



Minnesota Supercomputing Institute External Advisory Board Findings

Dan Reed - Chair (University of North Carolina)
John Quackenbush (Dana Farber / Harvard University)
Jeremy Kepner (MIT Lincoln Laboratory)


April 3, 2007

Executive Summary

- **Enhancing the research reputation of the U of M is essential to achieving its long term standing goal**
 - Within this context, the MSI is the only University wide organization which can substantially impact the research of every major department
- **MSI plays a key role in enabling collaborations across the University and with outside organizations**
- **In the last decade MSI has fallen from a Top 5 to a Top 25 academic supercomputing center**
 - Mostly due to the relatively large improvements made by peers
- **U of M must identify a *specific* core research agenda which exploits its organic advantages over its peers**
 - Engage in setting the national research agenda
 - This entails taking some risks
- **MSI mission should be aligned to this research agenda**
- **Formal MSI mission should recognize its existing expanded role**
- **MSI (perhaps renamed “Minnesota Research Computing”) should be a U of M wide organization reporting directly to the VP of Research with joint investment from AHC (as this is MSI’s major growth area)**
- **MSI should conduct a national search for an External Director**
- **Current 16 committee structure should be reduced to one**
- **MSI budget is ~1% of U of M research budget, this is too low**

Outline

- **Introduction**



- *Board Members*
- *Board Charter*
- *Board Inputs*

- MSI Today
- Research Mission
- Organization
- Operations
- Summary

External Advisory Board Members

- **Prof. Daniel A Reed (University of North Carolina) - Chair**
 - Director of the Renaissance Computing Institute
 - Chancellor's Senior Advisor for Strategy and Innovation
 - Professor of Computer Science
 - Former director of NCSA
- **Prof. John Quackenbush (Harvard / Dana-Farber Cancer Inst.)**
 - Professor of Computational Biology and Bioinformatics, Department of Biostatistics HSPH
 - Leading researcher in using genomic and computational approaches to reveal the underlying biology
- **Dr. Jeremy Kepner (MIT Lincoln Laboratory)**
 - Senior Research Staff Embedded Digital System Group
 - Leads Lincoln's Parallel Computing and Grid Computing Efforts

Board Charter

- *Provide advice and guidance to U of M on the best strategies, programs, services and future direction for MSI*
- **Our focus is on how MSI can improve its impact on the U of M goals**
 - Most of our comments will be on concrete improvements
- **Should not viewed as criticism of MSI and its considerable track record of accomplishments**
- **Deliverables**
 - This briefing
 - Summary of findings in a written report

Board Inputs - Documentation

- **Comments on the long range planning committees' report**
- **Report of the previous MSI external advisory committee**
- **Overview of MSI systems administration and technical operations**
- **Overview of MSI user support**
- **Detailed results of the survey of MSI users**
- **MSI budget history**
- **Overview of the Digital Technology Center (DTC)**
- **The Executive Summary and the full report of the University of Minnesota Task Force on Research Infrastructure**


Board Inputs - Meetings

- **MSI/DTC senior staff: Hellriegel, Odlyzko, Olesen, Runesha, and Ropers-Huilman**
- **MSI Systems and Operations and User Support staffs**
- **Chairs of the MSI Long Range Planning: Siepmann and Kaznessis**
- **U of M Department Heads**
- **Biomedical Steering Committees**
- **MSI Long-Range Planning and Procurement Committees**
- **Task Force on Initiatives in High-Performance Computing**
- **MSI Fellows and Mayo Clinic**
- **Dean Crouch and Associate Dean Kaveh**
- **U of M administrators**

Outline

- Introduction

- **MSI Today**



- *Current Climate*
- *Current Status*

- Research Mission

- Organization

- Operations

- Summary

Current Climate

- **The U of M long term goal is to be among the Top 3 Public University in the nation**
 - **A realistic near term goal is to be in the Top 3 in the Big Ten**
- **Enhancing the research reputation of the U of M is essential to achieving this goal**
 - **Within this context, the MSI is the only University wide organization which can substantially impact the research of every major department**
- **Nationally it is recognized that computation has emerged as a 3rd pillar of science (along with experiment and theory)**
 - **Among peer Universities 50% are planning or building major enhancements to their supercomputing infrastructure (> 10K sq ft)**
 - **None charge user for computing cycles**
- **Within this research climate it is inconceivable for the U of M to improve its research reputation without the MSI**

Current Status


- **MSI has a strong and distinguished history in this field**
- **The amount of external research revenue that can be directly attributed to MSI involvement is impressive**
- **MSI has a strong base of computing capacity**
- **MSI staff are highly professional and their contributions to research support are widely recognized throughout the campus**
- **MSI has taken the initiative to reach out to the high growth biomedical research area**
- **MSI plays a key role in enabling collaborations across the University and with outside organizations**
- **Minnesota has the strongest supercomputing industrial base in the world**
- **Minnesota has the strongest high performance storage industrial base in the world**
- **In the last decade MSI has fallen from a Top 5 to a Top 25 academic supercomputing center**
 - **Mostly due to the relatively large improvements made by peers**
- **It is realistic goal for MSI to regain its former leadership position**

Outline

- Introduction

- MSI Today

- **Research Mission**



- *Research Agenda*
- *Potential Partnerships*

- Organization

- Operations

- Summary

Research Agenda

- **U of M must identify a *specific* core research agenda which exploits its organic advantages over its peers**
 - This would allow MSI to, once again, help set the national research agenda
 - This entails taking some risks, but failure to do so risks stagnation
- **MSI mission should be aligned to this research agenda**
 - Biomedical and biological applications are obvious growth areas that MSI is currently pursuing in an informal way
 - Traditional areas of excellence (e.g. collaboration with scientists in the Institute of Technology) need to be maintained
- **Funding needs to be commensurate with a revitalized research mission**
 - To be nationally competitive in modern computing applications would require major state investment (~\$20M) across U of M (some of which would enhance MSI)
 - DoD is a major funder of supercomputing technologies and research
 - There are significant opportunities here for those willing to pursue them
 - Other sources of funding should be investigated
- **MSI should aggressively pursue strong relationships with the local computing and high-performance storage industry**
- **MSI should have a research arm to develop next generation solutions for its user base**
 - This is essential for attracting talented faculty and staff
 - This arm will also be a magnet for outside direct investment

Expanding the Research Agenda

- **Science has changed and the research mission of MSI must adapt to maintain its relevance and leadership in high performance computing**
- **Many new areas of investigation are in the biological and biomedical sciences**
 - **Adapting to this may require reaching out to new partners with different needs/strengths than physical scientists and engineers who are traditional collaborators**
 - **Building strength in this area may require forging an alliance with the Academic Health Center**
- **MSI is currently providing support to biomedical scientists, and this could be leveraged to open doors to research projects**
 - **Biomedical image analysis and diagnostics**
 - **Systems Biology and cellular simulations**
 - **Data integration across domains in the biomedical sciences**
- **Traditional areas of excellence (e.g. collaboration with scientists in the Institute of Technology) need to be maintained**

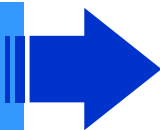
Industrial Outreach

- **MSI should aggressively pursue strong relationships with the local computing industry**
 - Strongest supercomputing industrial base in the world: IBM, Cray, SGI, Sun, Unisys, Lockheed Martin, GD-AIS, ...
 - Strongest high performance storage industrial base in the world: Hitachi, Sun, Quantum (StorNext), Seagate (Imation), Isolon, Panassas, Instrumental, ...
- **Relationship should focus on**
 - Research on core technology bottlenecks
 - New applications (e.g. medical imaging) for current technologies
 - Testing of Serial #1 products at large scale
 - Creating an academic community around these industries to support their needs: conferences, student internships, ...
 - The approach should be “Here is what we can do for you.”
- **Such relationships can eventually lead to external funding opportunities**
 - Helps set national research agenda
 - Strong research transition partners are important to (many) government funding organizations

National Security Outreach

- **Supercomputing and National Security are highly intertwined**
 - National security organizations are major funders of core supercomputing research
 - A primary goal of these organizations is to develop the national security workforce
 - Computational approaches are essential in developing security applications in the information age
- **National security organizations have significantly decreased their funding of Universities**
 - Largely foreign graduate student population
 - Increase in the need for workers who can obtain the highest level of clearance (typically natural born citizens whose immediate families are in the U.S.)
 - This also effects local industry who do a sizable business with national security organizations
- **A University that can create a strategy for increasing its U.S. Citizen graduate population will be have significant advantage**
 - Assume no citizenship constraints can be placed on RAs
 - Suggest MSI establish fellowships specifically targeting top Minnesota high school students who have left the state for their undergraduate education
 - Increases U.S. citizen grad student population
 - Decreases brain drain from the state

Outline


- Introduction
- MSI Today
- Research Mission
- **Organization** 
 - *Reporting Lines*
 - *Investment Levels*
- Operations
- Summary

Organization

- **Formal MSI mission should recognize its existing (expanded) role**
 - Ongoing support for supercomputing users
 - Support of research computing
 - This includes all computing technologies beyond the desktop necessary for research (e.g. current biomedical support)
 - Creating an on-ramp to data and computing intensive research
 - Collaborative infrastructure (inter and intra)
- **MSI (perhaps renamed “Minnesota Research Computing”) should be a U of M wide organization reporting directly to the VP of Research with joint investment from AHC (as this is MSI’s major growth area)**
 - DTC should be merged into this entity and perhaps become a research and industrial outreach arm
- **MSI should conduct a national search for an External Director**
 - Director should be given budgetary and policy authority
- **Current 16 committee structure should be reduced to one**
- **MSI budget is ~1% of U of M research budget, this is too low**
 - Actual budget should be tied to the research agenda and objectives

Outline

- Introduction
- MSI Today
- Research Mission
- Organization
- **Operations**
- Summary



- *Missions*
- *Operations*
- *Research*
- *Infrastructure*

MSI Operation Mission

- **MSI mission should broaden to include all of research computing beyond what OIT can provide**
- **User services**
 - **Software tools consulting**
 - **Software development (web, database, ...)**
 - **Access to licensed and public domain software**
- **Collaboration services**
 - **Common repositories, web services, ...**
 - **Should be offered as a part of all industrial collaborations**
 - **Should be offered to IMA visitors as a way to continue collaborations after visitors have left**
 - **Should write collaborative proposals with researchers**
- **All of these services should be provided in such a way that users can easily take advantage of MSI's supercomputing capabilities**
 - **Allows researchers to operate on larger data sets**
 - **Creates the next set of "big iron" users**

MSI Operations: Best Practices

- **MSI operation can be made significantly more efficient by adopting current best practices the field**
 - Alignment of services, use cases and policies to best meet user's needs will *significantly* decrease friction in system
 - Team leads should visit the leading supercomputing centers to “steal” current best practices, e.g. integrated incident tracking systems
 - Negative user experiences should be analyze by team to incrementally evolve practices and policies to eliminate these experiences
 - In nearly all cases these experiences can be eliminated when the users precise needs are understood in detail by the entire team
- **MSI needs a clear model of its services**
 - Clear description of its services to different classes of users
 - Clear set of use cases for satisfying different user needs
- **16 Committees should be replaced with one advisory board**
 - User's primary contact with MSI should be via well defined interactions (not committees or out-of-band channels)
 - MSI Staff take care of users, not committees
- **MSI policies are a tool by which the director improves service to users**
 - Policy and exception handling must under the direct control of the director

MSI Operations: Research Support

- **Develop a "Return On Investment" model to optimize dollar allocations to the most effective activities**
 - use queue logs to determine best cycle allocation policies
 - Identify promising research teams for indepth collaborations
- **Model should focus on parts of process MSI can influence**
 - Effort required to get code running on supercomputer
 - End performance of code on supercomputer
 - Strategic benefits connected to university goals
- **Identify the current desktop environments that researchers use most often (e.g. Perl, Matlab, IDL ...)**
 - Provide simple mechanisms that allow a large fraction of researchers to use supercomputing resources from their desktop to do embarrassingly parallel computations
 - Will dramatically increase the effectiveness of these researchers
 - Will dramatically increase time spent by staff to support these users

Physical Infrastructure

- **Physical co-location facility should be built to allow significantly larger systems to be acquired**
- **Facility should be planned to 2-3 Megawatt capacity**
 - These are typically integrated with a “green energy” plan
 - AHPARC site is an ideal location
- **Facility should also host smaller systems acquired by PIs (at no or minimal charge)**
- **Should consider a longer buy cycle (~3 years) to allow bigger buys and non-linear const benefits**
 - If possible, hardware funds should be “escrowed” to maximize “cash in hand effect”
 - For example, it is not unusual to receive one time offers to buy (on short notice) \$15M of hardware for \$5M
 - Opportunities are particularly plentiful after major award selections and the losing vendors are looking for another place to allocate their investment

Summary

- **Enhancing the research reputation of the U of M is essential to achieving its long term standing goal**
 - **Within this context, the MSI is the only University wide organization which can substantially impact the research of every major department**
- **MSI plays a key role in enabling collaborations across the University and with outside organizations**
- **In the last decade MSI has fallen from a Top 5 to a Top 25 academic supercomputing center**
 - **Mostly due to the relatively large improvements made by peers**
- **U of M must identify a *specific* core research agenda which exploits its organic advantages over its peers**
 - **Engage in setting the national research agenda**
 - **This entails taking some risks**
- **MSI mission should be aligned to this research agenda**
- **Formal MSI mission should recognize its existing expanded role**
- **MSI (perhaps renamed “Minnesota Research Computing”) should be a U of M wide organization reporting directly to the VP of Research with joint investment from AHC (as this is MSI’s major growth area)**
- **MSI should conduct a national search for an External Director**
- **Current 16 committee structure should be reduced to one**
- **MSI budget is ~1% of U of M research budget, this is too low**