BROWN

Cyberinfrastructure in Research and Learning at Brown University: Vision, Status, and Roadmap

Apr 8, 2008

Jan S Hesthaven
Professor of Applied Mathematics
Director of the Center for Computation and Visualization

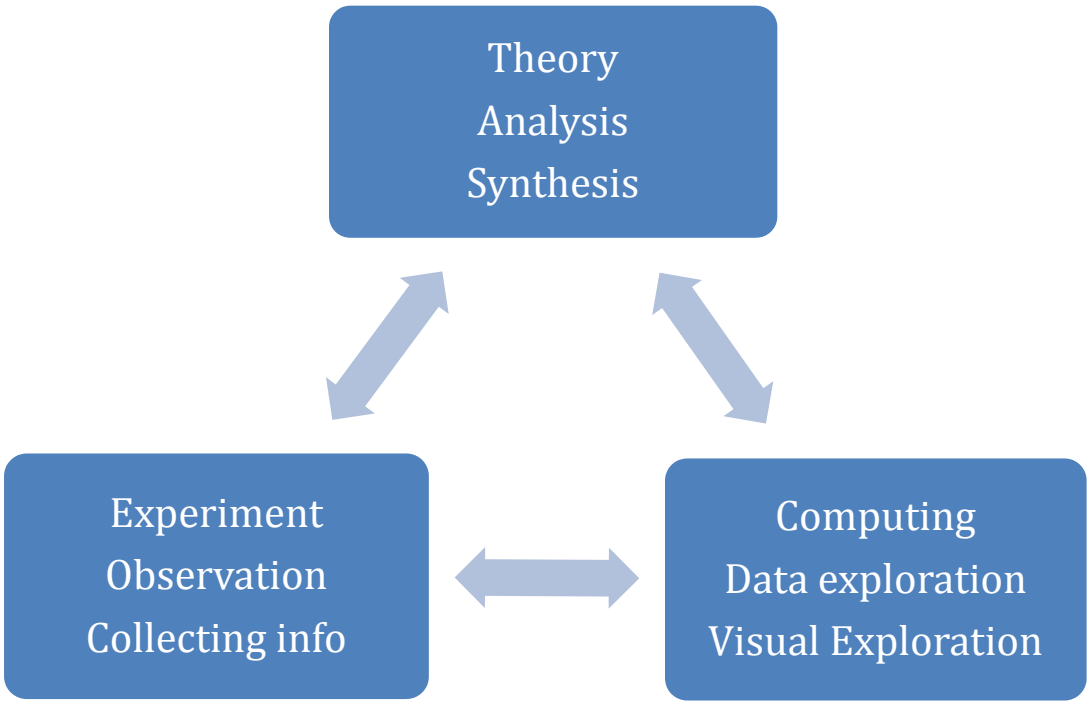

Note: This report represents a personal view although it is based on very broad input from faculty, staff and students from across Brown through the Working Groups for the IT Strategic Plan, input through CCV, and numerous focus group meeting organized by OVPR and the Library during the spring of 2008.

I. Introduction

During the last 10-15 years, universities worldwide have been transformed by the explosive and novel use of computers and IT throughout the organizations. Whereas much of this growth was initially seen in administrative and research activities; for the latter case mostly in the sciences and engineering; the impact has more recently become truly universal and computing now reaches deeply into teaching, learning, communication, social interaction, and most other areas of a modern university.

This transformation has been enabled by the spectacular growth in the speed and ability of computers we have witnessed during the last three decades. This growth, often described by Moore's law, is characterized as a doubling in performance every 18 months with a similar growth in performance-price ratios. Such sustained growth over decades is unprecedented in the world of science and technology and today's desktop computers are far more powerful than even the most expensive supercomputers were just a decade ago. This progress has, however, only wetted the appetite for solving even larger and more complex problems, and encouraged researchers to begin to consider problems previously thought to be out of reach for quantitative analysis.

Many recent activities represent examples of entirely new research areas emerging as a result of new possibilities offered by the innovative use of powerful computers and software, large displays, large data set analysis, network-based applications etc.

These innovations appear across the university in all disciplines and impacts teaching, learning, and research throughout the

It is equally important to appreciate how this growth and proliferation in the use computers has impacted and inspired researchers throughout the universities and not just in the sciences and engineering. During the last 5-10 years, we have witnessed the spread of novel, non-traditional, and innovative use of computers into research activities throughout any research university to the level where such activities now provide an essential and enabling component in the contemporary pursuit of knowledge, insight, and innovation. Such activities include the life sciences with its growing need for quantitative methods and bio-informatics, social and economic sciences with the analysis of demographics and financial data, the humanities with novel developments in visual arts and textual analysis and in art history with creation of virtual models of ancient cities and so forth. Whereas the traditional view of computers and their use was to consider this technology as way to make labor intensive research more effective, much of these more recent activities represent examples of entirely new research areas emerging as a result of new possibilities offered by the innovative use of the powerful computers and software, large displays, large data set analysis etc.

The prominence and impact of computing on research and learning across a modern research university is entirely comparable to that of laboratories and libraries.

Whereas the use of computers in research used to be directed mostly toward the sciences and engineering, the prominence and the impact of computing across a modern research university is now entirely comparable to that of laboratories and libraries. In fact, many of the issues we shall discuss shortly are shared with that of a research library system and the infrastructure and support required. This underlines the central observation that a leading university must ensure the availability of such computing and support infrastructure at adequate levels to enable students and faculty to ask questions of contemporary interest and to work at the frontiers in their respective disciplines. It also emphasizes that this requires a

substantial and sustained University commitment of capital investment and operating budget.

By **cyberinfrastructure,** we will in the following understand all computing infrastructure and support that enables faculty and students to pursue their research interests. This includes large scale high-performance computing, computationally intensive local or central research laboratories and large scale visualization/immersive/virtual environments as well as access to emerging models for large scale computing such as grid- or cloud-computing. It also includes data intensive applications, large-scale storage/backup/archival facilities and support for cyberinfrastructure applications to enable the efficient use of such facilities. The institutional network infrastructure is likewise an essential component of this, both in terms of internal connectivity and for off-campus connectivity for use of remote facilities, national resources, and data-intensive communication, e.g., multilink video-based interaction. We shall also generally use cyberinfrastructure to refer to the use of local resources in research down to the level of individual desktop/portable systems, data collection, and the analysis and modeling of data and non-data-driven research.

II. A vision for cyberinfrastructure to enhance research and learning

While the impact and the importance of cyberinfrastructure throughout the academe is undisputed at this time, it is reasonable to discuss what we should aim to have in place at Brown to fully benefit from this and to encourage and enable this transformation of how knowledge and innovation is generated at a modern research university.

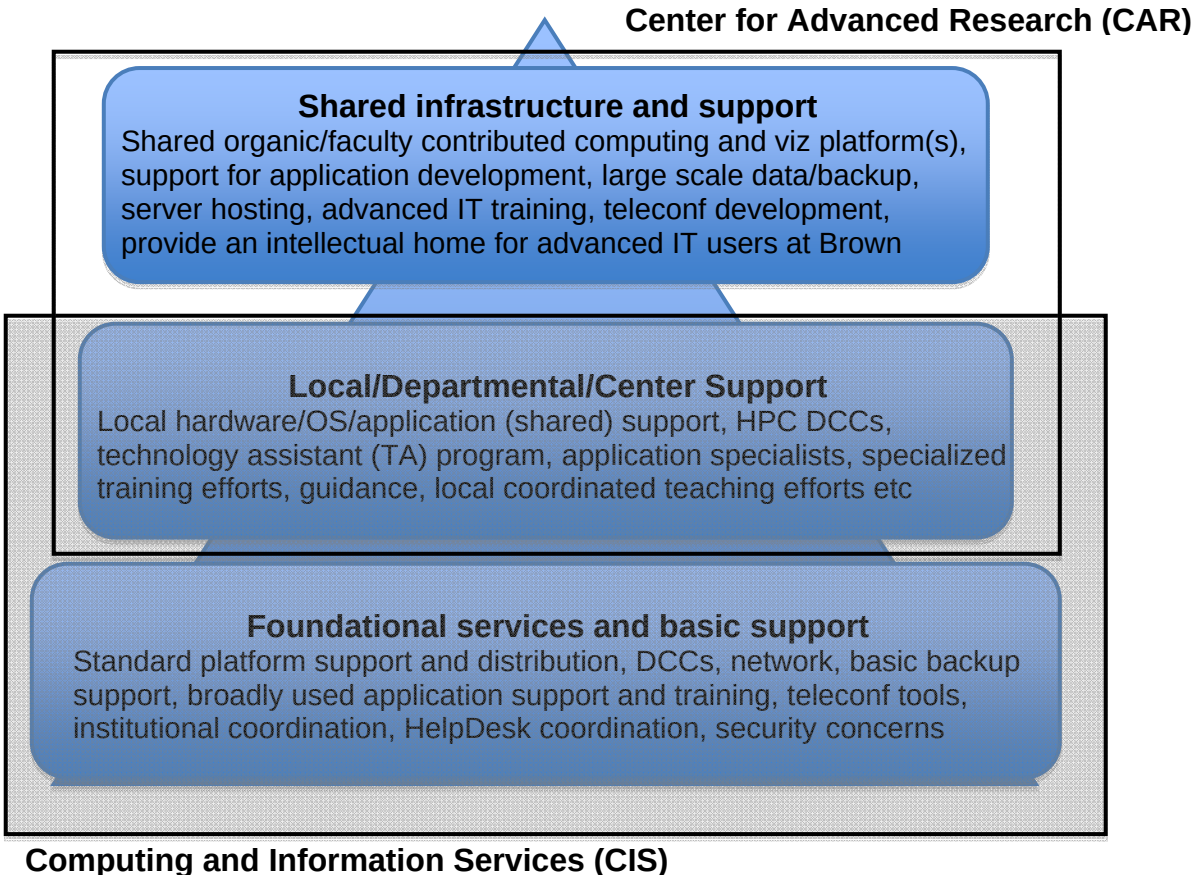

Foundational services and basis support.
These are services that should generally be expected from a central computing organization such as the Computing and Information Services (CIS). An essential institutional service is the network on which all cyberinfrastructure relies. This network must be maintained and updated to ensure a high-performance institutional network to enable data-intensive work and communication across the institution. However, high-performance off–campus connectivity is of equal importance to enable collaboration, access to remote instrumentation and national resources. Key foundational services furthermore include support and standards setting for widely used systems (Windows, Mac, Unix/Linux), coordination of desktop support in departments, basic backup infrastructure for desktop/laptop backup, and the support and training for applications used widely across campus. There must also be a central coordination for HelpDesks across campus, coordination for software purchases and license negotiations for widely used applications and the support and coordination of teleconferencing facilities across campus. All services at this level must be free for the user/department and should be considered essential components of the infrastructure.

A peek at history:

In the 70'-80' Brown University was a leader an innovator in the use of cyberinfrastructure

Brown was the first research university to install a campus-wide broadband network

Brown was among the first to adopt advanced computing and create software focused on need in the humanities

The research institute IRIS was a leading innovator in the development of modern hypermedia, now used daily in web-browsing

Brown was competing with MIT and CMU as one of the three leading high-tech institutions defining campus-wide innovation in the use of advanced computer technology

Local/Departmental/Center support
The nature of research/advanced computing is such that significant resources will have to be located in close proximity to the research activities. To support such activities, a centrally coordinated and university-employed and supported pool of specialized (shared) departmental support staff (High-Performance Computing Departmental Computing Coordinators (HPC DCC's)) must be available, assisted by a group of specially trained graduate students/technology assistants (TAs) to offer local/specialized support and training. This also provides opportunities for mentorship models for grad students who may be able to find "niches" that align with some of their research interests.

To enable such centralization to be efficient and practical it also requires some efforts in standardization and limits in the support for both software and hardware.

There should likewise be centrally coordinated application specialists to assist in larger scale developments and research efforts (Scholarly Technology Group (STG) and Center for Digital Initiatives (CDI) are examples from the humanities which could be expanded to the rest of university). Responsibility for these activities shall be coordinated between CIS, the Center for Research Computing (CRC), and the University Library. Whereas some of these services must be offered free of charge, there is clearly

the opportunity for some cost recovery through grants and contracts; e.g., by establishing a centrally coordinated group of advanced support staff which research groups and departments can purchase a fraction of for local support.

Shared infrastructure and support
A centrally located facility with substantial computing and visualization resources must be available, renewed, and consistently supported. This can be established partially (but not fully) through faculty-contributed hardware, but support staff must largely be University supported. The general infrastructure (storage, network, machine room etc) must be university funded. There must be support staff to assist in the use/training of advanced technology, including close collaborative efforts with faculty research projects. The shared facility also hosts servers and other advanced hardware for individual faculty groups and provides the infrastructure and support for large-scale data manipulation, large-scale visualization, and backup/archival of research data. The responsibility for these activities shall rest with CRC, which becomes a home and a focal point for research and learning activities in advanced computing and technology across campus.

Coordination and communication
CRC and CIS must be coordinated and engaged in a continued communication to enable and ensure an organic and flexible support structure, adapt to new challenges and the needs to revisions to supported standards. However, merging of the two organizations is likely not beneficial due to the many other activities and responsibilities in an organization like CIS and the specialized nature of CRC.

One way to ensure this close ongoing connection is by establishing a position as Assistant Vice President for Research Computing with a dual reporting line to the CIO and VP of Research. Research computing must be an important part of ongoing discussions in university committees such as Computing Advisory Board (CAB) as well as with entities across campus with related activities, e.g., the libraries, the medical school, and the hospitals.

Teaching and learning
To fully benefit from cyberinfrastructure in teaching and learning, these CIS and CRC must be closely connected to expertise in the development and use of advanced technology in the class-room and in teaching and learning in general. Prominent examples are the Sheridan Center for Teaching Excellence, the Undergraduate Resource Center, the Academic Technology Services group, and the University Library. These activities and cross-fertilization facilitate a growth in the use of advanced technology across campus, in inter-departmental/ interdisciplinary classes with participation from CRC staff and the expanded use of central facilities for computing and visualization at all levels of teaching and learning and across the university.

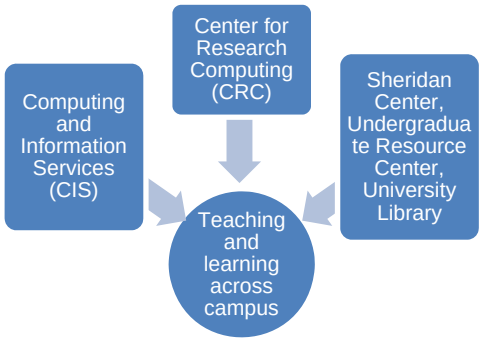

III. Brief overview of the current situation at Brown.

To appreciate the requirements needed to bring the above vision to life, let us briefly summarize the existing situation at Brown in the context of support and use of advanced computing across the disciplines.

Foundational services and basis support.

Some of these foundational services are currently provided by CIS, such as multi-departmental computing coordinators (MDCC), network, application support and training for a small number of applications, some support for Mac and Windows platforms etc.

However, there is no support for Unix/Linux, no efforts toward standardization and no training in this area. There is limited coordination between HelpDesk's across campus and many very widely used applications have no support at all, even for applications used extensively in classes (MATLAB is a prominent example with over 1000 users and no support). There is no coordination of software purchases and license negotiations and there are no teleconferencing facilities available for support, training, and development

No low-level backup and no shared resources for large file exchange is available.

While the institutional network backbone is currently undergoing an upgrade to 10MB, the out-going connectivity remains a limiting factor for many activities and novel modes of communication.

> ***Analogy: Imagine a research library***
>
> *There is a shared collection with renewal, growth, the development of new areas, and a physical home for the library. Departments may have local collections but can seek central advise and support. The collections continue to develop and users can seek help from knowledgeable staff. There is a centralized acquisition and extensive online collection as well as off site resources for less used resources ….*

Local/Departmental/Center support

This level of support and development of local resources currently does not exist. Local groups and centers are entirely on their own, receiving no support for hardware maintenance and support, there is no coordination/sharing of resources, no application scientists (with possibly STG(CIS) and CDI(University Library) being the exception) for local assistance and training.

The consequences of this complete lack of support has been to force research groups and departments to develop a variety of ways to deal with this problem – most often simply by putting the responsibility on faculty or graduate students. A few departments and centers across campus have the resources to hire support staff but often not the level required to fully benefit from the technology.

This not only leads to re-inventions and duplication of efforts across departments and research groups. It also exposes the university to significant security risks through the inadequate or non-existing support.

Shared infrastructure and support

All shared resources, currently located in the Center for Computation and Visualization (CCV), are aged and there are no resources/plans for renewal/growth etc. There is very

limited support for application development/support and some limited support for server hosting. Large-scale backup/archival capacity is available but no general data-storage.

CCV has not managed to provide a home for advanced computing, partially due to staffing limitations, partially due to poor physical facilities and partially due to historical circumstances. There is limited visibility of CCV, some coordination between CCV and CIS, and a general confusion of which services are offered by CIS, CCV and others.

All CCV resources and support are fee based except for purely educational activities.

Coordination and communication
Until recently, there was no communication between CIS, CCV, the University Library and similar entities across campus but situation has improved lately. Apart from regular meetings between CIS and CCV staff, the director of CCV has also been made ex-officio member of CAB and is in regular contact with the University Library. However, there is still much to improve in this, in particular in the area of HelpDesk support and coordination, efforts to avoid duplication, software purchases, long-term planning, visibility of IT services at Brown, support for teaching and learning, security management and control, emergency preparedness etc

… now imagine

There are no plans for growth and no funds for renewal and maintenance of existing resources. The library provides no support to local departments/groups and no central acquisition. The little available support is scattered and not visible. There are very limited online resources and no off-site space for less used material. There is no physical library site.

Occasional users have no access to the library and frequent users must pay.

Unimaginable perhaps, but close to the situation for cyberinfrastructure and its support at Brown.

Teaching and learning
CCV has been involved in many classes over the years, but these have mainly been reactive developments in response to specific requests. Very limited activity has been progressive to expand and develop new directions and initiatives. This has partly been due to staff limitations. There have been no interactions with other parts of the university in this regard. The exciting possibilities lying in extensive collaborations with RISD has never been fully explored.

IV. Roadmap.

As has been discussed in the above, the situation for support, use, and development of cyberinfrastructure at Brown is far from adequate and, in many instances, prohibits researchers and students from following their imagination in the pursuit of knowledge and innovation. Feedback from some faculty indicates alarming examples where research activities are indeed being driven by what resources are available rather than by what is possible and imaginable. Furthermore, an inadequate and uncoordinated support for such infrastructure exposes the university to significant security risks, ranging from loss of data to genuine security breach for restricted information by enabling back doors to the broad university systems.

It is worth emphasizing once again that this situation is not restricted to isolated parts of the university but it impacts the university across all disciplines in much the same way that a lack of modern laboratory facilities and excellent research libraries would. In this context, it is clear that a continued neglect of this integral component of a research university would directly collide with the vision of the Plan of Academic Enrichment.

In the following we will sketch a roadmap (approximately of five years) for how to bring Brown back on the track to the leadership position in the use and development of advanced computing technology it enjoyed in the 70's and 80's, but has since lost as discussed previously.

Phase 1: Year 1-3.
We recommend the following initiatives

Hardware/Infrastructure developments

- The seeding of a substantial centralized computing facility, possibly supporting several distinct computing platforms and models, and the development of a mechanism and incentives by which faculty can contribute to the growth of the hardware.

- Improved network infrastructure, both within the institution as well as to the outside through a lambda-connection.

- Deployment of one large-scale visualization facility for teaching and research purposes in a lecture hall on campus and a few smaller installations. Explore ways to support the development of a central fully immersive visualization facility as a unique facility, possibly with statewide utilization through an expanded network connectivity.

- The development of free basic foundational services, including access to shared resources, local hardware support, backup services for desktop/laptop, and support and training widely used application software. These services shall be developed and supported mainly by CIS.

- Expansion of video-conferencing capabilities to have at least one high-end facility for general use on campus. Development of collaborative tools and the establishment of a central development and training base for tele/video-conferencing and visualization over distances.

Goal

The formation of a physical and visible home for centralized computing, visualization, improved networking, and large-scale data-storage/ backup with adequate support and plans for renewal, and with substantial resources to assist and support local research, learning and teaching activities across campus.

Detailed strategies for partial cost-recovery and the development of incentives for faculty and research groups to use and commit to a central facility and support structure for research and learning.

- In a collaborative effort between CRC, CIS, and the University Library, develop and implement models for data archival and data curation facilities and services.

Support of research and learning

- The development of support of Unix/Linux desktop and locally hosted clusters. Standardization of supported systems and the establishment of the Unix/Linux HelpDesk, coordinated with similar local structures across campus.

- Creation of a cohort of research computing DCC's and application scientists, to enable shared support and application development for local research groups.

- The creation of a technology assistant program based on graduate students to complement the research computing DCC's, assist with HelpDesks, provide local training and basic application development, and engage in traditional TA services through CRC supported classes.

The current situation is the result of a decade long neglect and a lack of understanding of just how transformative advanced computing and visualization has become in our pursuit for knowledge, insight, and innovation across all university disciplines.

Advanced computing and IT-technology is an integral part of any modern research university and should be supported, renewed, and grown as any other essential parts of the university. It cannot and should not be expected to recover a substantial part of its cost from grants and contracts.

Coordination, communication and organization

- Coordinate IT support and research computing services across Brown and increase the visibility of these.

- Central coordination of software and license purchases.

- Develop a catalog of available services and facilities at Brown to help make make these visible and avoid duplication efforts.

- Develop incentives for faculty to use and develop classes extensively exploring the use of advanced computing technology. Commitment for sustained support for such classes.

- Develop incentives for faculty to use shared, collaborative services.

- Continued development of an appropriate organizational structure for advanced computing at Brown.

- Explore the opportunities for endowed support through the advancement office, including a naming opportunity.

Phase 2: Year 3-5.
This will consist mainly of implementing and expanding many of the initiatives begun during the first phase.

- The establishment of a physical and centrally located home for research computing and cyberinfrastructure at Brown with space to interact, seek information, and develop new collaborations. Consider possibility for a statewide impact.

- Expand the use of advanced visualization tools across campus to have several large-scale installations to support research and learning.

- Expand the support and development of tele-conferencing across campus in support of research, learning, and internationalization.

- Following a renewed evaluation of the needs, expand the programs for shared cyberinfrastructure, DCC's, application scientists, and technology assistants.

- Development of a cohort of interdisciplinary classes to support the use and development of advanced computing research and learning across campus.

- Explore the possibility for developing a Masters in Computational Science program.

- Secure long-term support through a capital campaign, in lime with the Plan of Academic Enrichment II, calling for improved support for research infrastructure.

.. by doing nothing we will experience:

Inadequate use of advanced technology and cyberinfrastructure in teaching, learning, and research.

Faculty forced to pursue what is the possible rather than what is imaginable.

Grants/contracts rejected, not submitted, or even returned due to inadequate support and infrastructure.

Students and faculty go elsewhere.

Faculty unable to develop novel modes of teaching and learning due to lacking infrastructure and support.

Junior faculty and new adopters are punished

.. some are happening now

Brown cannot afford this choice

V. Consequences for research, teaching, and learning of doing nothing.

While the list of initiatives listed above, as well as their projected cost, may seem

daunting and perhaps even unrealistic, it is paramount that we also be aware of the consequences of doing nothing.

There are already reports of faculty changing their mindset and their research from the imaginable to the practical as a consequence of lacking/insufficient resources. This is not only unfortunate -- it also goes against the very spirit of the Plan for Academic Enrichment and its very goal to maintain and grow the university as a world-class research and educational institution.

There are likewise instances where faculty express the hope that their grant proposals be unsuccessful as they otherwise find themselves unable to fulfill the hardware/ computing requirements of their proposed effort. This includes cases where the grant would provide the hardware. However, with no support from the university, the faculty member is left with a very significant workload to keep this hardware operational. This situation is made worse by recalling that cyberinfrastructure represents a rapidly evolving field in which obsolescence typically is measured in a few years.

The lack of adequate centralized cyberinfrastructure prohibits Brown form taking advantage of the economies of scale, leading to the development of silos, duplication of efforts and expenses, and wasteful use of resources. It also exposes the university to severe security problems by not having adequately trained staff addressing key issues of digital security and exposure through the network. The lack of such shared resources for leaves no space for experimentation. This particularly impacts young faculty and faculty in parts of the university where the use of computing and advanced IT is still nascent.

The current situation forces faculty to adapt their research and grant activities from the imaginable to the practical. There are likewise reports from faculty that funding has been denied due to a perceived lack of centralized computing and support by the funding agencies.

The complete lack of support for local activities forces faculty members to adapt to this situation. This often results in solutions, which involve the faculty member and/or graduate students maintaining advanced technology. This is impractical, problematic in terms of security, and, above all, a very expensive way of using the time of students and faculty. Furthermore, there could be legal implications both in terms of data-storage requirements and the time a faculty member spends on such activities when fully supported by grants and contracts.

The lack of adequate, indeed even basic, computing infrastructure to support research and learning already impacts the recruitment of faculty in several disciplines. It is likely, although only anecdotal, that it impacts the ability to attract technology proficient students, both at undergraduate and graduate levels, when our peers provide a superior resources and infrastructure for research computing and learning.

We must not miss this unique opportunity to develop new modes of teaching, learning and research across the disciplines by not having access to contemporary technology. We likewise must not lose opportunities to develop exciting interdisciplinary classes and research opportunities, which make use of advanced technology. The most severe consequence of inaction is that we will not become the research university that we aspire to be. A final chilling consequence is that our students will be unable to receive a

contemporary education at Brown due to lack of access to technology being used and developed within their disciplines at our peer institutions..