy. Our scholars and researchers consume knowledge, they produce new knowledge, and they transmit knowledge to our students. One need only reflect on the revolution that the Internet has brought to all scholarly activities, including research and education, to recognize the transformative role of cyberinfrastructure in the 21st century. Since the Atkins report was published in 2003, CI has broadened beyond science and engineering to include economics, social sciences, and the arts and humanities. It is with this broader view of CI that the UCSD Research Cyberinfrastructure Design Team (RCIDT) has approached its charge.

There are numerous reasons for UCSD to invest in RCI:

- Competitive Advantage: Just as universities with the largest library collections held a competitive advantage of attracting the best scholars in the 20th century, the universities that lead in managing and exploiting vast digital holdings will establish leadership positions in the 21st century. As a recent case in point, a young faculty member currently being recruited by UCSD in the area of human genomics needs the ability to store and manipulate 100 TB of sequence data in connection with his research. UCSD's ability to support this need may be a key factor in his decision whether to come here or accept an offer from a competing institution [2]. This competitive advantage applies to attracting the best research staff and students too.

- Campus Demand: A survey conducted by the RCIDT, described in detail in Appendix A, indicated a widespread need for data cyberinfrastructure and services. We are seeing explosive growth in the amount of data being generated by automated DNA sequencers and protein mass spectrometers at UCSD, and simply storing this data is a major driver for RCI (see Genomics at UCSD sidebar). Data set sizes are growing in other disciplines as well, including those from earth-observing systems and petascale numerical simulations. In addition to storage, the survey revealed a substantial demand for high level data services, including collection management, long-term preservation, and scientific analysis/visualization.

- Leadership opportunities: UCSD has a golden opportunity to take national and international leadership in developing the research cyberinfrastructure of the 21st century. RCI is built on top of high performance networking, which UCSD has in abundance. Thanks to vision and good planning, UCSD has laid a large amount of 10 gigabit fiber optic networking cables all over campus. This supports an emergent campus research network that will connect research labs needing such high bandwidth with centralized data and compute resources located at SDSC and elsewhere. Changes in SDSC's focus from national- to campus-level service, the presence of CalIT2 on the UCSD campus, and the willingness of UCSD Libraries and ACT to enter into productive partnerships make this the right time to announce a campus initiative in RCI development and deployment. Certainly this initiative will make UCSD the leader within the UC system, and will provide a model for other campuses to emulate.

- Complementarity with national programs: The National Science Foundation (NSF) has historically led the investments in national scale cyberinfrastructure, including the creation of the NSFnet (a forerunner to the modern Internet) and the supercomputer centers program in the 1980s, the Partnerships in Advanced Cyberinfrastructure (PACI) program in the 1990s, and various Office of Cyberinfrastructure (OCI) programs in the 2000's. As emphasized in a recent Educause article [3], the NSF investments have tended to focus on optimizing the utilization of unique resources (e.g., supercomputers) for "extreme researchers", whereas campus cyberinfrastructure should address the needs of the more "typical" scholar. As one of the DT members noted, the goal of the campus effort should be to make every scholar at UCSD more competitive. Moreover, the NSF increasingly views the development of RCI as a shared responsibility, and looks to leverage local investments when making its awards. Thus, it is in an institution's best interests to develop its local RCI in order to be more competitive for winning external grants. It also should be noted that this kind of RCI will make us a preferred partner for other national and international collaborators.

- Preservation of digital knowledge is vital to the scientific method: Independent verification of results is fundamental to the scientific method. Historically, books and journals have served as the preservation medium, and these are still adequate if data volumes are not too large and the data can be presented in tabular form. Frequently, though, digital data is of such size and complexity that it cannot be preserved in this way, leading to the creation of data repositories in many disciplines (e.g., the Protein Data Bank (PDB) housed at the SDSC.) However, in the absence of federally-funded data archives, the responsibility falls to individual researchers to preserve the digital knowledge.

- Institutional obligations: Not only is the preservation of digital knowledge vital to the scientific method, it is required by federal law. OMB Circular A-110/CFR Part 215 stipulates that federally funded research data must be preserved by the grantee for a minimum of three years after the expiration of the grant. This requirement is not broadly known among researchers, nor has it ever been enforced to our knowledge [3]. Inasmuch as a grant is awarded to the institution and not the individual, the institution is obligated to enforce or at least inform researchers of their responsibilities. While complying with this requirement would not be particularly burdensome for an individual researcher for small quantities of data, amounts significantly in excess of a Terabyte start to become both a financial and a logistical burden since off-site data replication is needed for truly secure data preservation.

- Escalating energy/space demands: Continuing technological innovation and market forces in personal computing, electronic entertainment, and telecommunications have made high performance computing, high speed networking, and data storage technologies cheap and ubiquitous. A well-funded research group can now afford to procure and operate what is essentially their own private data center. Many have done so at UCSD and at universities across the country. This has led to escalating energy and space demands which at UCSD is reaching critical proportions. Centralizing some of the common resources and services into a campus collocation facility achieves economies of scale in utilities, operation, and administration, and is a key element in the "greening" of UCSD.

- Integration with UC-wide initiatives: The demand for data storage and computing is escalating at all the UC campuses, and administrators are looking for system-wide solutions that offer even greater economies of scale. Under the leadership of David Ernst, UCOP is embarking on establishing a network of system-wide colocation facilities, the first of which will be at SDSC on the UCSD campus. To meet growing high performance computing needs, UCOP is funding the North-South Cluster—two HPC computing clusters to be located at LBL and SDSC. By investing in RCI, UCSD will better leverage these investments for the benefit of its faculty, staff, and students, but also have the opportunity to influence the direction of the UC-wide CI.

The RCIDT's proposed data-centric cyberinfrastructure directly addresses these opportunities and needs, as detailed in the following sections.

Why Invest Now?

We estimate a $5.5M startup cost, and an annual cost of $5.7M/yr (before Colo energy savings) to fully implement the recommendations of the CIDT. With state imposed budget cuts on the University of California, and the pressure that that imposes on UCSD administrators, one needs to ask the question "why invest now when funds are so scarce?" There are several answers to this question:

- <u>Doing more with less:</u> the most compelling reason is that diminishing state support means the university must use its resources more efficiently if it is to remain competitive and move forward: it is a way of doing more with less. The economies of scale inherent in centralizing functions like hosting and administering computer equipment in energy-efficient data centers will save the university money.

- <u>Empowering our researchers to do the creative work they were hired to do:</u> creative faculty and staff on the cutting edge of data-intensive research, such as in genomics or earth science, should not have to figure out how to build a data center and staff it in order to pursue their science objectives. Those who do run the risk of becoming bogged down in operational concerns that distract from their research. By investing in an extensible RCI, UCSD will boost the productivity of their researchers and increase their competitiveness relative to their peers while exposing the next generation of scientists to the research possibilities.

- <u>Exploit our under-developed synergies:</u> UCSD has all the elements to assemble a world class RCI that can serve as a blueprint for the entire UC system. We are at a unique moment in time with SDSC reinventing itself as a CI resource for UCSD and the UC system, with partnerships emerging between UCSD Libraries, CalIT2, SOM, JSOE and ACT, and UCOP launching UC-wide IT initiatives (co-location, UC3 cluster). Now is clearly the time to act.

- <u>SDSC deployment of the Triton Resource:</u> In spring 2009 SDSC will deploy the Triton Resource (see Sidebar) for the benefit of UC researchers. Triton will consist of a trio of powerful resources integrated together to provide an unparalleled capability for storing and analyzing massive amounts of data. Two of the three components--Data Oasis, and the Shared Resource Cluster--can be thought of as capitalizing part of the recommended RCI—specifically, the UCSD storage and condo cluster components, respectively. The Triton resource will be accessible to UCSD researchers via the campus production and research networks, and drive the growth in demand for the latter.

- <u>The longer we wait, the harder it will get, and the more opportunities will be lost:</u> research data volumes will continue to grow exponentially, and with it, more and more groups will find that they are devoting more and more time and resources to warehousing their data, when they should be focusing on scientific research, training their students, and disseminating their results. The longer we wait to put in place an integrated framework for meeting the basic needs of all data-intensive researchers, the harder and more difficult it will be to accomplish. Think of how interstate commerce blossomed with the development of the Interstate Highway System in the 1950s, and now imagine how economic development would have been retarded if its development had been deferred by several decades.

Finally, we note that the pain being felt in California is being felt everywhere in the US to a greater or lesser extent: it's a bad time for everyone: Business leaders instruct us that it is important to invest when you are least able to; those that do will come out stronger than their competitors who did not invest when the times were tough.

2. Emerging Needs

Results of the 2008 RCI Survey

In April 2008 the CIDT conducted a survey of UCSD faculty and researchers regarding their cyberinfrastructure needs. An online survey was constructed consisting of 19 questions covering demographic information, current usage, and future needs in computing, storage, networking, and data services with respect to both resource needs as well as human expertise. The questions and survey results are provided in Appendix D. Here we wish to summarize the principal findings.

The survey netted 45 respondents, corresponding to 17 individuals, 18 groups, and the remaining 10 representing centers, departments, and ORUs. Departments included Pediatrics, Bioengineering, Cognitive Science, Computer Science and Engineering, and Physics. Divisions/Schools included SIO/IGPP, SOM, JSOE, UCSD Libraries, Biostatistics/Bioinformatics, Social Sciences, and Geosciences.

The number one need expressed was in the area of data management. The entire stack of data management resources and services were requested, from data back-up to collection management and long-term preservation. Over 80% of the respondents cited data backup as their principal data need; 70% of respondents cited the need to store and analyze large quantities of research data; 64% of respondents expressed a need for long-term data preservation. Roughly 50% of respondents expressed concern about the ability to move their research data from where it is generated to their desktops where it is analyzed, and the ability to share research data with others.

9.	Data Management, Analysis, Storage, Archival, and Access Needs: What are your greatest data management, analysis, access, archival, and storage needs? (Please select as many as apply)		
Storage capacity, more short term (1-3 years) storage for research data		31	70%
Ability to process and manage large quantities of research data		31	70%
Data management software (Oracle, MySQL, SRB/iRODS, ...)		16	36%
Data analysis software (SPSS, SAS, Cognos)		15	34%
Transferring experimental data to storage facility		24	55%
Transferring data from storage to desktop or cluster		23	52%
Sharing your data collection with colleagues (via web or resources such as SDSC DataCentral)		22	50%

Access to national or community repositories (e.g., PDB, NVO, GenBank)	15	34%
Data Backup	37	84%
Long term data preservation of large data sets	28	64%
Metadata creation for large data sets for archival purposes	24	55%
Long term access via a common repository	20	45%
Data/format compatibility	16	36%
Meeting data privacy/security requirements (FISMA, HIPAA)	16	36%
Not applicable	1	2%
Other, please specify View Responses	6	14%

Figure 2.1: Responses to RCI survey question 9 about data management needs.

We surveyed for projected data storage needs over the next 2-3 years; respondents were asked to pick from a list of data volume ranges. The most common response (51% of respondents) was for a storage need of 2 – 100 TB, and seven respondents (17%) said they needed more than 100 TB each. None said that they needed more than 1 PB (1000 TB). Taking into account survey incompleteness, and the short doubling time of data demand, the CIDT estimates that the entire campus needs 5 to 10 PB of raw storage over the next 2 years, and 25 to 50 PB over the next 5 years.

The number two need expressed was for access to a broad spectrum of cyberinfrastructure services, including interface/portal development, database design, scientific visualization, and software development/optimization. Fig. 3 summarizes the responses from RCI survey question 16.

In the area of computing, demand for high-performance computing (HPC) was clearly present and growing. There was considerable interest in Co-Lo and Condo solutions for compute cycles; less so for an allocated campus "supercomputer". There was also a clearly expressed need for a large memory, I/O intensive architecture for massive data analysis, which is not surprising considering the data volumes anticipated.

Generally, the campus network was judged to be "good", but that it will soon fall behind bandwidth demand growth in general, and point-to-point bandwidth requirements in particular. Some inequities in connectivity were noted, which is to be expected since not all buildings are yet connected to the 10 Gb/s research network.

Finally, concerning how UCSD might improve its research cyberinfrastructure, there was general but not unanimous support for pursuing centralized solutions as they would:

- admit economies of scale
- allow more integration and coordination
- outsource technical problems to professionals
- relieve saturated facilities and understaffing problems
- solve some software availability problems

Considering the array of instruments and compute clusters already on campus, this suggests a highly networked CI architecture connecting distributed and centralized resources.

16. Cyberinfrastructure Services – Expertise/Programming: What types of cyberinfrastructure expertise or programming support would be most useful to your group, lab, department, ORU or other identifiable campus "unit"? This could be in the form of programming and staff support, managing a software stack, and others. What kind of advanced expertise would be of value (visualization, parallelization, database design)? (Select all that apply.)

Interface/Portal development (GUI, Web-based)	22	55%
Database and/or data management support (e.g., Schema design, management, tuning)	25	62%
Scientific programming/Modeling	15	38%
Visualization (Scientific, Medical, etc.)	23	58%
Managing a Software Stack (builds, revision control, ...)	11	28%
Statistical support (e.g., survey design, analysis)	11	28%
Software tailoring (e.g., porting code, scripting)	14	35%
Software parallelization/optimization for clusters	15	38%
Technical Documentation	13	32%
Not applicable	3	8%
Other, please specify View Responses	7	18%

Figure 2.2: Responses to RCI survey question 16 about cyberinfrastructure services needs.

Case study: High throughput DNA sequencing

The year 2001 saw the completion of the first draft sequence of the human genome. However, the genome sequencing efforts did not stop that year, but instead accelerated at an astonishing speed. This development is largely driven by several factors. First, the expense of DNA sequencing has decreased by several orders of magnitude over the last eight years. For example, the total cost to sequence the entire 3 billion base pairs of the human genome has been reduced from above $10 million in 2001 to less than $10,000 using today's technology. The time of genome sequencing has also been lowered from many months to days. Second, there has been an explosion of applications where new generations of DNA sequencing technologies are being used to address the fundamental problems of biology and human diseases. For example, sequencing of cancer cell genomes has revealed novel structural and sequencing variations that may underlie the development of different tumors. In addition, the combination of a standard biochemical technique, such as chromatin immunoprecipitation, and ultra-high throughput sequencing has allowed scientists to unravel the protein binding sites along the entire human genome at high resolution and determine the biological function of the vast amount of sequences that were previously deemed "junk DNA". Third, the development of personalized medicine calls for even more genome sequencing.

To stimulate the development of the genomics field and revolutionize biomedical research, the National Institutes of Health and other federal agencies have initiated a number of large-scale projects in recent years. For example, a $190 million epigenome roadmap project has recently been started, along with the ENCODE project with a similar funding scale. In addition, the National Cancer Institute has just announced a new $90 million program to characterize the genome, epigenome and transcriptome of over 10,000 tumor samples. In the future, such large-scale funding opportunities are likely to continue to exist, if not increase in numbers and scale.

There is a strong community of talents and expertise in genomics on the UCSD campus. One of the four epigenome centers in the Epigenome Roadmap project is led by UCSD researchers. In addition, several UCSD investigators have been involved in the cancer genome sequencing project and the ENCODE project. The ongoing genomic activities pose a challenge to the existing UCSD computing infrastructure, and the challenges are only going to increase as more genomics-oriented projects are funded.

Take the San Diego Epigenome Center, for example. The center is directed by Dr. Bing Ren, an Associate Professor of Cellular and Molecular Medicine at the UCSD School of Medicine. It consists of five participating laboratories from UCSD, Salk, University of Wisconsin-Madison and Cold Spring Harbor Laboratory. The main research goal of the center is to generate reference maps of the epigenome for a set of primary human cells types, including human embryonic stem cells and differentiated cell types. The term “Epigenome” refers to chemical modifications to the DNA and nucleosomes that regulate the expression of genetic information in the DNA. By combining biochemical methods with ultra-high throughput DNA sequencing, Dr. Ren’s group is going to determine the nucleosome modification states of the human genome. In addition, Dr. Joe Ecker of the Salk Institute, a participating member of the Epigenome Center, is also using ultra-high throughput DNA sequencing to determine the methylation state of every cytosine in the genome. To accomplish these goals, the center has installed six Illumina GA II analyzers, with each GAII capable of generating approximately 1 Terabyte of raw data in a single run. Since typically two runs are performed each week for each system, the total data output and storage need is approximately 12 terabytes per week, or 600 Terabytes per year. It is expected that about 3 Petabytes of data will be generated by the epigenome project in the 5-year research span.

Besides high data storage demands, the Epigenome Center’s activities also pose great challenges to computing. Currently, it would take approximately 30 hours using an 8 CPU 64 bit linux server with 24 GB of RAM to process the 1 Terabyte of raw data into a format that is easy to visualize and analyze. Clearly, faster and more powerful servers can reduce the time and shorten the data processing time.

Additionally, there is a need for the data to be preserved in a manner that can be referenced long after the 5-year research duration, or even after the Center ceases to exit. Currently, this task is preserved on local computer drives, and there is no clear mechanism for outside investigators to access the raw data.

The San Diego Epigenome Center is only one of many large and small federally funded projects in UCSD that use ultra-high throughput sequencing technologies. Beside the six Illumina GA II analyzers, there are an additional five similar systems in the UCSD campus at the time of this report. In the future, the sequencing capacity is expected to grow rapidly, as newer generations of sequencing technologies mature.

Case study: The Brain Observatory

At The Brain Observatory at UCSD, researchers combine multiple complementary neuroimaging techniques, including specialized applications of Magnetic Resonance Imaging (MRI), computer-controlled microscopy, and computational methods for image registration and 3-D reconstruction to elucidate the detailed structural design of the brain and how it is perturbed by neurological disease. Research ranges from designing novel anatomical and neuroimaging protocols to the application of extreme visualization environments that are necessary to interpret the extraordinarily large and complex data generated by our methodologies.

Because of its expertise and unique instrumentation, The Brain Observatory at UCSD is recognized nationally and internationally as a center where brain specimens of notable clinical importance are analyzed in detail. These are patients or individuals who were studied during their lives on account of very informative cognitive alterations, and whose generous donations allow us to study the neurological (microscopic) basis for their conditions. The Brain Observatory's anatomical collection is comprised of histological slides and anatomical digital models. In this respect, the long-term mission of The Brain Observatory is to create and maintain a prospective collection that will serve as a multimedia and multi-resolution library for the human brain. This repository of histological material, models, and visualization technology are meant to serve the scientific community as well as the general public. Thus it is essential to optimize computational techniques for the acquisition, storage and digital dissemination of the image data via the web.

As a notable example, The Brain Observatory has been charged with the examination of the brain of perhaps the most important study patient in medical history: patient H.M. At a young age, H.M. underwent an experimental operation in 1953 for the relief of medically intractable epilepsy; since that time, he has been profoundly amnesic, virtually unable to consolidate memories for facts or people. The results of experiments in which he participated during his life revolutionized the study of learning and memory. H.M. has taught us that the critical neural substrate for long-term declarative memory is located in the medial part of the temporal lobes. He also taught us that there are different kinds of memory with different neural substrates, such as declarative memory (including episodic and semantic memory, impaired in H.M.), skill learning (preserved in H.M.), and priming (also preserved in H.M.). Essentially, H.M.'s case first suggested that the establishment of memory has a distinct neural substrate. He is indeed the yardstick against which the severity of amnesia is measured in other patients. His case is cited in every psychology or neuroscience textbook and in thousands of scientific papers.

Henry Gustav Molaison (H.M.) died on December 2, 2008 at age 82. Researchers at The Brain Observatory will soon conduct the complete and systematic histopathological study of his brain in order to reveal the precise nature of his lesion and to quantify microstructural changes in the entire brain. A collection of many hundreds of large-format histological slides will be created, and each will be digitized at cellular resolution to create a comprehensive digital map of the brain. Each image representing a whole slice though the brain will actually be composed of thousands of frames that are acquired systematically at high magnification and stitched together into a single montage. The library of images generated by the project will be made publicly available via a dedicated web site to the neuroscience community and the general public. In addition to the brain of Mr. Molaison, a second map from an age-matched but neurologically normal individual will be available to users as a reference brain atlas on-line.

The project has attracted considerable attention from the media, and several news agencies will continue to follow the progress of the work. The project - the virtual dissection and digital illustration of the anatomy of the brain of HM and the control - will generate approximately 100 TB of image data. These data need to be stored permanently and, because of the visibility of the project, managed and distributed flawlessly via dedicated web servers. Building the massive storage and high-performance computing server infrastructure to propel projects such as this is not part of the mission of The Brain Observatory. Instead, it would be highly desirable to leverage centralized resources as proposed by the CI Design Team for the computing and storage needs and to connect our microscopy suite and scanners to campus via a dedicated 10GB pipe. The efficiency in image acquisition processing and analysis would be increased tremendously if data could be streamed directly into the SDSC servers and visualized on our own Optiputer power-wall set-up.

The brain of Mr. Molaison is only one of the many cases that The Brain Observatory at UCSD will receive in the relatively near future. We would like the case of H.M. to be supported and used as a pilot to illustrate the requirements and potential of such computationally demanding extreme anatomical projects. This is in the short term. In the long term, these projects will constitute a scalable Digital Library for the Human Brain at UCSD, which is an extremely innovative and unprecedented concept in Clinical Neuroscience.

In summary, the specific needs of The Brain Observatory in the context of the Digital Brain Library Initiative are: advanced computing power to improve digitization techniques for scanning and producing brain scans; a large amount of storage space for permanent residence of the digital collection for future generations of researchers; help with maintaining a web-based information portal for researchers and the general public to access the collections; and advanced information work for synthesizing and understanding the data.

Case study: Music, Visual Arts, and Theatre and Dance

UCSD has been one of the most forward-looking academic institutions in the world in relation to the application of the computing technologies to the arts, both in terms of research and production. The Center for Music Experiment (CME) was established at UCSD in 1972 and resulted in one of the first operating system-like software suites for audio processing and production. For many years, this system, which was called the Computer Audio Research Laboratory (CARL), and its variants were used all over the world as the most advanced framework for computer music research and production. Later, the facilities of CME were transformed to become the Center for Research in Computing and the Arts (CRCA), an Organized Research Unit (ORU) which brings researchers and faculty from various disciplines, such as Music, Visual Arts, Theatre and Dance, Computer Science, and Electrical Engineering together in many collaborative forms. The mission of CRCA is to facilitate the invention of new art forms that arise out of the developments of digital technologies. Current areas of interest include interactive networked multimedia, virtual reality, computer-spatialized audio, and live performance techniques for computer music and graphics. In 2005, after many years of planning, CRCA's facilities became a part of the New Media Layer of the California Institute for Telecommunications and Information Technology (Calit2).

One of the most important aspects of research in the arts at UCSD is the integration of advanced theoretical research with production. This requires a powerful and reliable computing infrastructure. The music department can boast to be in the forefront of real-time audio processing research, where computer and acoustic musicians collaborate on devising algorithms for synthesis, audio spatialization, signal processing, and compression, and on using such algorithms and software in large-scale and high-profile productions. The research in the Visual Arts department ranges from algorithmic production of high resolution visual artifacts and images, to developing new theories and frameworks for new cinema and new media, to the application of a wide range of technological discoveries to the critical discourse of performance. The Theatre and Dance department at UCSD is ranked as one of the best schools of theatre in the country. With the addition of two technology-driven programs, namely the MFA in Sound Design and MFA in Dance Theatre, the T&D department is committed to the use of algorithmic formal compositions and high bandwidth interactive technology in its professionally produced series of productions. Research areas include animation, algorithmic image and sound synthesis, real-time spatialization, and telematic performances/presence involving high fidelity audio and video network connections.

Research in the arts requires large amounts of storage, especially for the digitizing and archiving of the audio and video material. A typical storage requirement for a compressed HD stream with 24 channels of 24 bit, 88.2KHz audio is between 0.8 to 3 gigabyte per minute (depending on the compression level). A performance/piece requiring 10 HD streams with similar audio material would require between 4.5 to 25 gigabyte of data per minute. With the introduction and prevalence of 4K video technology, the bandwidth requirements will be multiplied by orders 5 to 10 in the near future.

While the storage is an important element in working with media, having access to the data through a guaranteed constant bandwidth network is also necessary. This is true both for streaming of the data during post-production and editing periods, and more importantly in performative context. Telematic performances are fast becoming popular in modern performance spaces. A considerable amount of research and production in this area has already been carried out at UCSD in collaboration with Stanford University, UC Berkeley, and Rensselaer Polytechnic Institute. A strong cyberinfrastructure with adequate storage and reliable constant bandwidth network access would not only provide the researchers in the arts a context to carry out their research more efficiently, but also allow them to devise new forms of unique and unprecedented production forms and aesthetics within professional and high-profile contexts. This will allow UCSD to continue to hold its status as one of the pioneers of the use of digital technologies in the arts.

Case study: The Center for Observations, Modeling, and Prediction at Scripps

(COMPAS, http://www.compas.ucsd.edu) was formed as a research core to address critical and cross-cutting issues in coupled ocean-atmosphere and ocean-atmosphere-land modeling that include physical process-oriented modeling for small-scale ocean physics, ocean modeling and state estimation. The mean state and variability of the coupled ocean-atmosphere-land system is a matter of crucial scientific and societal importance and is one of the strategic research directions at the Scripps Institution of Oceanography (SIO). Gaining fundamental understanding of the processes that result in changes to rain and snowfall patterns is critical in areas like the coastal desert region of Southern California. The changes can cause persistent droughts or flooding and have direct and often deleterious effects on agriculture and hydroelectric power supplies.

Precipitation-driven changes in hydrology (the actual distribution of water above and below ground) influence the susceptibility of arid regions to destructive wildfires or of hilly areas to mudslides. The ocean itself directly contributes to the long-term climate by influencing the transport of heat and moisture and by exchange and uptake of gases such as CO2. Small-scale physics such as the vertical mixing and upwelling in oceans play an important role in fisheries by determining the change and distribution of nutrients. Regional modeling allows COMPAS to address the concern of near-shore water quality and ecosystem health concerns that have driven the need to better understand the processes that move and mix water and constituents.

In 2008, COMPAS was awarded an NSF Major Research Instrumentation Grant to build a dedicated high-performance cluster supercomputer to support numerical modeling aimed at understanding and ultimately predicting physical processes and phenomena in the ocean, atmosphere, and land surface. The cluster itself consists of 128 dual-socket, quad-core nodes with a Myrinet high-performance network. Simulation runs last from hours to weeks depending on the problem being investigated. While this cluster is critical to meeting the scalable computation needs of COMPAS researchers, there are two notable areas in which our infrastructure would benefit greatly from additional campus capability. The first is long-term, high-performance storage that could be accessed directly by our dedicated facility, and the other is the occasional use of large-memory footprint nodes as envisioned in the design team architecture being deployed in the large-memory cluster of the Triton resource (see Triton sidebar).

In the arena of storage, the researchers in COMPAS already consume more than 30 Terabytes of space for model input files (e.g., bathymetry, temperature inputs, external forcing data, and more), intermediate files to support the forward-inverse iteration of our models, and simulation output. While the quantity of storage (with today's technology) is not out of reach, the storage performance is technically more difficult to achieve and requires a significant localized system expertise to achieve. Our current, local-only storage setup just meets our current needs but is likely to fall short as we continually explore more detailed scenarios.

Large memory systems, those with 10-100 times more memory than our COMPAS cluster nodes, can play a critical but infrequently-used component of our research infrastructure. A particular step in several scenarios involves a global optimization step that runs rapidly if enough memory (~100GB today) is available on a single machine. Today, we utilize single-CPU, large-memory runs at remote providers like NCSA. But this mode is quite cumbersome and manually driven with a researcher logging in and submitting jobs directly at the remote site. The CI design structure of including large-memory nodes in the proposed infrastructure will enable us to perform this critical step with significantly less effort and faster turn-around time.

3. Overview of Proposed RCI

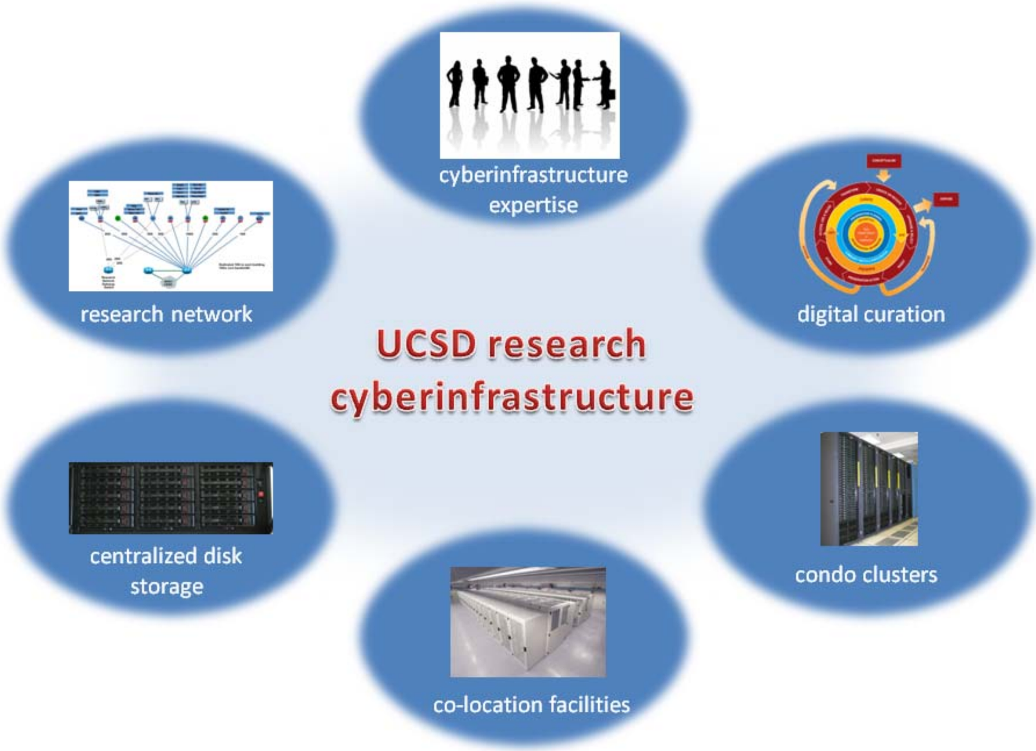

Figure 3.1: Elements of the proposed UCSD research cyberinfrastructure (RCI).

Elements of the RCI

The case studies presented in Section 2 highlight a number of different research cyberinfrastructure (RCI) needs at UCSD, including the ability to process and store very large data sets from instruments and simulations, build and share unique digital collections, and schedule dedicated high-bandwidth network paths for digital performance in the Arts. In addition, there is a widespread need for a broad array of advanced consulting services for applying RCI elements to specific needs (Cf. Fig. 2.2). In this section we introduce the elements of our proposed RCI, deferring a detailed discussion of each to sections 4-10. Figure 4 encapsulates the CIDT's vision for what is required to support digital scholarly activities at UCSD in the coming decades. We have chosen to focus on *general solutions to meet common needs*. In this way, it is hoped that UCSD investments in our proposed RCI will have the widest possible impact for all researchers.

A brief description of the five core elements, plus a sixth follows.

- Co-location facilities: Energy-efficient, centrally managed datacenter space for hosting computer equipment and related components from individuals, labs, departments, ORUs or other UC campuses, achieving economies of scale through capital and operating costs. SDSC has recently announced such a facility for the UC system (http://www.sdsc.edu/ucsd/colo.php) based on a cost-recovery model. The CIDT is recommending UCSD fund the use of this facility for campus researchers for their near-term needs and begin discussions on how to meet long-term needs.

- Centralized disk storage: A centrally administered disk storage farm for UCSD, to be housed in the co-lo facility, which features high performance, high accessibility, high reliability, and scalability. Its low cost/Byte will make this an attractive and cost-effective alternative to UCSD researchers using commercial storage clouds or ad hoc local storage solutions that do not scale and do not provide for data replication and backup. The CIDT is recommending an initial purchase of 2 PB of raw storage capacity to supplement the Data Oasis component of the Triton Resource, and operating funds to manage and scale up the UCSD storage resource to meet demand. Storage would be available to every UCSD researcher over the campus network for reliable data backup and long-term storage to meet federal requirements for digital data preservation. Researchers' compute clusters, whether housed in their labs or in the co-lo facility, would be connected to the storage farm via the Research cyberinfrastructure network.

- Digital curation and data services: Digital data is an important part of the intellectual capital of the University. The CIDT recommends the establishment of the Research Data Depot, a suite of three core services designed to meet the needs of modern researchers. The three services are 1) data curation, 2) data discovery and integration, and 3) data analysis and visualization. Drawing on the combined expertise of the UCSD Libraries, CalIT2, and SDSC, the Research Data Depot will provide software tools and services to store, analyze, preserve, and disseminate research results in digital form to knowledge consumers in academia and industry.

- Research cyberinfrastructure network: An uncongested, leading-edge network that facilitates research collaborations, high-performance data exchanges, access to co-location facilities, remote mounts of storage, and real-time communications. The CIDT recommends that a Research Cyberinfrastructure Network (RCN) be built using fiber optic cables already in the ground that will be used by every research lab whose requirements go beyond the standard production network. The RCN will complement the standard NGN production network and will be designed for ultra-high performance. The CIDT recommends that the current RCN pilot be expanded, and requests funds to connect 25 buildings using 10 Gb/s Ethernet networking within the next several years. Funding and access philosophy would aim to encourage usage of the network.

- Condo clusters: Many UCSD research teams and departments regularly procure and deploy small- to medium-sized compute clusters to support their research needs. These are typically deployed as stand-alone resources in "campus computing closets" and are used by a small team of researchers. Many universities have found more efficient and strategic ways to deploy these resources on their campuses. One of the most interesting is the concept of "condo clustering". Condo clusters provide a means of unifying these compute resources by integrating several clusters into one or more larger clusters.

Condo clusters allow the owners to leverage a pool of system administrators to manage the systems and other economies of scale (e.g., procurements, software licenses, security, and networking) that come with integrating into a single, centrally administered cluster. From a campus perspective condo clusters allow the campus to make use of the idle cycles on the cluster by opening these up to the broader campus community (e.g., students, researchers). Another benefit of condo clusters is that they are housed in a co-location facility and thereby reap the benefits of a "green datacenter". This central location of a condo cluster also provides a strategic hub for cyberinfrastructure activities that could result in a more competitive posture for educational and research opportunities. The CIDT recommends UCSD embrace the concept of condo clusters and exploit the deployment of the Triton Resource to launch the initiative.

- Cyberinfrastructure expertise: While not one of the five core components that form the focus of the CIDT's design recommendations and budget request, we acknowledge a campus-wide need for cyberinfrastructure expertise in areas not directly addressed by the Research Data Depot. As reflected in campus survey responses, there is need for advanced consulting services in software development, porting, and optimization; documentation; portal/interface development; database design; workflow systems; statistical analysis, etc. There is a considerable amount of expertise in these areas at SDSC, as well as in other large ORUs on campus. The CIDT recommends that a coordinating body be established to maintain a labor pool of such experts and work out mechanisms that would allow customers to pay for their services.

Summary of new services

Table 3.1 summarizes the new services that would be provided by our proposed RCI. These are assembled from the detailed discussions in Sections 5-9.

RCI component	Services
Co-lo facility	Basic services • Hosting of racks, servers • Standard in-rack networking • Remote hands Additional services available • System setup, trouble shooting, monitoring • System administration • Cross-platform backup • Storage solutions • Special networking
Centralized storage	• high reliability, persistent, replicated storage accessible from authenticated workstations and laptops on the UCSD network laboratory- and department-owned clusters; primary/secondary storage for data-intensive instruments, higher-level data preservation services, and very high-performance shared resource facilities such as the Triton Resource

	• 1 TB replicated storage for every researcher • Larger amounts of storage at incremental cost • Storage condo program (owner-purchased storage)
Research Data Depot	Data Curation Services • Assistance in the development of data management plans • Data transfer/ingest support services • Metadata creation support services • Data preservation services • Data replication services • Long-term data integrity check services Data Discovery and Integration Services • Data set registration • Data set provenance and possible reuse description • Data submission process support • Data set availability alerting services • Ontology development • Portal development and maintenance • Linking of data sets to resulting publications • Linking of data sets to resulting analysis, mining results, and visualizations Data Analysis and Visualization Services • Data set transfer and transformation • Database creation, management, and optimization • Data analysis support services • Data mining support services • Data visualization support services
Condo clusters	• Compute cycles on professionally managed HPC cluster • Strong systems support, networking, security • Economies-of-scale hardware procurement (1-2 annually) • Reduced software licensing cost • Access to peripherals (storage, etc.) • Free hosting for "condo owners" • Pre-emptive scheduling for "condo owners" • Cycle scavenging
Research network	• Uncongested high-speed access to researchers • High-speed access to central storage and computing • Buildings & labs connected as research needs arise • Infrastructure that is ready to flex as needs expand • Reasonable security controls

Table 3.1: Summary of services provided by the proposed RCI.

Budget summary

Table 3.2 summarizes the estimated startup and annual costs of the five core components of the proposed RCI, taken from Sections 5-9.

RCI component	Startup	Annual
Co-lo facility	$ --	$ 1,235,000
Centralized storage	$1,080,000	$1,606,400
Research Data Depot	$ 850,000	$1,123,000
Condo clusters	$ --	$1,173,030
Research network	$3,517,500	$ 625,000
TOTAL	**$5,447,500**	**$5,762,430**

Table 3.2: Cost summary.

Partners and Roles

The proposed RCI would be implemented through a partnership among SDSC, UCSD Libraries, ACT, and CalIT2. Table 3.3 summarizes the roles each institution would play in implementing and operating the proposed RCI. L indicates "Lead" organization, and P indicates "Partner" organization.

RCI component	SDSC	Libraries	ACT	CalIT2
Co-lo facility	L			
Centralized storage	L			
Research Data Depot	P	L		P
Research network	P		L	
Condo clusters	L			
Green cyber-infrastructure	P		P	L

Table 3.3: Lead (L) and partner (P) responsibilities for implementing and operating the proposed RCI.

SIDEBAR 1: UCSD Triton Resource

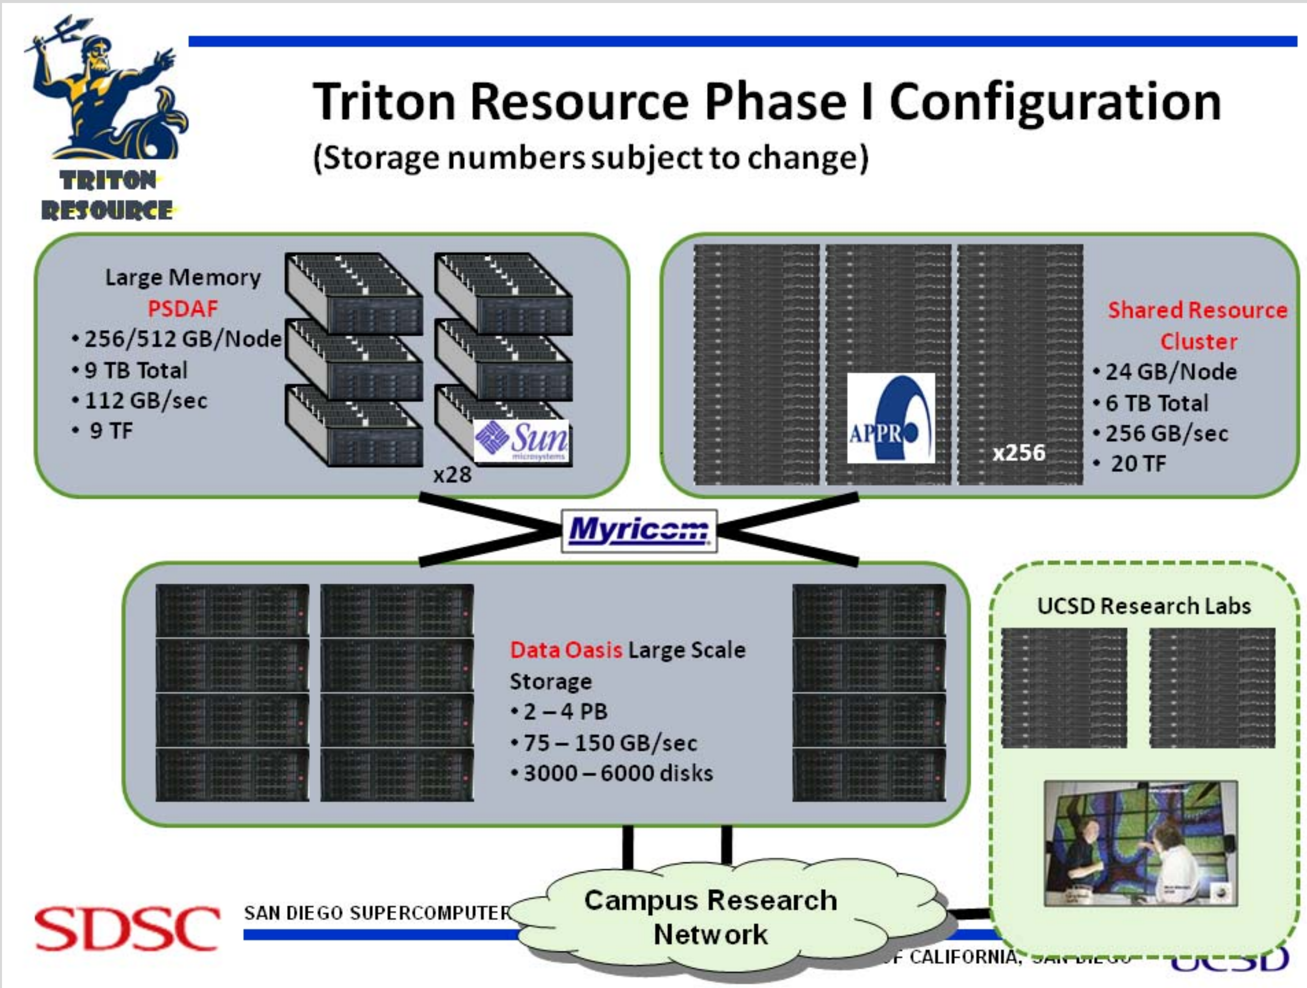

Fig. 1. Triton Resource Phase 1 configuration. Storage numbers subject to change.

Triton is a trio of computational resources being deployed in Spring 2009 for the use of UCSD and UC researchers. Triton consists of three elements: (1) the Petascale Data Analysis Facility (PSDAF)—a compute cluster with large memory nodes for massive data analysis applications; (2) the Shared Resource Cluster (SRC) –a standard HPC compute cluster with a larger number of smaller memory compute nodes; and (3) Data Oasis—large scale, high performance disk farm for staging large data sets for analysis on PSDAF, or for storing results computed on SRC. The three elements will be interconnected with a high speed network to maximize interoperability. The Triton Resource will be accessible from researchers' laboratories via the UCSD's 10Gbs research network. It is envisioned that each element could serve as a seed for a larger resource that would grow by accretion through faculty- and administration-funded procurements. For example, the SRC could be the seed for the Condo Cluster element of the proposed RCI.

4. Design Principles and Framework

Principles of Design

The design team recognized that there exists significant research unit, department, and individual laboratory cyberinfrastructure, and that new capabilities should complement and enhance the existing and inevitable diversity of needs across campus. The design presented here operates from the premise that *long-term stewardship of digital data is a critical responsibility of the University as a whole*. The traditional method for capturing scientific progress has been archival journal publications and conference proceedings. However, with modern instruments creating enormous volumes of data, scientific simulation creating data sets that require significant post-processing analysis, and sensors gathering the physical state of our environment, traditional journal publication is a woefully incomplete solution for experimental reproducibility. In short, *the digital data for modern research is the life-blood of the 21^st^ century University.*

Several members of the Design Team have been "in the trenches" for building clusters, file storage, networking, advanced display technologies, instruments, and data repositories at various campus sites. While some elements of this infrastructure are straightforward, the team was unanimous in its assessment that *large-scale data storage and data preservation represents the most difficult and person-intensive part of the infrastructure*. Campus laboratories are literally islands of storage. While some have reasonable reliability, an alarming number keep irreplaceable data on USB hard drives with no-backup and no redundancy. Further, even where data is handled "well", the treasure trove of digital data is often not available to others for inspection or new innovation.

Figure 4.1 shows a wiring diagram for our proposed RCI. This figure illustrates our first principle of design:

1. **Raw data storage, data discovery, and data preservation are critical missions of the university.**

In other words, our cyberinfrastructure is data-focused, but not exclusively. The second principle of design falls out easily:

2. **Centralize infrastructure when economies of scale are favorable, but recognize the inherent diversity of needs in cyber-enabled laboratories.**

This results in a RCI where various elements "plug into" or utilize a centralized storage facility. The elements could be larger community resources like a condo cluster or large-scale analysis facility, lab instruments and clusters, and data preservation/discovery services.

And finally,

3. **Centralized resources must be incrementally scalable with a base level of service.**

The dynamic range of needs of various labs is enormous. Some may need only small amounts of long-term data storage, while others may require truly huge resources. Financial reality requires a rational approach to supporting this range without all costs being borne exclusively by "community funds." The condo cluster and storage components represent practical approaches to having the base costs covered (such as administration, electricity, and first-bit hardware) while researchers can fund through extramural grants additional hardware to scale up these components.

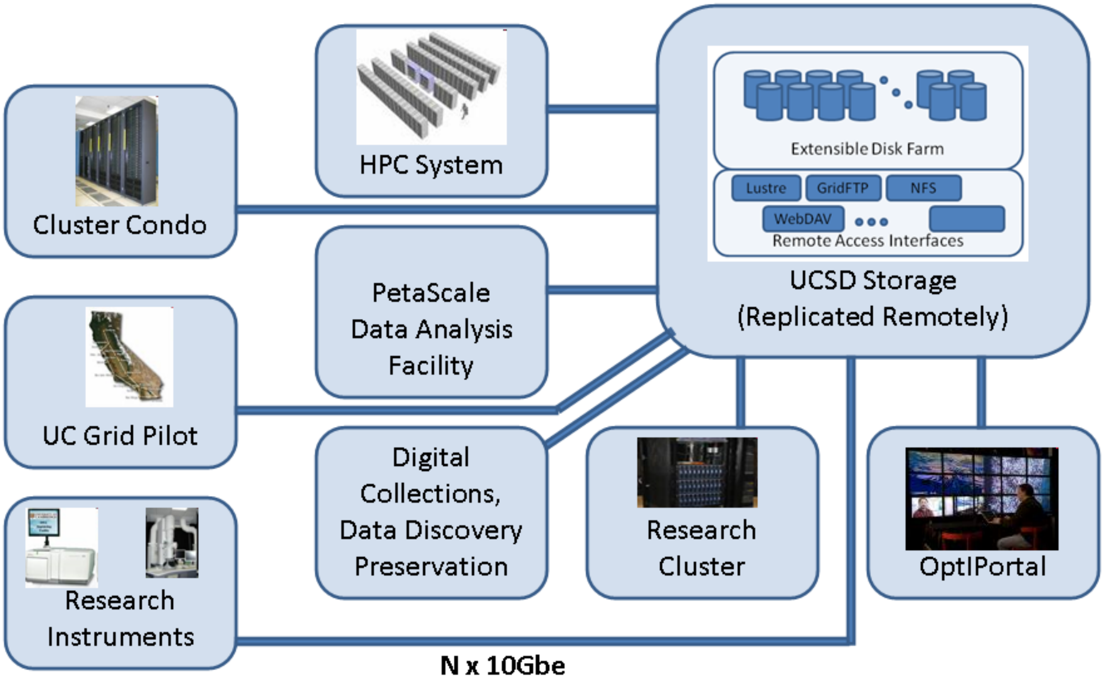

Figure 4.1 - Representative Components of Research Cyberinfrastructure (RCI). UCSD's investment in fiber-optic networking enables a practical focus on data storage, collections, discovery, and analysis.

Design Framework

Figure 4.1 shows a "wiring" view of the design team's proposed cyberinfrastructure. Based on experience throughout campus (and elsewhere), it is clear that management, analysis, preservation, and access to data is core to the research enterprise over a wide cross-section of campus researchers. However, highly- reliable, high-performance, and accessible storage eludes many due to the combination of 1) complexity of the underlying technology for ensuring data integrity and accessibility; 2) the relative simplicity for individuals to deploy single copy, unprotected, and unmonitored data-bit storage on their own; and, 3) the lack of critical mass of people and expertise within single projects to address the data problem. In short, *many researchers find it too difficult, too time-consuming, and too expensive to properly care for their data. This current state of affairs fundamentally puts the critical digital assets of the digital university in jeopardy.*

There are many examples where terabytes of precious data are stored on inexpensive external USB drivers. A single copy of the data exists and makes it vulnerable to both physical and software loss. The data is neither discoverable nor usable by others, making it difficult to reuse or, more importantly, be inspected by others as part of the scientific process. This same digital data is often the lifeblood of a research organization – it can hold irreplaceable output from lab instruments, the results of months or years of digital simulation, or the combined efforts of human and computer-aided data annotation. Unfortunately, the domain scientists often do not understand the risk of loss (or contamination) of their digital data, and even if they do, there simply often is not a time/technology practical way for them to address this risk.

The Design Team is making the bold statement that while single projects may not individually have critical mass to address data storage and analysis, the University, as a whole, does have critical mass. Moreover, the University *can* and *must* address the fundamental issue of data. Not doing this well represents a huge risk to the research enterprise. Doing this well will put the University in a leadership position of recognizing that *data handling as infrastructure* is a must for the 21st Century research institution.

Not the "Field of Dreams"

The Design Team was charged with an open-ended assignment: "come up with a practical roadmap for campus cyberinfrastructure." We recognized early in the process that we had to define the essentials of cyber-infrastructure so that it stayed within the confines of *supporting* peer-reviewed research and not become a mechanism to fund research on cyberinfrastructure itself. This is an important distinction to draw. For example, in our wiring diagram, UCSD RCI would support the data storage and analysis of data for an electron microscope but not the electron microscope nor the domain-specific research itself. The reality is that for RCI to work effectively, there is some software development and "gap-filling" that only dedicated research professionals can in reality achieve. RCI is not something you purchase, as a whole, off the shelf.

The team also focused on the items that we deemed to be essential/core elements of UCSD RCI

- raw data storage
- data collection, discovery, and preservation
- physical facilities (collocation)
- campus network extension for high-performance remote use
- condominium-style compute cluster

It should be noted that this does not include all elements of Figure 4.1. The structure is, by design, expandable and flexible to support a variety of modes of use.

It is important to realize that the cyberinfrastructure defined in this report is focused on the research mission (and the ways in which this educates the next generation of research scientists) and not the teaching or business components of the university. While the storage focus of the proposed design could be used to support these other key missions of the university, the additional complexity of data privacy, Sarbannes-Oxley compliance, and scaling to the *entire* campus community would add considerable expense and additional design time. It also would need to consider replacement of existing infrastructures, whereas our design is focused on what is missing to support research.

The Critical Importance of People

While we have presented a high-level wiring diagram, it essential to recognize the people component of any cyberinfrastructure for a top-tier research university. A reasonable rule-of-thumb is that for the core/essential infrastructure, approximately half the cost is hardware and other purchased items, the other half is the people to manage the infrastructure, assist users to get the most out of it, provide technical assessment of future technologies (both hardware and software), and generally to “fill the gap” to make CI a successful and integrated component of UCSD research.

Solid data cyberinfrastructure will require a great deal of thought and energy to answer fundamental questions of “What data is preserved and for how long?”, “Who pays for particular elements?”, “When is data publicly accessible?”, and “What happens when a researcher leaves the University?”

Other Notable Cyberinfrastuctures

UCSD participates in a substantial number of large-scale, national cyber-intensive research projects, including BIRN (Biomedical Informatics Research Network), OOI (Ocean Observing Initiative), GEON (Geosciences Network), NEON (National Ecological Observatories Network), TDLC (Temporal Dynamics Learning Center), ROADnet (Real-Time Observatories, Applications, and Data Management Network), CUAHSI (Consortium of Universities for Advancement of Hydrological Science), SEEK (Sharing Environmental Education Knowledge), and Teragrid. Many of these are domain-focused but have provided some level of national infrastructure footprint. However, Teragrid is the National Science Foundation’s National Computing Cyberinfrastructure

The Teragrid (www.Teragrid.org) consists of 11 major resource sites around the country, interconnected by a 10 gigabit/s dedicated network. Teragrid itself is where NSF invests $100M+ per year for supercomputers. According to the TeraGrid website:

> *“TeraGrid is an open scientific discovery infrastructure combining leadership class resources at eleven partner sites to create an integrated, persistent computational resource.”*

Using high-performance network connections, the TeraGrid integrates high-performance computers, data resources and tools, and high-end experimental facilities around the country. Currently, TeraGrid resources include more than 750 teraflops of computing capability and more than 30 petabytes of online and archival data storage, with rapid access and retrieval over high-performance networks. Researchers can also access more than 100 discipline-specific databases. With this combination of resources, the TeraGrid is the world's largest, most comprehensive distributed cyberinfrastructure for open scientific research.

TeraGrid is coordinated through the Grid Infrastructure Group (GIG) at the University of Chicago, working in partnership with the Resource Provider sites: Indiana University, the Louisiana Optical Network Initiative, the National Center for Atmospheric Research, National Center for Supercomputing Applications, the National Institute for Computational Sciences, Oak Ridge National Laboratory, Pittsburgh Supercomputing Center, Purdue University, San Diego Supercomputer Center, Texas Advanced Computing Center, and University of Chicago/Argonne National Laboratory”

In essence, the Teragrid is the supercomputer arm of the National Science Foundation, and while 30 Petabytes of online storage sounds “enormous”, split among hundreds of researchers and research universities it is actually quite modest. Teragrid stops at basic computing and bit storage for Teragrid calculations. Long-term data preservation is not explicitly funded. Data-bit archives or “backup” is funded only as part of a resource provider’s (RP) current 4-year award. At the end of the RP award, the data is orphaned or possibly lost.

SIDEBAR 2: Cloud Computing

Cloud computing is emerging as the next wave in computing models, combining aspects of the first four: mainframes, personal computers, client/server, and web applications. With cloud computing, storage and computing resources migrate into a set of managed infrastructure operated by third-party providers. Clients may instantiate computation on demand, paying only for the resources that they use at the times that they require them. Emerging technology allows clients to arbitrarily configure their machines with custom operating systems, applications, and even dedicated network resources. The promise of cloud computing can be summarized as delivering more flexibility, less cost, with a lower technical barrier to entry to large-scale computing.

Cloud computing has the potential to deliver significant operating efficiencies. For example, rather than requiring every research group with large-scale computational or storage needs to procure, install, and manage a complex infrastructure, they can "outsource" this cost and complexity to cloud service providers. Next, individual clusters often run at an average utilization of 3-10%. Statistical multiplexing of multiple workloads onto a single large physical infrastructure can overlap troughs in the workload from one application or customer with the peaks from another. Emerging virtualization technology from VMWare, Xen, Microsoft, and others, allow entire operating system/application machine environments to be multiplexed on physical machines and to be migrated on demand. Finally, centralized infrastructure can be operated more economically, both from a human perspective and from an energy and cooling perspective. Operating machines in a facility designed to support computing and storage can incur operating expenses that are a factor of 2-10 less than ad hoc deployments.

Today, cloud computing is relatively mature technology with stable product offerings from multiple companies, most notably Amazon, Google, and Microsoft. Virtually every large IT company has a product offering in this space. Today, one can configure a set of virtual machines on Amazon's infrastructure, using a variety of third-party tools to simplify the configuration process. However, existing environments are not geared toward the scientific computing requirements typical of universities. For instance, very large-scale storage and data archives are prohibitively expensive and it is currently not possible to access dedicated network resources among compute nodes. Further, the tools for configuring computing environments, instantiating and running applications, and finally managing and visualizing resulting output sets are still limited. Certainly, the state of the art is no worse with cloud computing relative to existing techniques, but there is opportunity to lower barriers to entry for many scientists.

Increasingly, the success of university research programs depends upon access to world class computation and storage infrastructure. Oftentimes, the research of a particular program is determined not by domain skills in the sciences but by the ability to procure, install, and manage large-scale computation infrastructure, a difficult, expensive, and time-consuming task. There is an opportunity for universities in general and UC San Diego in particular to act as a force multiplier for virtually all research programs on their campuses. With cloud computing technology, there is the opportunity to free individual researchers to once again focus on their particular domain expertise all while reducing aggregate computing capital and operating expenses. As an added benefit, building a world-class computation infrastructure will form yet another pillar in the University's ability to continue and to accelerate attraction of world-class scholars.

5. Achieving Economies of Scale: Campus Colocation Facilities

Overview and Recommendation

A data center is a facility designed and constructed specifically for the purpose of housing computing systems and related components. A colocation facility (also referred to as a collocation facility, colo, or co-lo), is a data center intended to house computing systems belonging to multiple customers. Such a facility will typically provide the specialized infrastructure required to support potentially large numbers of such systems: high density, reliable, and redundant power distribution, high bandwidth internal and external networking connectivity, high capacity cooling, fire protection, physical security, and in some cases a dedicated operations staff.

As one component of a next-generation cyberinfrastructure for UC San Diego, it is recommended that the University establish and support one or more computing resource colocation facilities to maintain and enhance the research and instructional mission of the University.

Large-scale colocation facilities can, when properly designed and implemented, achieve significant economies of scale over small- to medium-sized server rooms in both capital and continuing expenditures. Smaller server rooms and similar facilities frequently represent marked inefficiencies in energy consumption, allocation of campus space, and ultimately staffing costs. In addition to providing the foundation for the cyberinfrastructure required in all disciplines of the 21st century academic endeavor, the costs to establish and operate campus-wide colocation facilities will be offset substantially by the savings ultimately realized from the reduction in the number of smaller facilities.

Colocation Facility Infrastructure

A typical colocation facility is understood to provide a standard infrastructure to support the specialized environmental needs of large collections of computing systems:

A. One or more physical environments specifically engineered to support all infrastructural needs of compute, storage, networking, and similar cyberinfrastructure components. It is possible to implement the full set of recommended colocation features in a relatively small space (on the order of 1,000 sq ft). However, the economies of scale realized in substantially larger facilities argue against smaller ones unless they are required for technical or logistical reasons, and can demonstrate no significant loss of operational efficiency. For UCSD, the lower bound on a "large" facility would be on the order of 5,000 sq ft.

B. One or more rooms optimized for large numbers of standard, 19", 42RU racks, potentially densely populated with compute, storage, and networking equipment. Such racks and their physical dimensions will imply specific floor loading characteristics and specialized seismic mitigation structures. Common, but not required, is a raised floor with a surface consisting of removable tiles, some of which are perforated to permit the flow of cold air to computer system intakes. In this configuration, the subfloor area is used as the supply plenum. The rack systems, enclosure, flooring, etc. should be designed as an integrated system to support maximum cooling efficiency (see below).

C. Efficient, high capacity electrical distribution system sufficient to support overall room loads of 300W/sq ft. Common supply voltages for computing systems are 120VAC and 208VAC, typically via 30A circuits; the facility's electrical systems should be designed to support next-generation power infrastructures where possible, e.g., 480VAC and possibly DC systems. The power distribution system should supply a minimum of two circuits per rack, though high density rack configurations may require support for higher current loads.

D. Efficient, high capacity cooling systems to remove the heat loads generated by all computing systems, as well as other ancillary devices (e.g., electrical switching, transformer, power distribution, uninterruptible power supplies, room lighting, and cooling and humidification equipment). Per the current recommendations of the American Society of Heating, Refrigeration, and Air-Conditioning Engineers (ASHRE) the cooling system should have sufficient capacity to maintain a computer inlet temperature range of 18 °C – 27 °C (64.4 °F – 80.6 °F), and a relative humidity range of 35% - 60%. By way of comparison, most campus buildings are designed to handle a cooling load of 24W/sq ft, an order of magnitude less than what is typically required for high density server installations. High efficiency delivery of chilled air and removal of waste heat requires the use of thermal aisle containment systems. Both cold aisle and hot aisle containment have been shown to result in significant energy savings. In a recent study, Yahoo converted 8,000 sq ft of a 40,000 sq ft data center to use cold aisle containment. This mechanism, which also permitted increased server inlet temperatures, reduced energy usage by 21%, with estimated annual savings of $563K.

E. High bandwidth, low latency network connectivity within the colocation facility, to other systems on campus, and to research and production networks off campus. Increasingly, the services housed in colocation facilities support activities that are both compute and data intensive, and which routinely generate, store, and transmit hundreds of gigabytes and in some cases terabytes of data.

F. Electrical failover and redundancy systems. The extent to which such systems are implemented can vary significantly, but they often include uninterruptible power systems, backup generators, and possibly N+1 service replication.

G. Fire protection systems. There are several common fire suppression systems for colocation facilities. These include both water-based (e.g., wet pipe, dry pipe, pre-action) and chemical-based (e.g., FM200) systems.

H. A variety of physical security systems, including unique, time-based door codes, biometric readers, mantraps, video surveillance, alarm systems, and on-site personnel.

I. Supervisory, planning, facilities, and operations personnel. For both security and operational reasons, typical commercial colocation facilities are often supported 24/7/365 by personnel who design, build, maintain, and monitor all of the infrastructural components of the facility. In some cases, these personnel will also monitor the hardware and software state of the hosted computing systems. In academia, there is somewhat larger variation in the range and availability of support services associated with a colocation facility; while some facilities have staff levels akin to their commercial counterparts, at the opposite end of the spectrum, "dark" facilities are also not uncommon (i.e., minimal overall staffing, including no on-site personnel for some part of the day).

Motivation

The range of academic disciplines whose research is compute and data intensive has increased dramatically. In addition to engineering and the physical sciences, it is not uncommon for faculty in the life sciences, social sciences, and the arts and humanities to make use of significant computational and storage resources in support of their research. The naïve approach to meeting this resource need has been simply to add capacity when and where available, building out small and medium-sized server rooms in existing space, including those cases in which the term "server room" is generous at best. Even in those cases where these facilities are adequate as local solutions, they are far from optimal in the context of the University's capital and operating expenditures, and in the efficient allocation of space on campus.

Economies of Scale – Capital Costs

The costs to properly provision a legitimate server room are significant regardless of the size, but the overhead to construct a large number of small facilities is much higher than a small number of large facilities. Consider for example, the premise that, by virtue of value engineering, most campus buildings are designed and built with physical plant services sized for typical human occupancy over the life of the structure. In most cases, this does not include the potentially substantial electrical and cooling loads introduced by a server room facility. If we assume a commonly accepted electrical load value for outlets and lighting of approximately 6W/sq ft, a 100,000 sq ft building would be sized to support 600kW. If 1,000 sq ft of this space is used instead for a densely populated server room of 40 standard racks, the power requirements would increase to approximately 980kW, or by 63%, assuming currently available 1RMU servers running typical application loads. Similarly, assuming typical office cooling capacity of approximately 24W/sq ft, the cooling demands for the building would increase from 2,400kW to 2,761kW, or by 15%.

The theoretical power and cooling demands for 40 racks of computing equipment would be the same regardless of whether the racks are located in a small server room or a larger colocation facility. However, the lost economy of scale in the construction of many small server rooms versus a handful of large, dedicated facilities would be significant in terms of the ancillary infrastructure to support a more highly distributed aggregate rack count. The power and cooling needs of a server room are very different from those of office and most lab space, which leads to an increase in the complexity of in-building infrastructure. Further, the intra-building infrastructure required to support small server rooms distributed widely across campus is less efficient than the dedicated, high-capacity infrastructure to provision a very small number of large facilities. *Continuing to retrofit existing buildings and construct new ones with small server rooms results in significant capital costs that could be reduced via an efficient, large-scale support infrastructure within and among buildings.*

Economies of Scale – Operating Costs

The net decrease in capital costs described above assumed that operating costs are invariant as a function of size, but in reality, economies of scale provide an even stronger argument for the operational efficiency of large server rooms versus smaller ones, specifically energy efficiency. The design of any server room, data center, or colocation facility has one of the most significant impacts on overall energy efficiency. Two facilities housing essentially the same set of equipment can have very different energy bills based on how they are configured. While there is merit in assessing individual component efficiency, significant reductions in overall energy consumption are gained only via a holistic approach that considers the interaction of the components that make up the entire facility (rack placement, airflow dynamics, power distribution, etc.).

A comprehensive understanding of a data center's operating characteristics is possible regardless of the size of the facility, but the design process, underlying infrastructure, and ongoing updates are not only less efficient at a smaller scale, they are seldom implemented for want of technical expertise, and because very high efficiency power and cooling systems in smaller spaces are less compelling, not practical, or simply not available. The components of a high efficiency infrastructure have a larger upfront cost, but large-scale colocation facilities permit this cost to be amortized over tens of thousands of systems versus hundreds. A recent series of reports demonstrate how Google's large- scale data centers have been able to achieve "state-of-the-art efficiency by optimizing IT hardware and data center infrastructure end-to-end". In the case of Google and other service providers, data center efficiency efforts have been successful because a comprehensive approach is more effectively applied to larger facilities. Admittedly, UCSD will never have data centers comparable in size to Google, IBM, Microsoft, Yahoo, etc., but the finding that improved systemic efficiencies are best realized on a larger scale is nevertheless applicable to the University when comparing the relative merits of small, medium, and large data center operations.

As a component of the University's overall "green campus" initiative, the importance of maximizing the energy efficiency whereby UCSD houses its sizeable and growing computing facilities cannot be underestimated. Based on the early 2009 campus rate for electricity, a one megawatt power load costs approximately $745K/year; campus server systems already require multiple megawatts of electricity to satisfy their power and cooling needs, and a significant portion of this load results from inefficiencies that could be eliminated were the systems housed in a large-scale facility. Details of potential savings are discussed below.

Predicted Efficiencies at Scale

In a 2007 report to Congress on server and data center energy efficiency, the EPA presented an analysis of three energy reduction scenarios:

- Improved Operation – energy-efficiency improvements beyond current trends that are essentially operational in nature and require little or no capital investment. Represents the "low-hanging fruit" that can be harvested simply by operating the existing capital stock more efficiently. **Predicted 30% improvement in infrastructure energy efficiency.**

- Best Practice – efficiency gains that can be obtained through the more widespread adoption of the practices and technologies used in the most energy-efficient facilities in operation today. **Predicted 70% improvement in infrastructure energy efficiency.**

- State-of-the-Art: maximum energy-efficiency savings that could be achieved using available technologies. **Predicted 80% improvement in infrastructure energy efficiency.**

Two noteworthy features common to all three scenarios is that, regardless of the technologies used to improve energy efficiency, each scenario assumes that

- A significant base component of the predicted energy savings is the economy of scale to be gained by consolidation, both in terms of logical services aggregation and physical server aggregation into large facilities.

- The most cost effective and in some cases only physical configuration to implement fully the recommendations is a large server room.

The EPA report and similar studies are predicated on a wide variety of technical and organizational improvements in operational efficiency. Implicit in these is an understanding that the optimum model for the implementation of the recommendations is one that encourages uniformity across an organization, an assumption that is best realized and controlled via a centralized authority, and ultimately centralized facilities. In other words, large-scale, common data centers provide a superior mechanism for the implementation of energy efficiency best practices, an overview of which appears in Appendix F.

Sample Predicted Savings at Scale

Current trends indicate that site infrastructure costs, i.e., the costs related to powering and cooling IT equipment, have exceeded the cost of the equipment itself. The projected increase in energy costs predict that these will be from two to more than five times equipment cost over a 3-year life span. Continuing with a model in which the efficiency to be gained in large-scale, centralized colocation facilities is not realized is untenable from the perspective of a campus-wide energy management policy.

Consider, for example, the actual cost savings to the University if existing systems in inefficient, small- and medium-sized server rooms were to be migrated to a central facility with technologies ranging from “Improved Operation” to “Best Practice”. Assume a realistic scenario in which the heat generated by a single, 1U server running a typical application load is approximately 330W. Further assume that there exist, conservatively, 10,000 such systems distributed across campus in offices, labs, server “closets”, and small- and medium-sized server rooms (note that compute clusters of many hundreds and even thousands of nodes are not uncommon).

In the course of a year, based on the efficiencies of the cooling systems used at UCSD, the power needed to remove the waste heat would be 19,285MWh. Based on the effective rate paid for electricity by the University in early 2009 ($85/MWh), this represents an annual cost of $1.65M. Using the EPA estimates above, the range of savings to the University in annual energy costs for cooling alone would be $492K to $1.25M. Based on a 7-year linear regression analysis of wholesale energy prices in California, in 5 years the annual cost savings would be approximately$1M to $2M.

Other Considerations

Based on the discussion above, it is not practical to continue the build-out of smaller server rooms, but there is also an argument for dedicated, large-scale facilities that goes beyond the obvious savings in capital and operating expenses. As demonstrated elsewhere in this document, an increasing number of academics, including those in areas not traditionally considered compute and data intensive, now view such resources as an integral component of their research activities. As such, the need to supply and support facilities to house these resources in an efficient and effective way should be a mainstream concern of the University as a means to recruit and retain faculty. While a researcher may have funding to purchase compute and storage resources, the funding agency may not support the infrastructural costs of housing the equipment. In some sense, a colocation facility becomes the next logical component of an infrastructure to support the research and instructional mission of the University.

Further, in light of current and projected budgetary difficulties facing UCSD, mechanisms that centralize services where appropriate will ultimately reduce provisioning inefficiencies at the individual, department, and unit levels. Removing these inefficiencies in favor of the economies of scale available in large, centralized colocation facilities will ultimately result in savings that can be used to cover the capital and operating expenditures of such facilities, a cost-benefit analysis that will continue below.

Additional Cost Considerations

Power and cooling are substantial components of the operational costs of a colocation facility, but there are other expenses as well. Principal among these are construction and staffing costs.

The costs of building a data center from the ground up vary widely, ranging from $400/sq ft for low density facilities to $3,000/sq ft for very high density facilities, the latter including the current estimate for the effective construction costs of a typical Google data center. This wide price variation is largely a function of the design goals for a facility, primarily the ability to meet power and cooling load specifications, which in turn are driven by computing system densities. A commonly accepted average for data center construction is $1,000/sq ft, though this number will certainly go up as power and cooling densities increase.

The costs to retrofit existing facilities vary in a similar fashion. An analysis of several recent campus projects to convert existing space into small data center facilities resulted in an average cost of $850/sq ft. Two of the more expensive build-outs were constructed at costs of approximately $1,000/sq ft and $1,190/sq ft. As a practical matter, there already exist several campus facilities capable of meeting initial colocation needs with minimal or no start up costs. The principal facility in this context is the data center at SDSC, consisting of two rooms with a total of 18,000 sq ft of space and a current power capacity of 4MW.

As with construction costs, the personnel expenditures associated with a colocation facility can also vary significantly, primarily as a function of the extent to which a facility provides 24/7/365 support. Based on proposed recharge rates, SDSC estimates the following annual personnel costs for a fully staffed facility:

Staff Title	FTE	Salaries & Benefits
Operators	6	$356,495
Group Lead	1	96,720
Facility Manager	1	99,200
Networking	1	90,098
Security	0.5	56,060
Admin Support	0.25	70,680
TOTAL		$769,254

Table 5.1: Personnel costs for a fully staffed facility

As noted earlier, in academia it is not uncommon to operate a colocation facility in "lights out" or "dark" mode, in which the staffing needs are significantly less than a 24/7/365 facility. Consider the following as an example of a minimally staffed facility:

Staff Title	FTE	Salaries & Benefits
Operators (PAI average)	2	$131,795
Facility Manager (CRM2 average)	1	124,000
TOTAL		$255,795

Table 5.2: Personnel costs for a minimally staffed facility

In practice, the two staffing scenarios above constitute a range of possibilities that will be driven by the needs of the user community for a colocation facility and budgetary realities.

Funding Models

Any of the funding models described later in this document (Section 11, "Sustaining the CI: Funding Models") would work as a means of providing ongoing support for a research computing colocation facility; but some are more highly recommended than others. Specifically, feedback from IT directors and user communities at UCSD, elsewhere in the University of California, and at other academic institutions show that successful funding models for colocation facilities are those which encourage use by covering costs in a way that is transparent to the end user. Unless strictures are implemented to force acceptance, potential customers of a colocation facility will use it only to the extent to which it is perceived as "free", is at least as cost effective as any current hosting arrangement (including unit or departmental facilities), provides significant and desired added value, or if no other equivalent facility exists. This does not mean that the facility cannot have an underlying cost recovery model such as a recharge mechanism, but only that it is necessary to control end user exposure to such a mechanism in order to achieve broad acceptance and utilization.

Governance

Two principal components of a governance model for UCSD research computing colocation facilities are (1) the rules for oversight and management of the facility itself, and (2) one or more service level agreements that establish the rights and responsibilities of the facilities as service providers and the research community as users of facilities services. The rules for oversight and management ensure that the ongoing services provided by a colocation facility remain aligned with the needs of the user community it serves. Typically, the implementation of this process would include the establishment of a board or committee with representation from the management of a facility as well as from all relevant constituent user communities. This oversight group would ultimately be vested with sufficient power to direct the current operation and future plans of a colocation facility.

Service level agreements would record a common understanding among the service provider and customers and include a definition of services to be delivered, performance measurement, tracking and reporting mechanisms, problem management, legal compliance and resolution of disputes, customer duties and responsibilities, security and warranties, disaster recovery, and rules for agreement termination.

Examples of these two governance mechanisms can be found in Appendices G and H.

Implementation Plan

A natural starting point for a UCSD research computing colocation facility already exists in the resources available at the San Diego Supercomputer Center. Indeed, SDSC is actively engaged in a planning process to provide colocation facilities to other UC campuses, with a provision that this also be available to customers at UCSD. There are a number of details that remain to be resolved, including how to fund the UCSD component of SDSC's colocation facilities, as well as devising a model to encourage campus-wide use, or conversely to discourage continued reliance on small- and medium-size "server room" facilities. In addition, should needs warrant, SDSC should not be considered the sole colocation resource on campus. It may be necessary to use the SDSC model in conjunction with the recommendations in this section as a template for the implementation of additional, perhaps regional, colocation facilities elsewhere on campus.

6. Managing the Data Tsunami: Centralized Storage

Background and Recommendation

Reliable, high-performance, professionally managed raw data storage is a fundamental building block of UCSD research cyberinfrastructure. Simply stated, while users of computing (e.g., via their own clusters, through a campus condo cluster, an NSF computing center, or commercial clouds) can tolerate periods of downtime or limited availability, raw storage must be highly reliable and persistently available.

The Design team recommends that the campus invests in a centrally-managed raw storage facility (called UCSD Storage), the core hardware (disks, computing, networking, power, cooling) and a small dedicated, professional team to manage, update, monitor, and run the facility. Storage hardware must be renewed on a regular basis, making this a critical and ongoing commitment.

While exact governance of storage access and allocation needs significant additional effort to produce a tractable plan, the team suggests a two-pronged approach where a basic level of storage capacity (e.g., 1TB/researcher) would be provided to all campus researchers and additional storage could be purchased by those who have larger storage needs. UCSD storage would be made available through the campus research network (see Section 9) and would be accessible from specific workstations, departmental clusters, and the centralized computing facilities described elsewhere in this document. UCSD Storage forms the basic data “bit-warehouse” for higher-level services such as data preservation and stewardship (see Section 7), and would be sited in the SDSC collocation facility (see Section 5).

Our recommendation is for the campus to invest approximately $2M per year roughly equally divided between storage hardware and the people to run the UCSD Storage facility. Our estimates are that, if started today, approximately 1PB of highly-available, replicated storage (approximately 2PB raw space) could initiate the investment. Historical technology trend curves in storage capacity predict that with a recurring investment, storage space would double approximately every 2 years, including accounting for the retirement of older, less efficient, and beyond-warranty hardware.

The recurring yearly investment recognizes that raw data must be made available for many years (decades) and that the University itself must play a leading and significant role in providing access to historical digital data.

Technical Details and Support

We operate from the premise that valuable data must have at least two copies in geographically distinct locations to be considered reasonably secure and three copies forms the basis of data that is highly resistant to loss. A significant fraction of personnel resources will be spent in monitoring the integrity of stored data to look for and counteract malicious intent, and to guard against “bit rot” because data is stored on magnetic media.

At the most fundamental level, UCSD storage should provide access to the following types of clients:

- authenticated workstations and laptops on the UCSD network
- laboratory- and department-owned clusters
- primary/secondary storage for data-intensive instruments
- higher-level data preservation services (see Sec. 7)
- very high-performance shared resource facilities such as the Triton Resource (see Sidebar 1)

UCSD storage should form the data nucleus for campus so that standard research data workflows can be fundamentally enabled. For example, data from a short-read gene sequencing instrument needs to be stored in its raw form, then analyzed computationally (often using a cluster) and visualized on a workstation or high-end display wall, and finally "marked-up" by a data preservation system for future re-use. We believe that a properly designed storage system can provide these appropriate "access modes" of the raw data without requiring a different physical technology for each type of client.

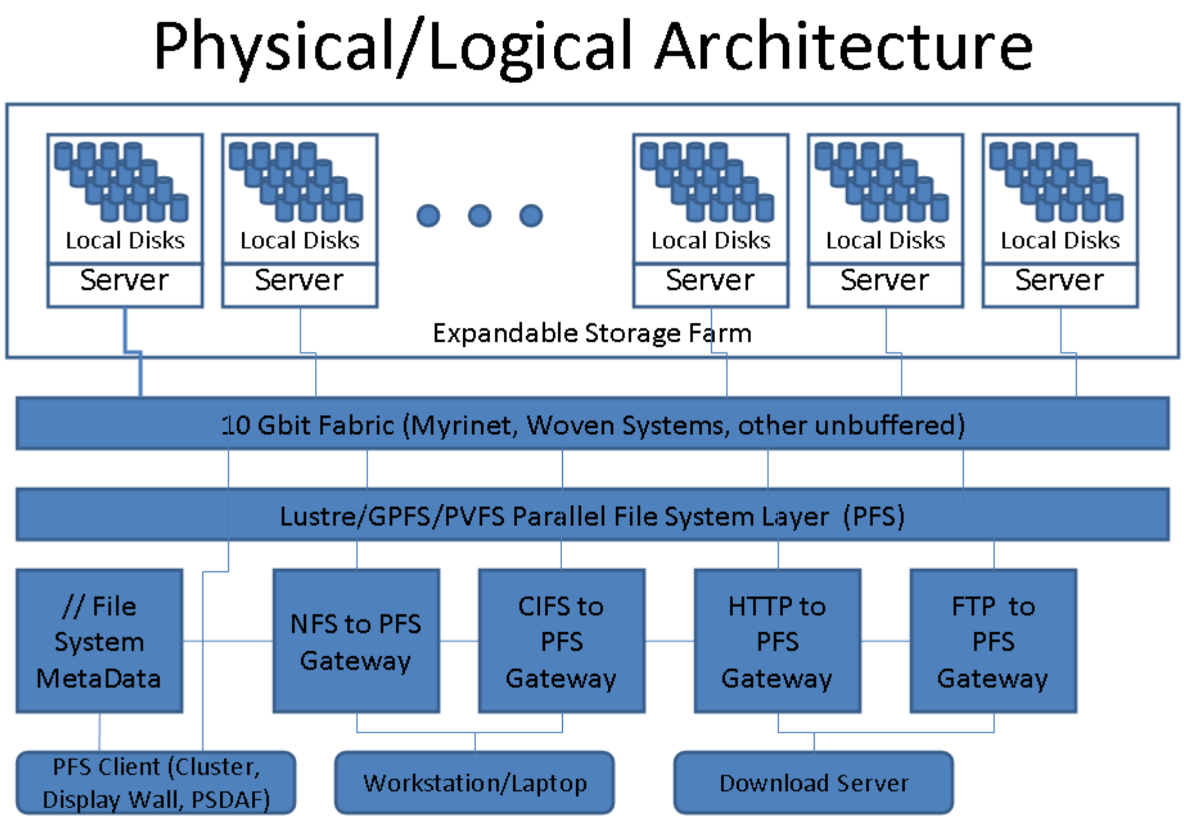

Figure 6.1: Physical Layout of Storage. A Parallel File System (PFS) forms the basic storage technology. Service intermediaries or "gateways" enable clients to access raw data using existing software

To meet the above needs, a parallel file system (PFS) is required to meet the storage performance requirements for data analysis using clusters and higher-end resources. We would propose that all UCSD data be stored in a parallel file system, and the "service intermediaries" or "gateways" (see Fig. 6.1) export the parallel file systems in other suitable network formats. For example, workstations and laptops would use their existing built-in clients (e.g., Windows File Sharing, Unix-based NFS, Mac file sharing) in conjunction with the appropriate UCSD Storage-operated intermediary to provide secure and authenticated data access. Clusters and higher-end resources would need to use native PFS client software to achieve acceptable performance.

Current technologies for the PFS include Lustre, PVFS2, GFS, GPFS, and Panasas. The first three are open source, while the latter two are commercial offerings. The physical hardware will likely support several "instances" of a PFS to differentiate elements of data integrity (e.g., single copy scratch space vs. two-copy secure data vs. three-copy loss-resistant data). This technical approach for multiple access methods to a PFS has been achieved at Harvard's Arts and Science's central IT facility and Tokyo Institute of Technology's TSUBAME cluster.

Data must be replicated to be physically secure. Single copy data (even if on physically reliable hardware) has an unacceptably high risk for loss. Figure 6.2 illustrates the physical layout of two storage sites on campus and gives possible locations for these sites. While this diagram shows two UCSD sites, it is highly recommended that a remote partner site be sought for true geographical distribution of data. This is consistent with the recommendation in the Data Preservation section.

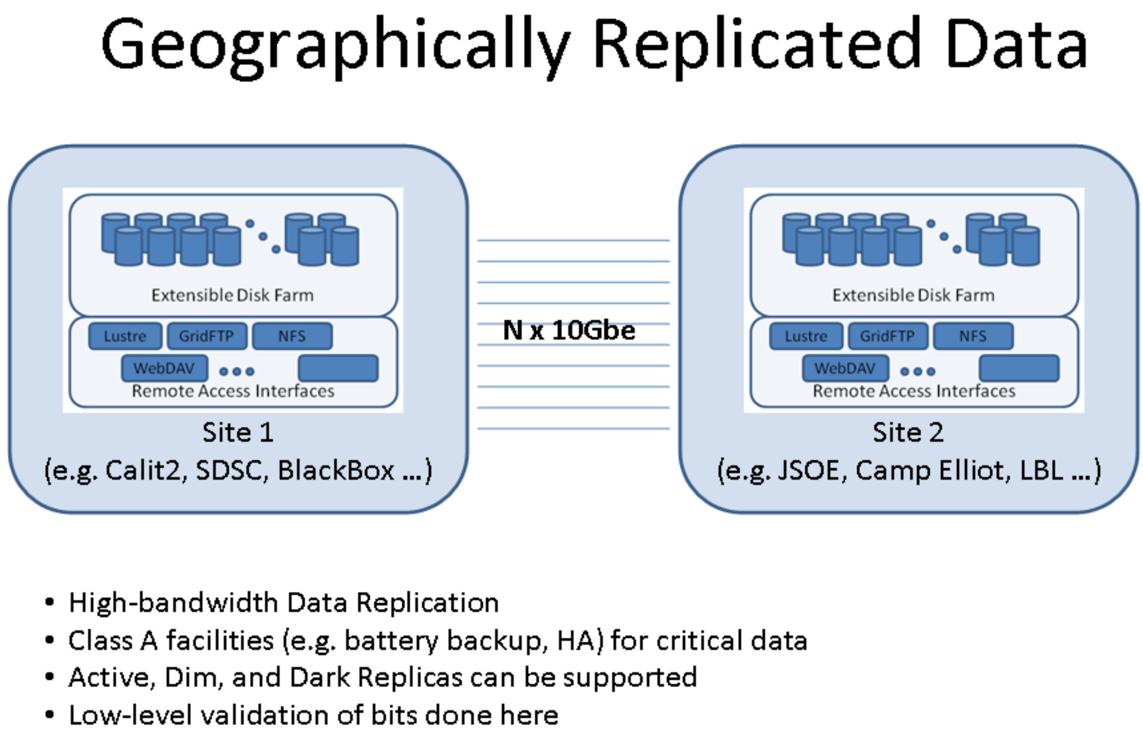

Figure 6.2: Replicated Storage for Integrity. A minimum of two copies of data is needed for data reliability. Geographic replication as shown here improves reliability. For even higher levels of data integrity, a third "offline" copy (e.g., via tape or optical media) should be defined. At the basic level proposed here, Third copy/Tape storage is not included in the plan.

Personnel Costs

It is critical to the ongoing storage commitment to provide appropriate levels of personnel funding to match the physical storage. The Design Team has identified the following critical activities and estimated staffing levels that are required to provide ongoing 7x24 storage service:

- Parallel File System Engineers/Administrators (2 FTEs)
    - Responsible for design and implementation of the Parallel File Systems needed to meet the needs of campus researchers

- Service Intermediary/Gateway Specialists (2 FTEs)
    - Responsible for design, debugging and deep technical assistance for providing access to storage

- Accounting, Monitoring, Security Specialists (2 FTEs)
    - Storage allocations will be given based upon the policies set forth in the storage governance. Both the policies and how data are being accessed must be actively monitored.

- Help-Desk, Website, Operator Policy Procedures (2 FTEs)
    - UCSD users must be able to find a person for assistance, as well as self-help through a website.
    - This group also responsible for writing operator procedures for non-business hour problem resolution

- Project Oversight, interface with governance committee (1 FTE)

The above represents minimum levels to provide a robust service. The critical areas require at least 2 FTEs to allow for vacation time of staff while still providing 7x24 Service. In addition, we identify other services that must be provided, but are not wholly unique to UCSD Storage:

- 7x24 Machine Room Operators
    - This is expected as part of a co-location recharge. UCSD Storage must provide the appropriate policies for off-hour operator procedures

- Network Design and access via UCSD Research and/or UCSD Production Network
    - Network access is essential to make UCSD Storage usable. Cost for providing physical connectivity of UCSD Storage to the campus and/or research network would be appropriate for a fraction of the Storage Hardware budget. Ongoing network engineering costs or laboratory connection costs are expected to be provided elsewhere.

We emphasize that the above personnel allocations are good *estimates*. The actual governance of storage may require additional personnel to implement policy.

Start-up costs

In this section we outline estimated start-up costs to initiate UCSD Storage:

- <u>Core Infrastructure</u>: Raw disk Infrastructure at approximately $500/TB of raw disk. Total Initial storage of approximately 2PB Raw or $1M.

- <u>Access/Parallel file System servers</u>: Servers that provide access to basic storage from clients. Initially approximately $80K for parallel file systems metadata servers and CIFS gateway servers.

- Networking at Core Sites: Site1 and Site2 would need significant network access via research network(s) and/or production network. Estimate at the beginning at least 8x10GbE per Site.

- Remote Cluster Access: Remote clusters using core storage as high-performance data storage will need access via the Research network.

- Multiple 10Gb research links: cost is highly variable and will be determined on a case-by-case basis.

We believe that some of the start-up costs can be mitigated with the storage component of the Triton Resource. For example, Triton Resource storage is approximately of this scale, will run a parallel file system for high-performance access and will be made available to some research labs in addition to the Triton large-memory and Shared Resource Clusters.

The fundamental difference between Triton storage and UCSD storage would be policy. Today, there are no funds for Triton storage to be considered over the long term or for growing. In the absence of such funds, data stored in Triton runs the risk of being lost when the warranty expires, when funding for data center space/electricity becomes unavailable and/or when no funds for personnel are identified on a continuing basis. This underscores that the long-term commitment to data requires ongoing core funding to ensure data integrity and availability.

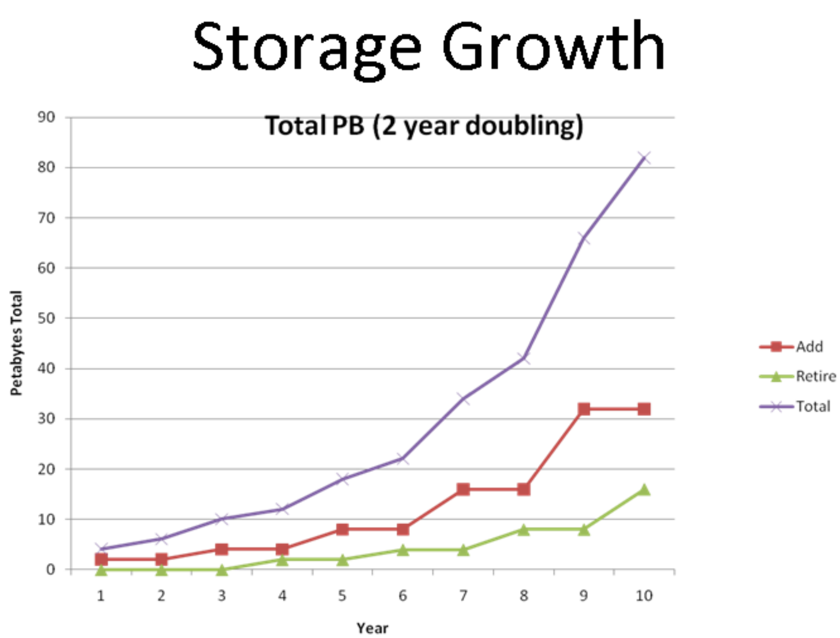

- **Intialize at 4 PB add $X. Spend 1/3X each year for new storage.**
- **Retire 3 year old hardware. Expect doubling rate of storage/$$ at 2 years.**

Figure 6.3: Storage growth over 10 years with 4PB of Raw Storage at Initialization. Assumes a recurring investment of storage hardware and retiring out-of-warranty hardware. Capacity growth in this model is 20X/decade

Scaling Storage – Storage Condo

It is clear that data storage requirements are not uniform across campus, researchers or labs. UCSD storage with stable base funding provides for the critical infrastructure, know-how, monitoring, and maintenance for a defined volume of storage that can grow over time. Figure 6.3 illustrates a conservative growth strategy. If for example, each faculty member were allocated 1 TB of long-term storage this year, at the end of 10 years, that number could grow to 20TB/faculty member. However, this does not answer the question of "How does a single researcher store 100TB today?" To solve this issue, we propose that UCSD storage could be operated as a storage "Condo." By that we mean the following:

Researchers could write extramural grants to fund the acquisition of physical storage to meet their needs and a fraction of administration, security, and monitoring that scales above the basic personnel costs outlined above. It would be part of the final governance to determine appropriate cost-recovery rates, but an advisory committee should consider at least two different cost scenarios:

- Condo storage (above basic allocation) has a lifetime limited to the warranty of the project-purchased storage.
- Condo storage has a lifetime equivalent to the core storage.

In other words, costing should be calculated on the basis of both limited lifetime and infinite lifetime for long-term data preservation of large-scale research data.

7. Institutional Stewardship of Research Data: The Research Data Depot

Background and Recommendation

Members of the UC San Diego research community routinely produce large amounts of data that need to be stored, analyzed, and preserved[1]. These research data sets and their derivative output (e.g., publications, visualizations, etc.) represent the intellectual capital of the University; they have inherent and enduring value and must be preserved and made readily accessible for reuse by future researchers. Today's interdisciplinary research challenges cannot be addressed without the ability to combine data from disparate disciplines. Researchers need to know: (1) what relevant data exist, (2) how to retrieve, (3) how to combine, and (4) how to mine and analyze them using the latest data-mining, analysis, and visualization tools. Granting agencies understand this fundamental scientific need, and are increasingly making it a condition of funding that researchers have a plan for preserving their data and for making it discoverable and available for reuse by other researchers. *If UC San Diego is to remain competitive, we need to invest in baseline data services that respond to these new realities.*

It is recommended that the University invest in the development and operation of an advanced data services facility, which we term the Research Data Depot. Implementation would involve a tight collaboration between the UCSD Libraries and SDSC, building on their respective expertise in digital curation, data management, and data cyberinfrastructure tools.

Research Data Depot

The proposed Research Data Depot is a suite of three core services designed to support the needs of modern researchers:

- Data Curation
- Data Discovery and Integration
- Data Analysis and Visualization

These services complement each other and provide a horizontal stack of data services that covers both active contemporary use and preservation for future use.

[1] As per CIDT survey: 70% of respondents expressed a need to process and manage large quantities of research data; 60% expressed a need for metadata creation resources; 55% expressed a need for interface/portal development; 62% expressed a need for database and/or data management support; and 58% expressed a need for visualization services.

Data Curation

Data curation encompasses the following three concepts:

- **Curation**: The activity of managing and promoting the use of data from their creation, to ensure they are fit for contemporary use and available for discovery and reuse. For dynamic data sets this may mean continuous updating or monitoring to keep them fit for future research. Higher levels of curation can involve maintaining links and annotations with published materials.

- **Archiving**: A *curation activity* that ensures that data are properly selected, appraised, stored, and made accessible. The logical and physical integrity of the data are maintained over time, including security and authenticity.

- **Preservation**: An *archiving activity* in which specific items or collections are maintained over time so that they can be accessed and remain viable in subsequent technology environments.

It is important to note that archiving and preservation are subsets of the larger curation process, which is a much broader, planned, and interactive process. It is also important to note that the curation process is well-understood and applied to non-digital scholarly materials by the UCSD Libraries.

Data curation is critically important for a research institution because it provides two vital services needed to ensure data longevity:

- Data are not merely stored, but are preserved to overcome the technical obsolescence inherent and unavoidable in any storage system.

- Data are documented in such a way that they can be linked in scientific publications and meet the requirements of funding agencies.

The curation service component of the Research Data Depot would be provided jointly by staff of the UCSD Libraries and SDSC. The UCSD Libraries would provide curatorial oversight and bibliographic control and integration services. SDSC staff would provide the back-end technology services needed to actively maintain the data and the storage systems holding them. Staff from both organizations will provide the metadata services necessary to ensure that data remain discoverable and accessible.

The data itself would be housed on campus, in the proposed UCSD Cyberinfrastructure campus storage facility. It should be noted that this is merely the first level of storage needed. For true long-term preservation, it is essential to plan for storage that is not on campus. If this is not done, data are always dependent on a single point of failure, and are thus highly vulnerable. Baseline investments are required to establish geographically distributed replicas of data. Since data are inextricably dependent on a mediating technological infrastructure and subject to loss occasioned by either environmental, organizational, or technological disruptions, it is imperative that vital campus research data be replicated in at least two remote sites—geographically, organizationally, and technically independent of each other—and that the entire enterprise be anchored within a reliable and predictable baseline source of revenue, as even a temporary interruption of proactive curation activities can lead to irreparable loss.

For this reason, another layer of service is required that stores *exact duplicates* of the data offsite. This important service would be modeled on Chronopolis, a ground-breaking project started by the UCSD Libraries and SDSC and initially funded by the Library of Congress. Specifically, Chronopolis is a joint partnership between SDSC, the UCSD Libraries, the National Center for Atmospheric Research (NCAR), and the University of Maryland (UMD). These sites have joined to provide the largest-scale preservation environment in the United States with 50 terabytes of federated storage available. The defining characteristic of this storage is that it is actively monitored, maintained, and managed at sites widely dispersed across the country on differing hardware platforms.

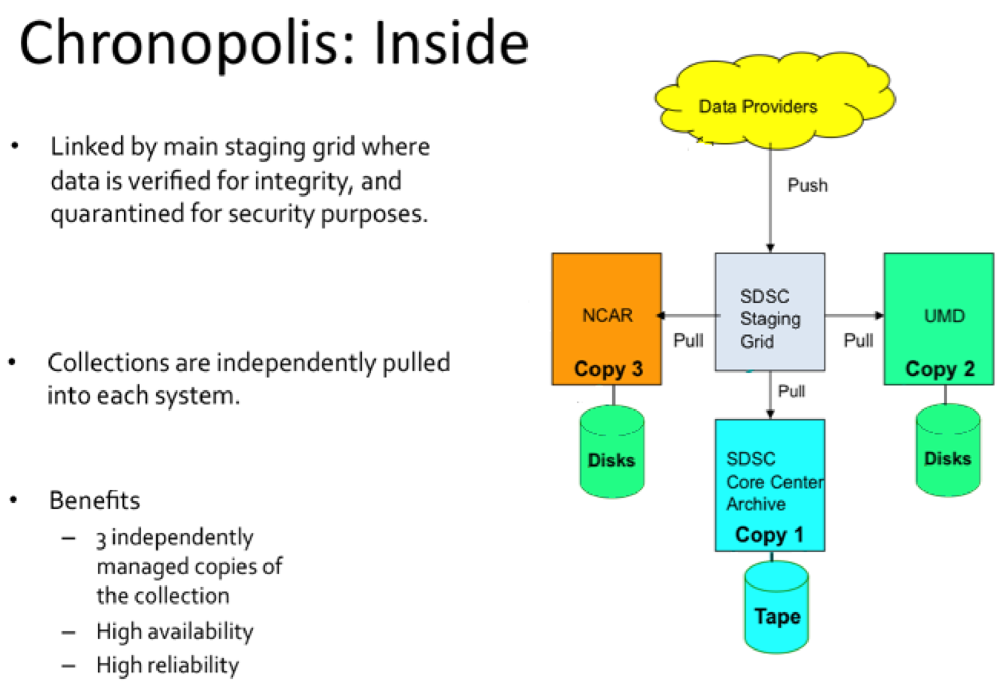

Figure 7.1: Chronopolis replication model

The UCSD-based curation services would be based on Chronopolis-style replication services using a minimum of 2 replicas, but preferably 3 replicas (the standard for long-term preservation).

The proposed curation service would leverage the investments made in the UC Shared Research Computing Services Pilot by depositing and linking replicas at the shared regional data centers.

Data Discovery and Integration

The Research Data Depot would provide a portal to facilitate the discovery of, and access to, the research data held in the Depot. This service would include facilities for the registration and description of collections, services to support the submission of collections, and assistance for the use, reuse, and amalgamation of data sets for further research.

The portal would assign persistent identifiers to each of the data collections, provide the ability to search across all collections that have been registered, link the data collections to their author(s) by leveraging the Calit2 Research Intelligence Portal, link the data collections to the resultant analyses and visualizations, and link the data collections to their published output through the integration of portal content with traditional library discovery tools and databases. Where appropriate, the contents of the Research Data Depot would be offered for harvesting and crawling by external discovery tools such as Google or disciplinary content aggregators.

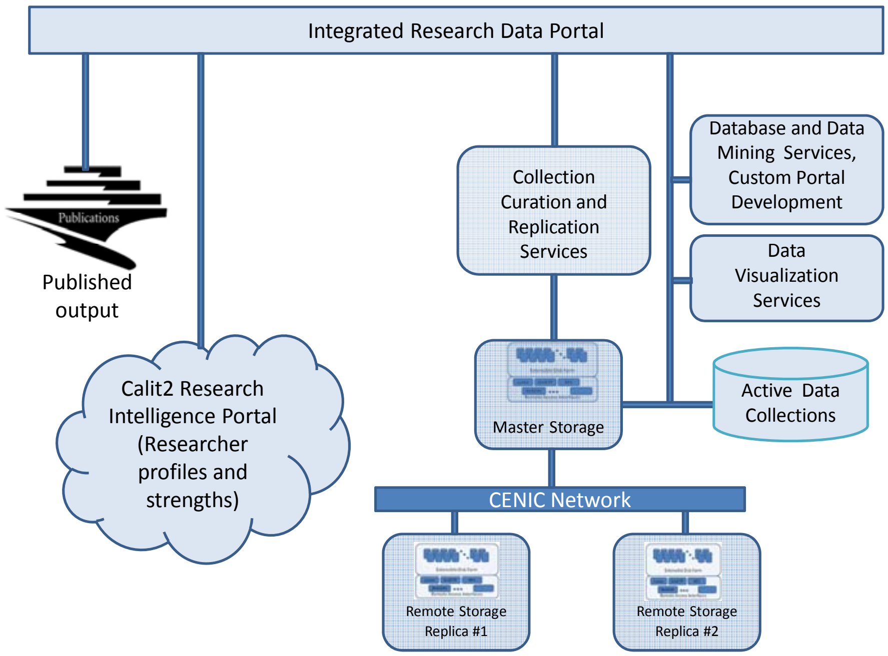

Figure 7.2: Research Data Depot service model

Data Analysis and Visualization

Baseline data analysis and visualization services would be provided to all UCSD researchers as one of the core services of the Research Data Depot.

Discovering the knowledge that is buried in large data collections is a team effort that includes the researchers who created or gathered the data, the staff who host the data, and the specialists who can analyze and visualize the data. There are several aspects to this work:

- Data migration, upload, metadata creation, and management: bringing data into active disk areas where they can be accessed, synthesized, and used.
- Interface creation: adding front-ends for either the data owners or their designated audiences to access and manipulate the data.
- Data analysis and mining: providing services that use advanced statistical and database processes to create usable data sets out of raw data.

- Database/management tools selection (Oracle, MySQL, SRB, etc.): helping data owners and users understand the options at their disposal and helping them choose the most appropriate tools for their needs.

- Distributed data management: working with data owners and researchers who have data scattered across different sources and in different locations, synthesizing it to form a more coherent working environment.

- Database application tuning, and database optimization: providing ongoing advanced database support for a myriad of activities.

- Schema design and SQL query tuning: helping with advanced data searching services for a wide variety of data.

These tasks are all necessarily *active* in nature, and involve researchers and service providers working directly with the data on a nearly continuous basis. Only by doing this can they provide users with the ability to organize, process, and manage large quantities of research data into data collections for data-driven discovery.

The visualization services at SDSC and CalIT2 provide users with a wide range of tools and services to aid in their scientific research. Their work with the UCSD Cancer Center, the Southern California Earthquake Center, and the Center for Astrophysics and Space Sciences, among others, has allowed these researchers to see their data in a new light, leading to innovations and discoveries in their respective fields.

Summary of services

The three (3) core sets of services will be composed of the following detailed set of services:

- Data Curation Services
    - Assistance in the development of data management plans
    - Data transfer/ingest support services
    - Metadata creation support services
    - Data preservation services
    - Data replication services
    - Long-term data integrity check services

- Data Discovery and Integration Services
    - Data set registration
    - Data set provenance and possible reuse description
    - Data submission process support
    - Data set availability alerting services
    - Ontology development
    - Portal development and maintenance
    - Linking of data sets to resulting publications
    - Linking of data sets to resulting analysis, mining results, and visualizations

- Data Analysis and Visualization Services
    - Data set transfer and transformation
    - Database creation, management, and optimization

- Data analysis support services
- Data mining support services
- Data visualization support services

While the 3 categories of services are somewhat independent of each other, UCSD researchers requesting Data Discovery, Integration, Data Analysis, or Data Visualization services will be strongly encouraged to use the Data Curation Service to ensure the long-term preservation of the data and associated analyses for discovery by future researchers.

What kinds of data will be eligible for data depot services

In principle, all and any data of importance to UCSD researchers, in any discipline, should be eligible for depot services. For financial, legal, and practical purposes, however, some level of appraisal and selection must take place. These decisions will be guided by the mission and goals of UCSD, intellectual property considerations and rights, and legal and contractual obligations. Beyond that, additional criteria for selection and appraisal might include:

- Are the data important or relevant to other constituencies in the discipline?
- Are the data significant, unique, and influential?
- What is the potential for re-use of the data in subsequent scholarship and research?
- Are the data complete and authentic?

Benefits

The benefits of the Depot to the UCSD research community are both quantitative and qualitative.

- **Reduced expenditures**, provided by economies of scale gained through central purchasing.

- **Increased researcher productivity**, provided by:
  - Availability of automated workflows for the ingest of data sets
  - Assistance in the provision of metadata

- **Improved data sharing and opportunities for interdisciplinary research**, provided by:
  - Improved discovery of data sets
  - Creation of interdisciplinary ontologies
  - Application of common standards and protocols to enable use of data by a wider range of user communities
  - Automated matching of researcher interests through the CalIT2 Research Intelligence Portal

- **Improved competitiveness** for grants, provided by availability of baseline data preservation and dissemination services.

- **Improved researcher recruitment and retention**, provided by availability of baseline cyberinfrastructure services.

- **Improved research activity assessment measures**, provided by monitoring of portal and data reuse activities.

Governance

Historically, libraries have served as the repository of the scholarly record. Over the past centuries, libraries have developed a solid record of trust in preserving that record and making it discoverable and accessible to researchers, students, and the general public. They have deep and well-established collection appraisal, accession, description, arrangement, storage, archiving, preservation, and access policies and practices to guide the process. These policies and practices apply as well to the data challenge as they do to the written record. Data centers such as SDSC have substantial experience in the creation and maintenance of massive data stores, as well as in the ingestion, analysis, and manipulation of large-scale science and engineering data sets. Combining the expertise of both of these organizations is a solid recipe for success. The UCSD Libraries and SDSC recognize each other's strengths and have developed collaborative partnerships that have proven highly successful in the arena of distributed digital preservation.

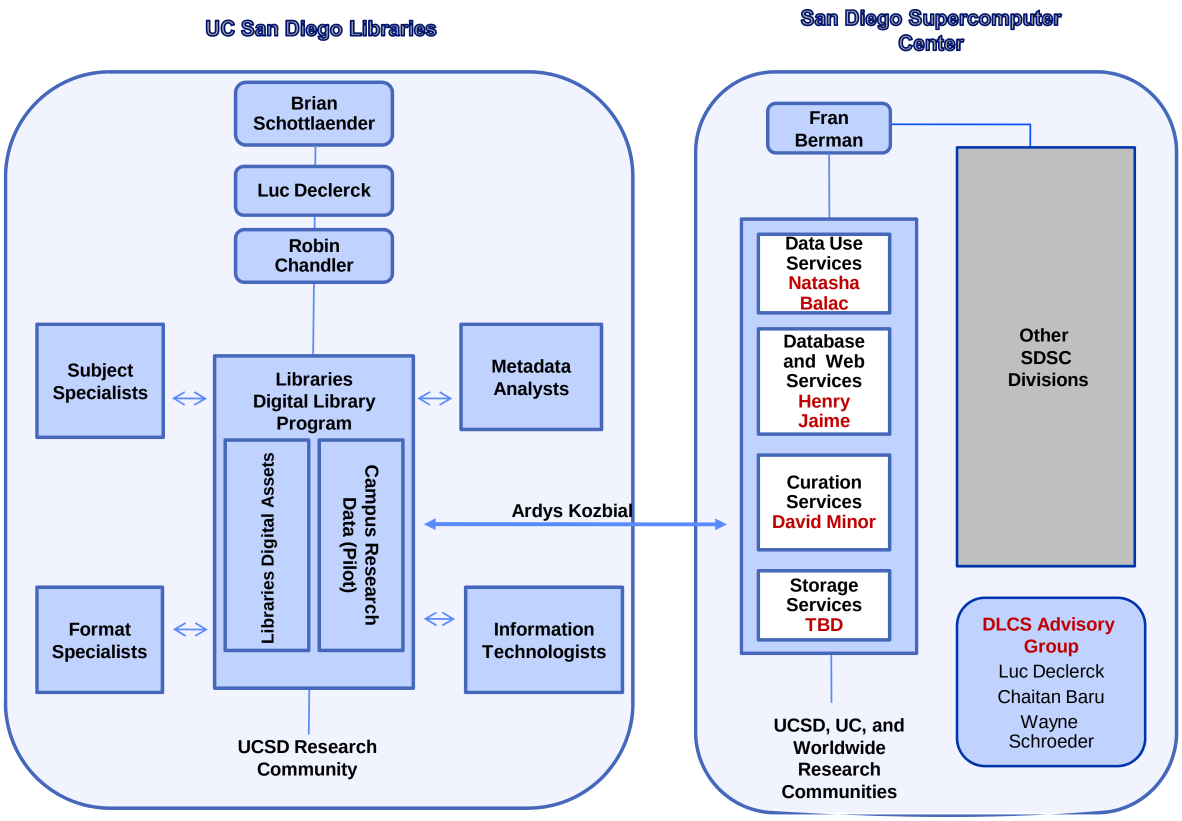

Figure 7.3: Libraries/SDSC collaboration organizational structure

This model would be used as a starting point for the management of the Research Data Depot and will be adjusted to cover activities beyond data preservation and to add other partners, including CalIT2.

A Steering Committee would be formed to prioritize the areas of focus and to make key decisions about all areas of the program, including products and services, partnerships, funding, and performance of delivered services. The Steering Committee will be composed of representatives from all key constituencies, as follows:

(a) A chair appointed by the VC Research
(b) One representative appointed by each of the Depot partners
(c) Two to three data providers and researchers
(d) Additional members such as data policy or other specialist representatives
(e) Research Data Depot Director, *ex officio*

Baseline costs:

The baseline costs associated with operating and maintaining the Research Data Depot include:

- **Personnel costs:** Depot partners can provide the foundation for the service and ensure that adequate management is provided, but require a personnel augmentation of a minimum of 7.0 FTE, as follows:

  1.0 FTE Data specialist
  1.0 FTE Metadata specialist
  3.0 FTE Systems/software specialists
  1.0 FTE Web Developer
  1.0 FTE Systems administrator

- **Data Storage costs:** Initial data storage needs, estimated at 500 TB (250TB dedicated to preservation and 250TB used for data staging, data analysis, and data visualization) will be provisioned by the proposed UCSD RCI centralized storage facility (Sec. 6 below).

- **Data replication costs:** Data replication costs for 250 TB at two remote sites are calculated based on actual costs incurred by the Chronopolis project and noted in section 6, below.

- **Software license and server costs:** Software license and server costs are estimated at $75k based on actual SDSC experience.

- **Administrative costs:** Administrative costs such as desktop computers, communications charges, and other overhead will be absorbed by the partners.

Budget request

The additional budget increment needed to seed the Depot is estimated to be as follows:

- One-time cost: $850k
  Funds to purchase 250 TB of storage space at each of two remote sites ($75K for servers + $350K for 250 TB of storage per site). Note that this cost could be reduced by leveraging the investments made in the regional data centers included in the Shared Research Computing Services Pilot.

- Annual personnel costs: $840k (including benefits)
  Covers the incremental 7.0 FTE needed to operate the Depot.

- Annual technology refresh costs: $283k Covers replacement of servers and data storage hardware at two remote sites on a 3-year replacement cycle, i.e., $850K divided by 3. (Cost of refreshing the local Campus Storage facility hardware is covered elsewhere)

Incremental personnel and data storage additions may be needed from time to time, depending on the rate of adoption of the service by the UCSD research community.

Cost Summary

	$	Comments
Start-up Costs		
Storage	$700,000	250 TB Storage for each of the 2 remote sites
Servers	$150,000	Server for each of the 2 remote sites
Subtotal	***$850,000***	
Annual Cost		
Labor	$840,000	(including benefits) for 7 FTEs
Technology Refresh	$283,000	Replacement of servers and storage on 3-year cycle
Subtotal	***$1,123,000***	

Ongoing sustainability

Beyond the baseline budget level commitment requested from the UCSD Administration, we intend for the Depot to secure long-term sustainability through a variety of other budget sources, including:

- <u>Grant funding</u>
  We will seek grant-funded opportunities to augment operating costs, conduct research, and hopefully sponsor the growth and development of new facilities and resources. For example, recent NSF proposals such as DataNet, and projects funded by the Library of Congress and the Mellon Foundation could all use and help fund the Depot.

- <u>Grant budget line revenues as required by VC</u>
  In an increasing number of cases, research grants are allocated with the expectation or requirement that research data sets and their associated papers and media will be made available to other researchers for additional scholarly activities. At the moment, these requirements are managed inconsistently across the campus. Facilities provided by the Depot would allow the Vice Chancellor for Research to build in standardized budget lines within each research grant request where the need to provide access to and preservation of data collections exists.

- <u>Partners</u>
  As stated above, costs could be reduced through reciprocal storage agreements with another UC campus or the Lawrence Berkeley Lab. We intend to pursue such opportunities and to seek other cost-sharing partnerships where feasible.

- Service revenues from external customers
  Once established, Depot services could be offered to external customers for a fee. Examples would be the proposed UC Grid, other UC campuses, campuses from the California State system, and other universities nationally.

8. Scaling Up Research Computing to Meet Needs: Condo Clusters

What is "Condo Clustering?"

Many UCSD research teams and departments regularly procure and deploy small- to medium-sized compute clusters to support their research needs. These are typically deployed as stand-alone resources in "campus computing closets" and are used by a small team of researchers. Many universities have found more efficient and strategic ways to deploy these resources on their campuses. One of the most interesting is the concept of "condo clustering". Condo clusters provide a means of unifying these compute resources by integrating several clusters into one or more larger clusters. Once configured this way they offer many benefits to the owners of the system as well as to a larger campus community.

Condo clusters offer several benefits to the owners of the system. They allow the owners to leverage a pool of system administrators to manage the systems and other economies of scale (e.g., procurements, software licenses, security, and networking) that come with integrating into a single, centrally administered cluster. Currently, significant resources (labor of graduate students and postdoctoral scholars and other money) are spent on disaggregated clusters that could be significantly optimized in a condo cluster setting. Another benefit of condo clusters includes the ability for a "faster-time-to-solution" in research given that a larger resource is available overall.

From a campus perspective condo clusters also offer several advantages. Foremost, they allow the campus to make use of the idle cycles on the clusters by opening these up to the broader campus community (e.g., students, researchers). Another benefit of condo clusters is that they are housed in a co-location facility and thereby reap the benefits of a "green datacenter" as described in the co-location section. This central location of a condo cluster also provides a strategic hub for cyberinfrastructure activities that could result in a more competitive posture for educational and research opportunities.

The CIDT recommends UCSD embrace the concept of condo clusters and exploit the deployment of the Triton Resource at SDSC (see Sidebar 1) to launch the initiative.

What are Other Campuses Doing with Condo Clusters?

Many leading research universities have deployed condo cluster capabilities on their campuses. An informal survey of the CASC (Coalition for Advanced Scientific Computing, www.casc.org) membership included the following campuses that have deployed condo cluster solutions: U Buffalo, SUNY; U South Florida, USC, Clemson, Indiana, Purdue, Virginia Tech, Utah State University, Penn State, and U Arkansas. Also a couple of UC campuses have either initiated condo cluster programs (UCLA) or are in the planning stages (UCR). One overall interesting observation is that all of the campuses have support from their VCR or Provost to cover all or most of the costs related to condo clustering. There was significant variation in how the overall programs were implemented (scheduling algorithms, procurement strategies, size of program); nevertheless, it is clear that this is an important asset to having a strong vibrant cyberinfrastructure community on a university campus.

Below are a couple of case studies from the CASC survey:

Case Study #1

Penn State University is a leader in condo clustering within the CASC community. They have a 3000+ core condo cluster that is integrated for use by the entire campus community. This program has been very well received by both the researchers procuring clusters ("owners") via this program and by the campus user community that has access to idle time on this system ("orphans"). The Penn State program has significant campus support at $2M/year (with 15 FTEs) with an increase in support planned to $5M/year.

The condo cluster capability is provided at no cost to the "owners" or the "orphans". Cluster owners purchase the hardware using their own start-up or research funds, which they then contribute to the condo. The campus contributes about 25% of the condo gear (e.g., interconnects, file systems, nodes) to the overall system to provide an even larger and more capable resource. Procurement is done centrally; however, "owners" maintain their ownership rights. A most interesting feature of the Penn State condo model is their scheduling/queuing policies for the system. In particular they use a "next to run for owners" and do not pre-empt "orphan" users. A major strength of the success of this program appears to be the real engagement of the leadership on the condo cluster management with the "owners" of the system to manage expectations and create an atmosphere of real partnership.

Case Study #2

Clemson University also has an active condo cluster program. They receive $1.5M/year from their campus plus extra help that is leveraged from their campus cyber-institute. Their condo system includes 30-40% individual ownership ("owners") and the rest of the system is purchased with central funds. This integrated resource is provided to "owners" and "orphans" at no cost. The Clemson representative said "recharge dooms the condo to fighting over nickels. A University has to believe it's strategic… and not a plumbing thing". This comment seemed to echo that of the broader CASC community that was surveyed. Note: Campus cluster users at Clemson that don't participate in the "condo program" are charged a co-location recharge. This encourages strong condo participation. As with Penn State, "owners" continue to own their gear at Clemson; however, scheduling is done differently. "Owners" have the right to run their jobs "right away" (jobs submitted by "orphans" are preempted). Overall there is strong support for this program by their users, who see this as a partnership and not just a customer-service relationship.

Recommendations for UCSD Condo Clustering

The following recommendations address two elements of the charge to the RCIDT

1. Outline possible operations and business models for condo types of cluster arrangements

2. Determine how best to meet the needs of the UCSD HPC community

The RCIDT recommends that UCSD deploy a strong condo cluster program. UCSD is in a unique position with an active cluster and HPC program currently on campus that could be leveraged to seed the beginnings of this new program. In particular it is recommended that:

1. **Getting Started:** A UCSD Condo cluster program should leverage the SDSC "Shared Research Cluster" (part of Triton that is being procured) and the recently approved "Southern UC3 cluster" to create an initial resource to being a condo cluster program.

2. **Oversight**: It is important that this program be kicked off so as to engage campus users. It is recommended that an advisory committee made up of owners and key campus "orphan" representatives be formed to develop a program based on a strong partnership and shared goals.

3. **Operational Recommendations:** The following operational recommendations are made based on the understanding of what other campuses are doing. These should be refined and reviewed in consultation with the advisory committee.

    a. Service Level: It is recommend that the condo cluster(s) be operated at a level that is more formal than a typical departmental cluster (e.g., strong systems support, networking, security) and yet not at the level at which national supercomputers are supported (modest user support, documentation, training, software).

    b. Procurement Cycle: Procure 1-2 times per year. Make a campus "call" so that campus owners can leverage costing opportunities and that the condo cluster not be unnecessarily heterogeneous. Develop a "value statement" for this program so that owners understand the benefits (e.g., access to larger system and more cycles, economies of scale, access to peripherals (filesystem, archive), system administration, backups, software licenses).

    c. Ownership: The "owners" should continue to maintain ownership of their gear. If for some reason they no longer want to participate in the condo program they can "take their gear back". This has never happened in the two case studies mentioned above but it was an important aspect in getting "owner" buy-in into the program.

    d. Scheduling: Scheduling should support pre-emption by "owners". "Orphans" should be taught how to make jobs re-startable should they be pre-empted.

4. **Business Model:** It is recommended that all "owners" participating in the program get services at no cost since they are providing the idle cycles of their system to the campus user community. No condo participants should be charged the co-location rate for their clusters.

5. **Budget**: By leveraging the current investments by SDSC, no initial procurement investment is needed by UCSD. However, increased operational support (while leveraging the current Triton and Southern cluster staff) will be needed to provide the following functions:

    a. Management oversight, coordination, procurement (0.5 FTE)

b. System administration, creation of user accounts (2 FTE)

c. Scheduling system and management (0.5 FTE)

d. Software Licenses ($30K)

e. User support, documentation and training (1.5 FTE)

At $120K/FTE this comes to $570K for annual operating labor and software expenses.Also, it would be expected that some contribution to refresh the system (e.g., 25% of the total) would be needed on a periodic basis (estimate $1M every three years – or $333K/year). These three categories total $1.2M/year to operate, run and refresh a condo cluster program at UCSD. It should be noted that the utility costs (est $170K of this total) would be incurred by campus irrespective of whether the condo program existed.

Metrics

To ensure overall success of this program, it is important to establish metrics that align with the program goals. Some metrics to consider include:

- How many “owners” participate in the program? A range of 3-6 would be considered successful.
- How much idle time is delivered to “orphans”? If overall system utilization is >80% and it is serving at least 50 campus “orphan” users with 25% of the resource, this would be considered successful.
- How happy are users? An annual user survey would be conducted to assess user satisfaction.
- What science impact is the system having? An annual report to document science advances supported by the system would be provided.
- Has the system provided increased competitiveness on proposals? How could this be measured?
- Has the system had an educational impact? How many students used the system?

9. Providing High Speed Access: The Research Cyberinfrastructure Network

Recommendation

UC San Diego should provide all campus researchers with a leading-edge network that meets their needs and facilitates collaboration, high-performance data exchange, access to co-location facilities, remote mounts to storage, and real-time communications.

To that end, we recommend a Research Cyberinfrastructure Network (RCN) that will be used by every research lab whose requirements go beyond the standard production network. The RCN will complement the standard NGN production network and will be designed for ultra-high performance. It should be the first campus environment for implementation of newer technologies before being adopted into the standard production network. Funding and access philosophy should aim to encourage usage of the network.

Brief description including support needs:

- Uncongested high-speed access to researchers
- High-speed access to central storage and computing
- Buildings and labs connected as research needs arise
- Infrastructure is ready to flex as needs expand
- Reasonable security controls

Start-up costs (one-time)

Semi-dedicated lab connections are used for these cost projections — each lab will have its own 10g switch and a dedicated 10g connection to the building switch. From there, the lab would share the building's 10g research connection. If needed, a lab could also opt for a dedicated-to-the-core connection, but that configuration will carry a higher price tag. Buildings included in this estimate are:

- Applied Physics & Mathematics
- Atkinson Hall (CalIT2)
- Blackbox #1: CalIT2/Greenlight
- Blackbox #2: Medicine
- Bonner Hall
- Cancer Center
- Center for Molecular Genetics
- CMM/E
- Computer Science and Engineering: IIIb
- Dunhill Street
- Engineering Building I
- Engineering Building II
- George Palade Laboratories (CMM/W)
- Holly/SOM Modular Building Complex
- Hubbs Hall
- IGPP
- La Jolla Professional Building
- Leichtag Building
- Mayer Hall
- Muir Biology Building
- Music (new building)
- Pharmaceutical Sciences
- Powell-Focht Bioengineering Hall: IIIa
- Powell/SERF
- Visual Arts

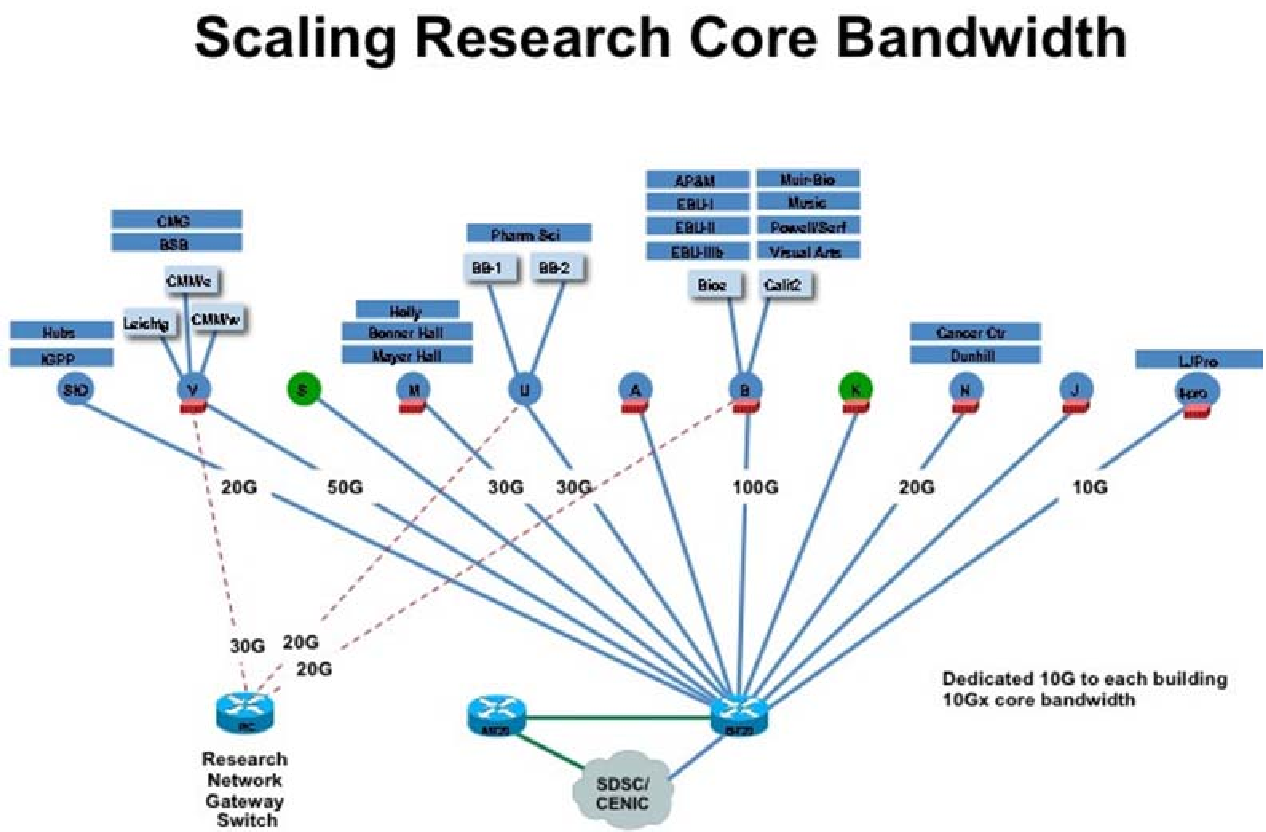

Figure 9.1: Proposed expansion of the Research Cyberinfrastructure Network

- $500,000 - $700,000 for core Infrastructure. This includes core network components, including items such as fiber links, research distribution switch, backbone WDM gear, and central support staff.
- $1,000,000 - $1,500,000 to connect the 25 listed buildings, depending on whether we prepare them for one or four labs per buildings.
- $20,500 per lab for semi-dedicated or $37,700 for dedicated connections (no incremental cost for labs that will simply share the upgraded building bandwidth) plus $1,000/year maintenance

Notes:

1) Model assumes 25 buildings in the RCN and the addition of labs ratably over 4 years.
2) Initial staffing: one Senior Network Engineer, one Senior Security Engineer, and one Research Engineer.
3) Security will be placed at the distribution nodes only.
4) These costs do not include any additional off-campus network connections such as those required for the UC co-location service.

Continuous costs (refresh cycle, scaling, staffing)

- $550,000 - $600,000/year in annual operational costs
- Electronic devices need to be replaced every 3-5 years.
- Optronics need to be replaced every 4-6 years.

Funding model

- Negotiate an increase to the campus overhead rate to support these costs.
- Increase the NGN rates, possibly differentiating between researchers and the rest of the community.
- Ask the Chancellor and the academic Vice chancellors to subsidize these costs.
- Charge the researchers for their direct costs.
- Combination of any or all of the above.

How do we get there?

- Identify specific labs that need to be connected to the research network.
- Prioritize connections.
- Develop cost projections depending on estimated connections.
- Determine a funding model with VC's, Academic Senate, NGN committee, Budget Office, UC, etc.

Deployment phases and proposed approach

The proposed model deploys high-speed infrastructure — fiber, switching, optics etc. — to an upgraded network core and to strategic buildings (see proposed list above) with research needs. The network core upgrade has already been planned and must be completed prior to RCN rollout. Required fiber and switches will be in place, offering researchers in these buildings immediate increases in bandwidth and all the benefits of the production network services. The minimal additional equipment necessary for fully reserved 10g lines will be kept in stock so that such a connection can be quickly deployed on demand to any of the 25 identified buildings. The *diagram* on the last page of this document shows the substantial research bandwidth available under this model.

Under this model, researchers who need a reserved-bandwidth, private connection can have their requirements satisfied promptly, without time-consuming security discussions and with minimal configuration. Alternately, researchers whose requirements are high bandwidth and access to the campus environment can have all the services they expect without providing laborious specifications. Finally, in-between needs can be accommodated swiftly, allowing a researcher to request the type of connection that exactly serves his/her needs and budget with a fast turnaround time. This model will also allow use of dynamic circuit switching when appropriate, and will allow us to incorporate that into the production network as the technology matures.

The security model in this network allows for sharing the security infrastructure of the production environment, which is closer to the protected machine. This allows us to provide similar security features and protections to most researchers without adding impediments and without installing expensive security equipment at the network center. At the same time, we can remove constraints from those who want dedicated private pipes.

Existing Pilot Phase

The first pilot phase of the research network was jointly funded by ACT, the School of Medicine, and CalIT2. The goal of the pilot is to test the network design and develop plans for the future. Buildings included in the pilot phase are: Leichtag, CMM/W, Atkinson Hall, and the Sun Blackboxes.

Implementation

Deployment process and schedule will depend on identified needs, funding model, and available funds.

10. Implementation Plan

Table 10.1 presents, in bullet form, a 3-year implementation plan beginning in FY10 for the 5 core elements of the proposed RCI.

RCI element	FY10	FY11	FY12
Colocation facility	UCSD leases 15 racks for "free hosting" at SDSC colo UCSD allocation committee formed; Chair appointed to UC oversight board Free hosting service announced UC service level agreement adopted	UCSD leases 15 additional racks for "free hosting" at SDSC colo Appraisal and capacity planning by allocation committee	UCSD leases 15 additional racks for "free hosting" at SDSC colo Appraisal and capacity planning by allocation committee
UCSD storage	UCSD augments Data Oasis with 2 PB to launch UCSD storage Faculty terabytes program announced Governance committee formed SRB/IRODS gateway to UCSD storage	Storage condo program announced Additional gateways (NFS, CIFS) added as needed Direct mount of parallel file systems on lab clusters over RC	Ongoing scale-up and replacement of UCSD storage
FTE UCSD storage	2.5	7	9
Research Data Depot	Create Data Depot core instance on UCSD storage Select replication sites Launch 3 pilot data collections Data Depot steering committee formed Initial staffing hired	Activate first replication site Add data collections Add staff	Activate second replication site Add data collections Add staff

FTE Data Depot	3	5	7
Condo Cluster	Leverage Triton Shared Resource Cluster and UC3 South to launch condo cluster program Form campus advisory committee 1st campus "call" for procurement	Scheduler improvements 1-2 campus "calls" for procurement Assessment and policy refinement	1-2 campus "calls" for procurement Impact analysis
FTE Condo Cluster	3.5	4	4.5
Networking	Upgrade network core Identify labs needing RCN connections Prioritize building connections Develop cost projections Determine funding model Connect additional buildings/labs to RCN	Review/update priorities list Implement funding model Connect additional buildings/labs to RCN	Review/update priorities list Connect additional buildings/labs to RCN
FTE Networking	3	3	3
FTE TOTAL	**12**	**17**	**19.5**

Table 10.1: A 3-year plan for implementing the proposed RCI.

11. Sustaining the RCI: Funding Models

Models listed here are presented in *ascending order of viability*. Hybrids of these models are also possible, such as employing campus subsidies to provide start-up funding, followed by an NGN-like model to sustain ongoing activities.

Recharge

This self-supported, cost recovery model is used widely at UCSD. The main advantage is the relative ease of establishing a service. Income varies based on demand for the services offered. Equipment is funded through a reserve fund; initial expenditures result in a preliminary deficit. It may take several years to produce adequate income and to clear the initial equipment deficit. Annual rate reviews are required, leading to varying, unpredictable costs to customers. Billing overhead increases the cost of the service. Encourages users to seek alternatives that may not be as cost effective. The services can end in deficits stemming from waning demand or cost inefficiencies.

Overhead Rate Increase

Funds from indirect cost recovery (ICR) could legitimately be used to create and sustain these services. These funds are currently over-committed.

Campus Funding

Funding is obtained directly from the Chancellor and/or the SVCAA, VCHS and VCMS are assessed for their fair share of the costs. The assessed areas choose a fundraising method. One technique is to identify and redirect funds currently earmarked to researchers for the creation of facilities that could instead be provided by the CI project. Extensive consultation with the funding sources would be needed on an ongoing basis to maintain what would be a long-term commitment.

NGN-like

Puts an assessment on a population of users. Examplar populations:

- all researchers
- all researchers with grants
- all researchers with grants above a certain dollar amount
- all researchers with projects that could potentially use the facilities and services

This model encourages use and discourages users from creating duplicative facilities and services. The concept has been validated by the successful eight-year experience of funding network services by this method. Some of those assessed will not be able to use the services provided due to the nature of their projects and the types of services offered by the facility; they will object to having to pay for something they can't use. Despite the campus experience with this model, extensive consultation would be needed before its implementation.

12. Governance and Oversight

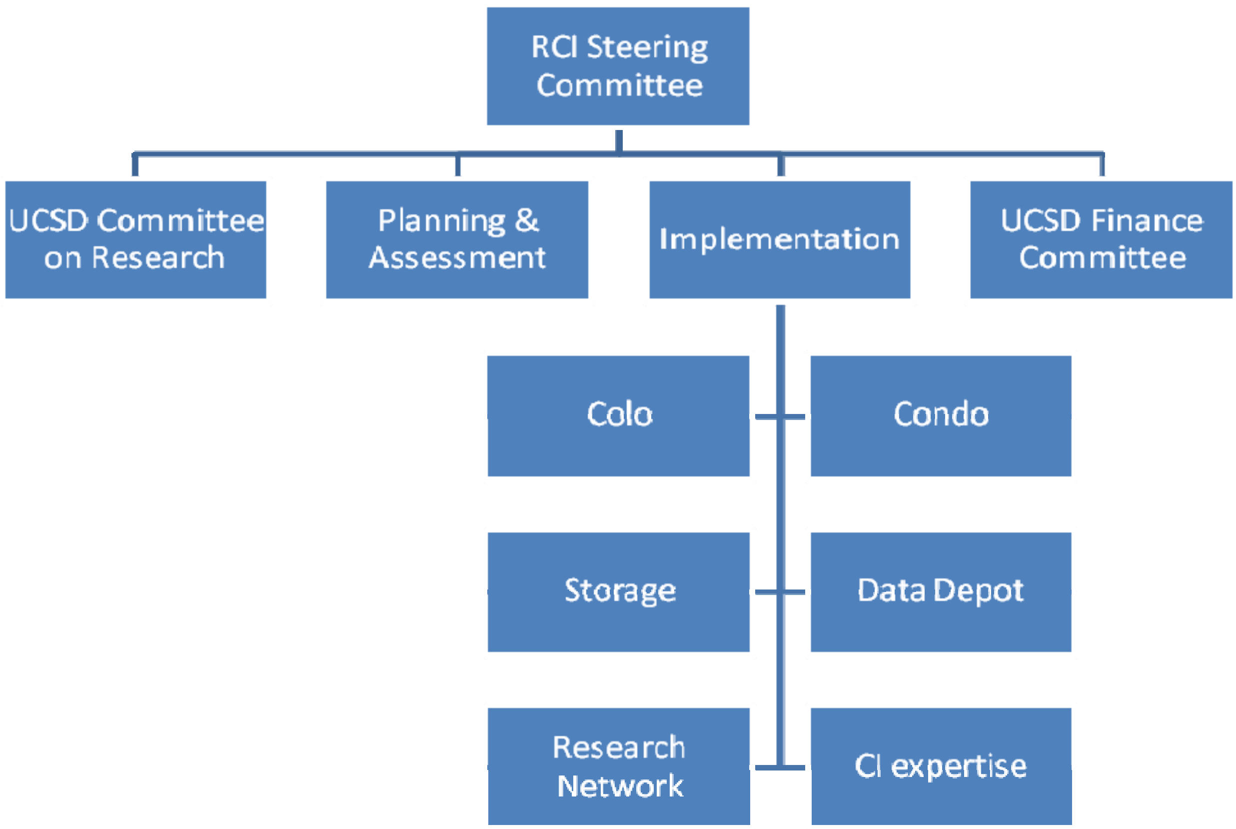

Table 12.1: Possible governance structure for proposed RCI.

Table 12.1 shows a possible governance structure for the proposed RCI project. A standing steering committee composed of the VCR, 3-4 faculty members, 2 external advisors, and the chairpersons of the planning and implementation committees would set direction, track progress, approve proposed access and allocation policies. They would coordinate with the UCSD Academic Senate's Committee on Research on research needs and satisfaction, and with the UCSD Budget Committee on budget matters. The project itself would be led by a RCI Program Manager who could be a senior staff member in any of the participating service provider units, or could be a new hire within the Office of Research Affairs. The RCI program manager would chair the Implementation Subcommittee. The latter would be composed of project-level leads in each of the five core RCI elements; s/he would also be responsible for documenting and coordinating the CI expertise labor pool. It is envisioned that there would be a working group for each of the five core RCI elements, which would be composed of project personnel as well as some faculty who could help develop access and allocation policy. Alternatively, there could be a Policy

Subcommittee reporting to the Steering Committee to do this function. The Planning and Assessment Subcommittee would further develop the implementation plan, track progress, develop metrics, and work closely with the Implementation Subcommittee on technical matters, and with the Steering Committee on direction, policy, and budget.

References

[1] Atkins Report, "Revolutionizing Science and Engineering through Cyberinfrastructure", Report of the National Science Foundation Blue Ribbon Advisory Panel on Cyberinfrastructure

[2] Dr. Jonathan Sebat, private communication

[3] Clifford Lynch, "The Institutional Challenges of Cyberinfrastructure and e-Research", Educause Review, November/December 2008.

[4] "Sustaining the Digital Investment: Issues and Challenges of Economically Sustainable Digital Preservation", Interim Report of the Blue Ribbon Task Force on Sustainable Digital Preservation and Access. F. Berman et al., http://brtf.sdsc.edu

[5] Szalay, A., Gray, J. "Science in an Exponential World", *Nature,* 440, 413-414 March 23, 2006.

[6] Hey, A.J.G., Trefethen, A.E. "The Data Deluge: An e-Science Perspective," pp. 809-824 in Berman, F., Fox, G.C. and Hey, A.J.G. Eds., *Grid Computing – Making the Global Infrastructure a Reality*, Wiley and Sons (2003). http://eprints.ecs.soton.ac.uk/7648

[7] Anderson, C. "The End of Theory: The Data Deluge Makes the Scientific Method Obsolete," *Wired Magazine*. 16.07 (June 23, 2008). http://www.wired.com/science/discoveries/magazine/16-07/pb_theory

[8] "Harnessing the Power of Digital Data for Science and Society", Report of the Interagency Working Group on Digital Data to the Committee on Science and the National Science and Technology Council, January 2009

Appendix A: Research Cyberinfrastructure Design Team Members

Prof. Michael Norman (co-chair), Physics Dept. and San Diego Supercomputer Center

Dr. Philip Papadopoulos (co-chair), San Diego Supercomputer Center and Calit2

Luc Declerck, Libraries

Declan Fleming, Libraries

Prof. Terry Gaasterland, Marine Biology Research Division, Scripps Institute of Oceanography

Dr. Elazar Harel, Administrative Computing and Telecommunications

Dr. David Hutches, Jacobs School of Engineering

Prof. Trey Ideker, Bioengineering Dept.

Dr. Reagan Moore, San Diego Supercomputer Center

Prof. Bing Ren, Cellular & Molecular Medicine Dept.

Prof. Barbara Sawrey, Chemistry & Biochemistry Dept. and Undergraduate Affairs

Prof. Amin Vahdat, Computer Science & Engineering Dept.

Tony Wood, Academic Computing Services

Prof. Shahrokh Yadegari, Theater Dept.

Appendix B: Charge to Committee

Dear Mike and Phil,

Thank you for agreeing to serve as co-chairs of the Research Cyberinfrastructure Design Team (RCIDT).

The creation of the RCIDT began with the San Diego Supercomputer Center Working Group (SDSCWG) Report, which was issued in March, 2008 and is attached. That report endorsed a campus-integrated scenario for SDSC and provided a rough blueprint for achieving integration. The Chancellor and Senior Vice Chancellor – Academic Affairs requested additional details, including operations and business models, for how UCSD could develop the pre-eminent university research cyberinfrastructure (RCI) given the campus's traditional strengths in high performance computing (HPC), data storage, and networking and the outstanding foundational elements represented by SDSC, Calit2, ACT and the University Libraries.

In order to develop an appropriate set of recommendations to the campus, I am asking you to lead a Design Team that will comprise key individuals from the faculty, IT staff and administration. The goal, as one of our RCIDT members put it, should be to ensure that all of our scholars at UCSD will become more competitive as a result of investments the campus will make in research cyberinfrastructure.

Specific elements of your charge include the following:

- Survey the campus to determine scholar needs related to RCI.
- Identify synergistic opportunities across SDSC, Calit2, ACT and the University Libraries.
- Outline possible operations and business models for co-location and condo types of cluster arrangements that incorporate principles of "green" cyberinfrastructure.
- Determine how best to meet the needs of the UCSD HPC community.
- Characterize needs and opportunities for curation of data.
- Explore networking infrastructure that will be needed to support RCI.
- Consider how campus RCI can interface with UC-wide needs and external providers like CENIC and I2; UC connections should be informed by the December 2007 UC Information Technology Guidance Committee (ITGC) report, http://www.universityofcalifornia.edu/itgc/ITGC_final%20report.pdf,
- Although not a primary focus, consider educational opportunities provided by RCI.

- Determine metrics of success for RCI.
- Outline possible strategies for meeting future RCI campus needs.

The last bullet recognizes the highly dynamic nature of RCI. Please have the RCIDT give some thought to the kind of structure(s) the campus should put in place to ensure that UCSD can respond in a timely manner to the exciting developments that lie ahead.

Thank you for taking on this important leadership role for the campus.

Best regards,

Arthur B. Ellis
Vice Chancellor for Research

Appendix C: Presentations to the Design Team

Prof. Larry Smarr, Director, CalIT2, “Green Cyberinfrastructure”

Prof. Fran Berman, Director, SDSC, “Next Generation SDSC”

Dr. Brian Schottlaender, Librarian, UCSD Libraries, “Digital Collections Management”

Dr. Reagan Moore, SDSC, “Data Management and Preservation”

Prof. Trey Ideker, Bioengineering, “Cal-IT^2 / School of Medicine Cyberinfrastructure Pilot Project”

Prof. Bing Ren, Cellular & Molecular Medicine, “Ren Lab Computational Infrastructure”

Prof. Michael Norman, Physics, “Data Analysis Challenges in Petascale Numerical Cosmology”

Dr. Philip Papadopoulos, SDSC and Calit2, “A Possible Research Cyberinfrastructure”

Dr. Philip Papadopoulos, SDSC and Calit2, “Strawman Storage”

Dr. Philip Papadopoulos, SDSC and Calit2, “PSDAF”

Anke Kamrath, SDSC, “Condo Clusters”

Dr. Rosio Alvarez, LBL, “Shared Computing Resources at LBL”

Appendix D: 2008 UCSD Research Cyberinfrastructure Survey

In April 2008 the CIDT conducted a survey of research cyberinfrastructure needs at UCSD. A survey questionnaire, was formulated by the team covering current usage and perceived needs for both resources and services in the core areas of computing; data storage, management and curation; networking; and advanced user services. The survey was posted online for 3 weeks at the commercial survey site www.zoomerang.com. A total of 45 responses were received, with roughly half responding as individuals, and the other half for labs, departments, or schools. The questionnaire and responses are reproduced here in their entirety.

Questionnaire

2008 UCSD Research Cyberinfrastructure Survey

The UCSD Research Cyberinfrastructure Design Team (RCDT) was established by the Vice Chancellor for Research to help identify current and future faculty research cyberinfrastructure needs. The team's membership is listed at the end of this document.

For the purposes of this questionnaire, *Cyberinfrastructure* (CI) consists of *computing systems, data storage systems, advanced instruments and data repositories, visualization environments, and people, all linked by high-speed networks to enable scholarly innovation and discoveries not otherwise possible.*

The team's work will proceed through four phases: (1) needs assessment (this survey); (2) conceptual design; (3) detailed design and technology assessment; and (4) implementation plan. The team will use the results of the survey to develop a campus-wide CI strategy that will be communicated to the University in the form of a report by the end of 2008.

This survey will help identify faculty research cyberinfrastructure needs, including software programs, compute cycles, data storage, data management, data archiving, networks, facilities space, expertise and programming. *We hope that you will choose to complete the survey.* Your responses will be confidential and any contact or identifying information you provide will be removed from the response database. The survey should take only about 15-20 minutes to complete.

Please feel free to skip or mark as "N/A" (not-applicable) any questions that are not relevant or applicable to your research needs. If you would like a member of the RCDT to contact you directly to discuss your needs please complete question #20 or email Michael Norman at *mlnorman@ucsd.edu.*

1) What is the "business" of your group, lab, department, ORU (or other campus unit)? And what are the predominant uses of cyberinfrastructure?

____

____

2) Please enter the name of your academic unit.

____

____

3) Type of response: Are you responding as: (Select one)

An Individual
For a Group
For a Center
For a Department
For a College or School
Other

If you selected other, group, or center, please specify:

____

4) Compute/CPU Needs
How are your research computing needs for "compute cycles" currently being met? (Select all that apply.)

Desktop system (Local personal computer, workstation)
Faculty/Research Group Cluster *
Department-owned Cluster *
San Diego Supercomputer Center (Allocations or other)
CalIT2
Grid Computing (Open Science Grid)
UCGrid Computing
National Systems (e.g., National Labs, Non-SDSC TeraGrid Allocations)
Not applicable
Other (please specify) *

* If you selected other, please specify. If you selected Department or Group cluster, please describe its configuration, where it is located and how it is supported.

____

Additional comments:

____

5) Compute/CPU Needs
If your research needs for "compute/CPU cycles" are not being met, please specify the unmet needs now and into the future (e.g., RAM, CPU, I/O BW etc. requirements). What is the driving need for this growth? Please enter "N/A" if this question is not relevant to your research needs.

____

____

____

6) If cluster computing services were provided on campus, which of the following would be of interest to you? (Select all that apply.)

Support to locate "rack units" in a central campus computer lab (e.g., with basic networking, power, facilities support)

Support to run a commodity cluster in a central campus computer lab (e.g., order/build your cluster, provide basic system and account administration).

Provide "Condo Cluster" services (e.g., integrate campus clusters to leverage idle cycles and provide a larger aggregate system for owners to use. Would include advanced scheduling options to ensure owners get optimal access, economies of scale for user support, software procurement, system administration, etc.)

Provide "Computing Cycles" (e.g., campus procures a large cluster and allocates cycle hours (e.g., via vouchers) to researchers/students across campus. Would include user support, centralized software, storage, etc).

Other (please specify)

If you selected other, please specify.

____

____

7) Software Needs:
What software packages needed by your research are currently not available to you and why? (e.g., cost, lack of campus site license, etc.) Please enter "N/A" if this question is not relevant to your research needs.

____

____

____

8) Software Needs
What are the major commercial or open-source software packages you are using for your research (e.g., matlab, sas, starP, gene annotation software, quantum chemistry, data mining, similarity searching, statistical analysis, bioinformatics, other analysis software)? Please do not include the personal productivity and communication software you use, such as word processing and spreadsheet tools, email clients, etc. Please enter N/A if this question is not relevant to your research needs.

____

____

____

9) Data Management, Analysis, Storage, Archival, and Access Needs
What are your greatest data management, analysis, access, archival, and storage needs? (Please select as many as apply)

Storage capacity, more short term (1-3 years) storage for research data

Ability to process and manage large quantities of research data

Data management software (Oracle, MySQL, SRB/iRODS, …)

Data analysis software (SPSS, SAS, Cognos)

Transferring experimental data to storage facility

Transferring data from storage to desktop or cluster

Sharing your data collection with colleagues (via web or resources such as SDSC DataCentral)

Access to national or community repositories (examples? PDB? NVO? GenBank))

Data Backup
Long term data preservation of large data sets
Metadata creation for large data sets for archival purposes
Long term access via a common repository
Data/format compatibility
Meeting data privacy/security requirements (FISMA, HIPAA)
Not applicable
Other (please specify)

If you selected other, please specify:

Additional comments:

10) Data Storage
Are your current research needs for data storage being met?

Yes
No
Not applicable

Additional comments:

11) Data Storage Needs
What do you anticipate your research data storage requirements will be for the next 2 years? (Please select one option)

1-500 Gigabytes
500 Gigabytes - 2 Terabytes
2-100 Terabytes
More than 100 Terabytes
More than 1 Petabyte
Not applicable
Other (please specify)

If you selected other, please specify:

Additional comments:

12) Data Archival Needs
What do you anticipate your long term research data storage requirements will be for the next 2 years? (Please select one option)

1-500 Gigabytes
500 Gigabytes - 2 Terabytes
2-100 Terabytes
More than 100 Terabytes
More than 1 Petabyte
Not applicable
Other (please specify)

If you selected other please specify:

Additional comments:

13) Networking Needs
Does the current campus network meet your current needs? If not, please be as specific as possible as to your particular areas of need, including peak and sustained data rates if possible. Please enter "N/A" if this question is not relevant to your research needs.

____

____

____

14) Networking Needs: Collaborations
Does your current or near-term future (next 2 years) research require connectivity to any national laboratories, research centers, or international collaborations?

Yes
No
Not applicable

Additional comments:

____

15) Cyberinfrastructure Services – Resources/Infrastructure/Administration
What types of cyberinfrastructure, resources or system administration would be most useful to your group, lab, department, ORU or other identifiable campus "unit"? This could be in the form of deployment of a cluster, providing robust storage, providing long-term data management, machine room space, and others. (Select all that apply.)

IT system administration
Cluster system administration (e.g., via Rocks, other)
Machine Room Space and/or Co-Location Facility
Compute Resources
Storage Resources
Networking Resources
Archival planning Resources
Metadata creation Resources
Not applicable
Other (please specify)

If you selected other please specify:

____

Additional comments:

____

____

16) Cyberinfrastructure Services – Expertise/Programming
What types of cyberinfrastructure expertise or programming support would be most useful to your group, lab, department, ORU or other identifiable campus "unit"? This could be in the form of programming and staff support, managing a software stack, and others. What kind of advanced expertise would be of value (visualization, parallelization, database design)? (Select all that apply.)

Interface/Portal development (GUI, Web-based)
Database and/or data management support (e.g., Schema design, management, tuning)
Scientific programming/Modeling
Visualization (Scientific, Medical, etc.)
Managing a Software Stack (builds, revision control, …)

Statistical support (e.g., survey design, analysis)
Software tailoring (e.g., porting code, scripting)
Software parallelization/optimization for clusters
Technical Documentation
Not applicable
Other (please specify)

If you selected other, please specify:

____

Additional comments:

____

____

17) Please enter any additional comments you would like to add about your research computing needs.

____

____

____

18) What would attract you to change from local/departmental clusters/storage to other providers (on or off campus)?

____

19) What are the pros and cons of the services that you currently use?

____

20) Responder Name and Contact Info: Completion of this field is necessary only if you would like someone to contact you about your research computing needs.

____

Survey Responses

Question 1. What is the business of your unit?

Research	Cyberinfrastructure R & D
Clinical	Clinical and biomedical
biology, biostatistics, and bioinformatics	Sensor nets
Computational biology	Shared access to compute & data
Climate and ocean modeling	Bibliographic data
Environmental observing systems	Cog. Sci data acquisition
Behavioral science	HPC
Functional MRI	Analysis and visualization
Business & economics	Application-driven CI
Seismic data collection	Website and teleconferencing
Computer science	Virtual organizations
Computational neuroscience	Massive data movement
Translational medicine	Metagenomics community
Experimental high energy physics	Data mining
Geophysics	Spatial information systems
Astrophysics	Academic library consortium
Ubiquitous computing	Digital collection management
	Data grid deployment

Question 2: What is the name of your academic unit?

Departments	Divisions/Schools
Pediatrics	SIO/IGPP
Bioengineering	UCSD Libraries
Cognitive Science	Biostatistics/Bioinformatics
Computer Science & Engineering	Social Sciences
Physics	Geosciences
	Pharmacy
Organized Research Units	Labs/Groups
SIO/IGPP	Behavioral Medicine Lab
White Mountain Res. Sta.	Nat'l Biomedical Computation Resource
CalIT2	Temporal Dynamics of Learning
SDSC	Natural Computation Lab
CASS	COMPAS
CRBS	SIO Marine Physical Lab
	Laboratory for Computational Astrophysics

Question 3: Are you responding as a (select one)

3.	Are you responding as (Select one):		
An Individual		17	38%
For a Group		18	40%
For a Center		2	4%
For a Department		4	9%
For a College or School		0	0%
Other, please specify View Responses		4	9%
	Total	45	100%

Other:
WMRS Assoc. Dir.
Research group and School of Pharmacy
Pacific Digital Library Alliance (PDLA)

Question 4: How do you meet current computing needs? Multiple responses

4.	Compute/CPU Needs: How are your research computing needs for "compute cycles" currently being met? (Select all that apply.)		
Desktop system (Local personal computer, workstation)		32	80%
Faculty/Research Group Cluster *		15	38%
Department-owned Cluster *		7	18%
San Diego Supercomputer Center (Allocations or other)		14	35%
CalIT2		8	20%
Grid Computing (Open Science Grid)		5	12%
UCGrid Computing		0	0%
National Systems (e.g., National Labs, Non-SDSC TeraGrid Allocations)		6	15%
Not applicable		0	0%
* Other or please describe. View Responses		13	32%

Other:
PRAGMA Grid, BIRN Grid, TeraGrid Data servers Instrument host computers Non-UC resources *fightsaidsathome.scripps.edu* CalIT2 visualization servers *Note:* no use of commercial cloud computing was reported

Question 5: What are your unmet computing needs? Free response

Basic desktop support, including HW refresh

Ever-increasing scale of simulations outstrip compute, storage, & viz. needs

- Scalable HPC is needed
- Large memory
- Single and muticore systems
- High I/O bandwidth for data analysis

More network bandwidth in/out bldgs. (Cog. Sci.)

Multi-terabyte data storage and analysis (multiple hits)

Powerful, Matlab-enabled computers

Education friendly computer labs in CSE
"we would desperately love to have a campus-provided data server, even on a recharge basis"

Question 6: If computer services were provided, which would be of interest?

6. Compute/CPU Needs: If cluster computing services were provided on campus, which of the following would be of interest to you? (Select all that apply.)

Support to locate "rack units" in a central campus computer lab (e.g., with basic networking, power, facilities support)	15	54%
Support to run a commodity cluster in a central campus computer lab (e.g., order/build your cluster, provide basic system and account administration).	12	43%
Provide "Condo Cluster" services (e.g., integrate campus clusters to leverage idle cycles and provide a larger aggregate system for owners to use. Would include advanced scheduling options to ensure owners get optimal access, economies of scale for user support, software procurement, system administration, etc.)	13	46%
Provide "Computing Cycles" (e.g., campus procures a large cluster and allocates cycle hours (e.g., via vouchers) to researchers/students across campus. Would include user support, centralized software, storage, etc).	11	39%
Other, please specify View Responses	9	32%

Other:
Commercial offerings/outsourcing
On-demand, large memory, data-intensive supercomputer
Dedicated server for proteomics

Question 7: What software packages are unavailable to you and why?

Clinical study database SW (cost)

Combined proteomics/statistical analysis (cost and availability)

Commercial statistical and mathematical analysis software due to cost/absence of site licenses

- SPSS
- SAS
- NCSS/PASS
- Matlab
- Mathematica

Spitfire, MpCCI, Neuroexplorer Analysis Program,….

Question 8: What major commercial or open-source SW are you using?

Numerous, including SPSS, SAS, NCSS/PASS, MATLAB, SRB, iRODS, BLAST, OpenSees, Abaqus, Ansys, LS-Dyna, Fluent, IDL

Many visualization, mapping, and animation packages

Question 9: What are your greatest data management needs (multi response)

9.	Data Management, Analysis, Storage, Archival, and Access Needs: What are your greatest data management, analysis, access, archival, and storage needs? (Please select as many as apply)		
Storage capacity, more short term (1-3 years) storage for research data		31	70%
Ability to process and manage large quantities of research data		31	70%
Data management software (Oracle, MySQL, SRB/iRODS, ...)		16	36%
Data analysis software (SPSS, SAS, Cognos)		15	34%
Transferring experimental data to storage facility		24	55%
Transferring data from storage to desktop or cluster		23	52%
Sharing your data collection with colleagues (via web or resources such as SDSC DataCentral)		22	50%

Access to national or community repositories (e.g., PDB, NVO, GenBank)		15	34%
Data Backup		37	84%
Long term data preservation of large data sets		28	64%
Metadata creation for large data sets for archival purposes		24	55%
Long term access via a common repository		20	45%
Data/format compatibility		16	36%
Meeting data privacy/security requirements (FISMA, HIPAA)		16	36%
Not applicable		1	2%
Other, please specify View Responses		6	14%

Other:
Web-based clinical research data and study management
Better bandwidth and security infrastructure
Software development tools (profilers, run-time…)
Real-time streaming at up to 7 Gb/s
BIRN, CCDB, NCMIR do all this – don't need
Enforcing management policies

Question 10: Are your current data storage needs being met? (Y/N)

10.	Data Storage Needs: Are your current research needs for data storage being met?		
Yes		25	60%
No		17	40%
	Total	42	100%
	View 20 Responses		

Response summary:
Yes, but barely; multi-TB soon (multiple responses)
Need reliable back-up (multiple responses)
SDSC too expensive, rolled our own
Use SDSC
Need high I/O bandwidth to storage at SDSC
Departmental solutions not adequate
Use distributed data grids, digital libraries

Question 11: What are your data storage needs over the next 2 years?

11.	Data Storage Needs: What do you anticipate your research data storage requirements will be for the next 2 years? (Please select one option)		
1-500 Gigabytes		5	12%
500 Gigabytes - 2 Terabytes		5	12%
2-100 Terabytes		21	51%
More than 100 Terabytes		7	17%
More than 1 Petabyte		0	0%
Not applicable		0	0%
Other, please specify View Responses		3	7%
	Total	41	100%

Other:
Around 500 TB
Data storage for increasing proteomic research
We currently manage more than a petabyte of data

Question 12: What are your long-term data archival needs?

12.	Data Archival Needs: What do you anticipate your long term research data storage requirements will be for the next 2 years? (Please select one option)		
1-500 Gigabytes		4	10%
500 Gigabytes - 2 Terabytes		3	7%
2-100 Terabytes		22	54%
More than 100 Terabytes		7	17%
More than 1 Petabyte		2	5%
Not applicable		2	5%
Other, please specify View Responses		1	2%
	Total	41	100%

Question 13: Does current campus network meet your needs?

1	No, researchers based in CTF down at Hillcrest Hospital have to cope with very strange 'security' requirements since we access network via the Hospital backbone. AS a result, numerous important sites are blocked for us.
2	Almost...I need to access several different network drives from multiple locations, including home
3	Calit2 has 1 Gbps link to every port. We need to have 10 Gbps links to key workstations/file servers/data centers for individual groups or workstations for visualization when large amount of data is visualized.
4	usually
5	Yes
6	there are web-crawling experiments i do not do because i feel they would be too network intensive
7	more or less. vpn/proxy falls way short of being easy to implement, explain and use. In library public service areas, we frequently experience wireless service overload (where patrons get 'locked out' of network due to huge demand). Personally, I think the wireless footprint could be fuller on campus; there are locations I've tried to work/meet with people over coffee and demo info products, and had variability of connectivity (mostly outside or near public areas where people commonly congregate). Improving footprint would mean it would be better for me to meet with faculty and students 'where they live' and provide my services seamlessly.
8	Yes.
9	Yes
10	n/a
11	yes
12	Need better QoS: less hanging.
13	Yes, we are able to use the optiputer network for fast data transfers to upper campus. Would like faster on/off campus data transfers.
14	10Gb ethernet would be helpful.
15	yes
16	yes
17	Campus network meets current needs. When I was running my lab, we needed 1 Gb from CSE (in AP&M) to SDSC which we got.
18	We have not tried to transfer large 3D image data sets ~10 GB each, perhaps sets of 10-20. If that can be done quickly, it would be very useful.
19	Need better wireless access and always faster networks
20	I need dedicated 10G optical paths to various sites on campus and to national and global collaborators
21	Yes, however we also made use of 10 Gb links from ESNET and TeraGrid with our cluster at SDSC. We can use large fractions of 10 Gb links when we are moving data.
22	Ok for now. I am unsure about data rates and other numeric data to meet our needs. Will defer this question.
23	We would definitely benefit from improved bandwidth between Ritter Hall at SIO and the SDSC machine room for transferring large data sets to our Digital Library storage.
24	No. We need, and are building, 10GE layer 2 networks into all our servers and clusters so that we can move data much faster and more reliably (that is real-time) than the campus/regional/state/national/international Internet2-style services allow.
25	Some parts of campus still have only 100 megabits to the desktop; for instance, the Social Sciences Building.
26	Some parts of campus still have only 100 megabits to the desktop; for instance, the Social Sciences Building.
27	Yes.
28	Speed and capacity of electronic communications will probably need to be increased in the near future.
29	we use Calit2 10G network links extensively
30	yes
31	Yes, moving a petabyte of data per year is only 35 MB/sec
32	yes

Question 14: Need for broadband external connectivity (not just email)

14. Networking Needs - Collaborations: Does your current or near-term future (next 2 years) research require connectivity to any national laboratories, research centers, or international collaborations?

	Count	Percent
Yes	28	76%
No	9	24%
Total	37	100%

View 14 Responses

1	Via web-based tools
2	We have many international and national collaboratories and they would all benefit from better network infrastructure.
3	pacific rimcollaborations
4	The TDLC GRID will need to be nation and international.
5	a good place to start would be the LJ Mesa! ie, interactions with TSRI networks, security, etc are currently difficult
6	CDL (http://www.cdlib.org/inside/projects/preservation/webatrisk/ and collections listed here http://www.cdlib.org/inside/resources/, particularly eScholarship. I support researchers from SD Supercomputing and Econ & Rady who need financial datasets from commercial vendors (WRDS, GIS).
7	Assume you mean at higher bandwidth than currently available...
8	No
9	We need sufficient bandwidth and network hardware/software to support teleconferencing. Because we are at a remote location (Bishop, CA) this is especially important to us. We also would like to increase bandwidth (currently 100 BaseT) to our remote high elevation labs in the White Mountains
10	Several ongoing collaborations with outside resarch groups
11	many international and national laboratories as well as other uni8versities.
12	Japan, Korea, Taiwan, New Zealand, etc...
13	International long term access required for SIOExplorer Digital Library and IODP Site Survey Data Bank
14	We are building what we call a Global Planning Grid in partnership with SDSC and the Worldwide Universities Network (WUN), for details see: http://gpeig.org/GRID-splash.htm

Q15: What kind of CI services and resources would be most useful?

15. Cyberinfrastructure Services: Resources/Infrastructure/Administration What types of cyberinfrastructure, resources or system administration would be most useful to your group, lab, department, ORU or other identifiable campus "unit"? This could be in the form of deployment of a cluster, providing robust storage, providing long-term data management, machine room space, and others. (Select all that apply.)

IT system administration	20	48%
Cluster system administration (e.g., via Rocks, other)	16	38%
Machine Room Space and/or Co-Location Facility	20	48%
Compute Resources	20	48%
Storage Resources	34	81%
Networking Resources	24	57%
Archival planning Resources	25	60%
Metadata creation Resources	16	38%
Not applicable	1	2%
Other, please specify View Responses	9	21%

Question 16: What kind of CI expertise would be most useful for your unit?

16. Cyberinfrastructure Services - Expertise/Programming: What types of cyberinfrastructure expertise or programming support would be most useful to your group, lab, department, ORU or other identifiable campus "unit"? This could be in the form of programming and staff support, managing a software stack, and others. What kind of advanced expertise would be of value (visualization, parallelization, database design)? (Select all that apply.)

Interface/Portal development (GUI, Web-based)	22	55%
Database and/or data management support (e.g., Schema design, management, tuning)	25	62%
Scientific programming/Modeling	15	38%
Visualization (Scientific, Medical, etc.)	23	58%
Managing a Software Stack (builds, revision control, ...)	11	28%
Statistical support (e.g., survey design, analysis)	11	28%
Software tailoring (e.g., porting code, scripting)	14	35%
Software parallelization/optimization for clusters	15	38%
Technical Documentation	13	32%
Not applicable	3	8%
Other, please specify View Responses	7	18%

Question 17: Please enter any additional comments you would like to add about your research cyberinfrastructure needs.

Comments ran the gamut from:

- Need campus-level integrated CI (majority opinion)

  to

- Don't spend money on centralized solutions; give it to the campus PIs (minority opinion)

General expression of need for more and better integration and interoperability

Need for CI domain expertise (e.g., data base, data mining, visualization)

24x7 operational support for global collaborations

Question 18: What would attract you to change from local/departmental solutions to other providers?

- Cost and cost effectiveness
- Upgrade path
- Improvements in uptime, performance, reliability
- Ease of use
- On-campus resources are preferable to off-campus solutions
- Make my problem someone else's problem

Question 19: What are the pros and cons of current solutions: local solutions

	Pros	Cons
Local solutions	- Control (customizability) - Dedicated access to compute - Cost?!	- Short staffed or not at all - IT support parasitic on researchers' time - Facilities max'ed out - No sharing of common solutions and sometimes data
Centralized solutions	-Economies of scale in housing, operating, and resource capitalization - Stability and continuity of expertise - Higher level integration - Range of CI expertise at SDSC	- Level of service (consultant responsiveness) - Who do I talk to about….? - Compute job turnaround suffers if resources are oversubscribed - Control/customizability

Appendix E: Consolidated Budget

RCIDT Budget Summary

RCI component	Startup	Annual
Co-lo facility	$ --	$ 1,235,000
Centralized storage	$1,080,000	$1,606,400
Research Data Depot	$ 850,000	$1,123,000
Condo clusters	$ --	$1,173,030
Research network	$3,517,500	$ 625,000
TOTAL	**$5,447,500**	**$5,762,430**

CoLocation

	$	Comments
Start-up Costs		
	0	See networking
Annual Cost		
Rack Co Lo	$292,500	SDSC Rate - $6500. 45 UCSD racks
Utilities	$585,000	Using 10kw rack * 45 racks
	$877,500	

Centralized Storage

	$	Comments
Start-up Costs		
Core Infrastructure	$1,000,000	$500/TB, 2PB (Raw) Purchase
Management Servers	$80,000	Parallel File System Support Servers Gateways
Networking	-	See Networking Costs
Subtotal	***$1,080,000***	
Annual Cost		
Labor	$1,080,000	Ramp over several years. 9FTEs in limit
Refresh Hardware	$526,400	Storage and server refresh. Expand storage annually.
Subtotal	***$1,606,400***	

Data Depot

	$	Comments
Start-up Costs		
Storage	$700,000	250 TB Storage for each of the 2 remote sites
Servers	$150,000	Server for each of the 2 remote sites
Subtotal	***$1,080,000***	
Annual Cost		
Labor	$840,000	Labor (including benefits) for 7 FTEs
Technology Refresh	$283,000	Replacement of servers and storage on 3-year cycle
Subtotal	***$1,123,000***	

Condo cluster

	$	Comments
Start-up Costs		
	0	Leverage SDSC Triton and UC3 South
Annual Cost		
Labor	$540,000	4.5 FTEs
Software	$30	
CoLo	$130,000	$6500 * 20 Racks
Utilities	$170,000	$8500 * 20 Racks
Capital	$333,000	For small refresh, interconnect, etc.
	$1,173,030	

Networking

	$	Comments
Start-up Costs		
Core infrastructure	$700,000	Network, fibre links, distribution switch, backbone WDM, support staff
Building Connections	$1,250,000	Connect 25 Buildings (1-4 labs/building). Assume 50 connections.
Lab Connections	$1,567,500	25 Labs semi dedicated, 25 labs dedicated
Subtotal	***$3,517,500***	
Annual Cost		
Operating costs	$625,000	Refresh, scaling and staffing (includes $1000/year maint. Of lab equipment)
Subtotal	***$1,123,000***	

Appendix F: Energy Efficiency Best Practices

The *EPA Report on Server and Data Center Efficiency* lists a number of techniques for improving the aggregate energy efficiency of typical data centers. Of these, some are particularly relevant in the context of a UCSD research computing colocation facility and are summarized below. For complete details, reference the linked document above.

- Improved physical and logical server consolidation.
- Elimination of unused or underused physical servers.
- Common and comprehensive control of power management systems for and in all systems.
- Improved airflow management, particularly via cold aisle/hot aisle containment systems. In simple terms, hot aisle containment is designed to evacuate heated air and direct it back to the air conditioning equipment as efficiently as possible. Cold aisle containment focuses on directing the outlet air from the air conditioning equipment as efficiently as possible to the front of the racks in the cold aisle. These systems yield improved efficiencies by
    - Preventing the mixing of supply (cold) and return (hot) air. Cold aisle and hot aisle isolation systems reduce cooling load and improve efficiency by supplying a high inlet temperature within recommended limits without the excessively low supply temperatures required to compensate for non-isolated, single mass airflow models.
    - Allowing the temperature of the return air to the data center cooling system to be as high as possible within the design limits of the cooling equipment. Fluid thermodynamics indicates that the higher the return temperature, the more efficient the heat transfer and thus the cooling process; all cooling systems will yield higher capacities (efficiencies) when return air is warmer. Without supply/return isolation (see above), mixed airflow forces return air temperatures to be significantly lower than those that would permit highest efficiency heat transfer.
    - Implementing a centrally controlled air distribution system that supplies cold air to and removes hot air from systems dynamically as a function of generated heat loads, via variable speed fans and other technologies. Historically, data center managers could use rack load design points and assume that these design points applied to every rack. This allowed them to calculate cooling and other infrastructure requirements using averages and relatively simple math. The subsequent cooling system design simply flooded the entire data center with cooled air. Today's reality is very different. Virtualization, blade servers and other technologies have created a situation where one

rack can have a 3kW load and the next rack a 30kW load. In a large data center these high density racks can be distributed in an unpredictable pattern across the floor, and further, the power and cooling patterns change over time as a function of system loads and equipment changes. A dynamically adjustable system ensures that the volume and possibly the temperature of cold supply air are matched with finer granularity to varying equipment cooling requirements. Such a model reduces the aggregate cooling load as well as the net power required to move air through the system.

- Improved transformers and uninterruptible power supplies. For example, best-in-class UPS systems have 70% less energy loss than legacy UPS systems at typical loads.
- Improved efficiency chillers, fans, and pumps.
- Cooling and heat reclamation systems integrated with the use of fuel cells and other distributed generation systems, in which waste heat is used to provide cooling.
- Common, energy efficient standards for rack and equipment installation, including blanking panels, cable routing, power distribution units, etc.
- Equipment acquisition and replacement policies that implement a uniform standard for energy efficiency that changes over time as technologies improve. These include common protocols for the recommendation and introduction of energy efficient compute, storage, and network equipment, and a governing model for the retirement and replacement of old and inefficient equipment. Such policies can be particularly effective in large data centers where there is a relatively regular flow of equipment that is replaced as a result of age, failures, or support expiration. The policies may further identify which energy efficient systems should be used as replacements. For example, a retired four-way server can be replaced with a two-way dual-core server.

Appendix G: Bylaws for the Governing Board of UC Regional Data Centers

Appendix H: SDSC Colocation Facility and Service Level Agreement