Why provide high performance computing (HPC) services?

- Computation is increasingly important for research in and practice of many disciplines
- In spite of increases in performance of desktop computers, many problems require more compute capability to be solved in reasonable length of time (or at all)
- Computational capability is important for all universities, but is particularly critical resource for STEM intensive universities – for support of both research and instruction
- Increase efficiency of campus research computing investments – pool resources to create a more significant capability
- Increase research group productivity – allow research groups to focus on their research instead of dealing with distraction of operating their own HPC infrastructure
- Attract and retain computational researchers
- Increase university grant funding – by having computational capability and productive computational researchers
- Graduate students skilled in use of computational methods

What services are currently offered?

- In 2003 NC State decided to provide intermediate level of HPC services (agreement between Vice Provost for IT, Chancellor, Provost, and Vice Chancellor for Research as NC Supercomputing Center was closing)
- This intermediate level is intended to fill gap between computational capability of desktop and resources available at national supercomputing centers
- Currently there are two HPC platforms available from OIT for campus use:
    - a shared memory POWER5 system that supports jobs using up to 8 processors and 32GB of shared memory (the system has total of 32 processors for running jobs)
    - a distributed memory Linux cluster that supports jobs using up to 128 processors and 256GB of distributed memory (the system has total of about 1500 processors)
- In addition to compute cycles, both shared and distributed memory compute services also include
    - Storage
        - Home directory – backed up space for scripts and source code
        - Scratch – not backed up space for input and output data sets during code execution and data analysis
        - Mass Storage – backed up, hierarchically managed file system for large data files when not in active use
    - Storage software (Tivoli Storage Manager – handles file system backup and mass storage HSM)

- System software (eg operating system, compilers, debuggers, resource manager and scheduler, …)
- Application software
- User Support (answering basic questions about resource use and resolving problems with resources)
- Training (periodic – about 3 per year – few day classes on use of HPC resources and parallel programming)
- Consulting (more in-depth assistance with application/algorithm selection, code porting, code optimization, etc.)
- Collaboration (extensive involvement with project using HPC resources – usually funded as part of a grant)

Who is able to use the services?

- OIT HPC resources are available to any NC State faculty member to support their research or instruction
- A project is established for each faculty member using HPC resources. Associated with their project they can add access for any other active unity ids
- Faculty members with HPC projects are asked to renew annually by completing an online report that collects information about how the resources were used during the previous year and how many publications, grants, and student projects were related to use of the resources

How have the services been funded?

- In 2003 a budget of $1.5M was requested as minimum funding level to provide HPC services (as a point of reference NCSC had been operating with a budget of more than $5M with NC State using about 60% of the NCSC resources). HPC funding of $800K was provided ($110K from an existing open ITD position and $345K each from the Provost Office and Research Office)
- To leverage the base HPC funding a partner's program was developed to allow faculty members to purchase compatible hardware and add it to the central HPC resources. Partners have unlimited, priority access to the resources they have purchased. But when they are not using their compute resources the resources are available to other users
- HPC has also been able to utilize IBM's Shared University Research grant program to significantly augment the available compute resources
- Over time, legislative increases have increased the (now OIT) continuing funding to about $150K for FY08-09, the Provost Office funding has remained consistent - $350K for FY08-09, and the Research Office funding has declined ranging from $120K-$180K the last three fiscal years – anticipating $180K for FY08-09
- Partner program participation has exceeded $500K each of the past two fiscal years

- See attached chart showing HPC expenditures by funding source since program initiation in July 2003
    - note that some Research Office funds were carried over from FY07 to FY08
    - and that in FY04 there was additional funding from the Research Office for HPC participation in the university storage management system purchased that year
    - IBM SUR grant in FY08 was utilized entirely to support VCL with upfit of donated Intel blades – but these are expected to provide significant cycles for HPC use even though their primary purpose is to support VCL

Who are current users and partners?

- For calendar year 2008 there were 360 unique logins that utilized HPC compute services and 107 unique projects
- These users represented 25 departments and 6 colleges
- There are currently 28 faculty partners from 15 departments and 4 colleges
- See attached chart for breakdown of CPU hours and storage used by department and college

What are immediate issues for the committee?

- Research Computing Committee
    - Membership
        - Are there other areas that should be represented? If so, what are they?
    - Operational policies and procedures
        - Selection of Chair
        - Terms, reappointment
        - Meeting frequency
        - Communication with CIO
        - Communication with constituencies
- HPC Strategy
    - Is the HPC strategy developed in 2003 – the provision of a central, intermediate level HPC capability – still appropriate? If not, what is appropriate HPC strategy for NC State?
- Funding Model
    - Is there more effective model for HPC base funding
    - Are there opportunities for additional funding
        - Charge back for some services or services beyond some threshold
        - Infrastructure grants
        - …
- Data Storage

    - How should computational research data be managed (and how should that be funded)?
    - What policies should we implement to encourage the desired data management?
- Resource Management and Scheduling
    - Are there other factors that should be considered in management and scheduling of limited HPC resources.. eg
        - University strategic thrust areas – should these have higher priority for cycles, scratch space, consulting, application porting/licensing
        - Faculty hiring/retention – are there specific faculty members who department heads or deans would like to have higher priority for cycles, scratch space, consulting, application porting/licensing
- Unmet needs
    - What are the HPC needs not being met by current services?
    - What is relative priority of possible services to address the unmet needs versus existing services (should the offered services be changed)?

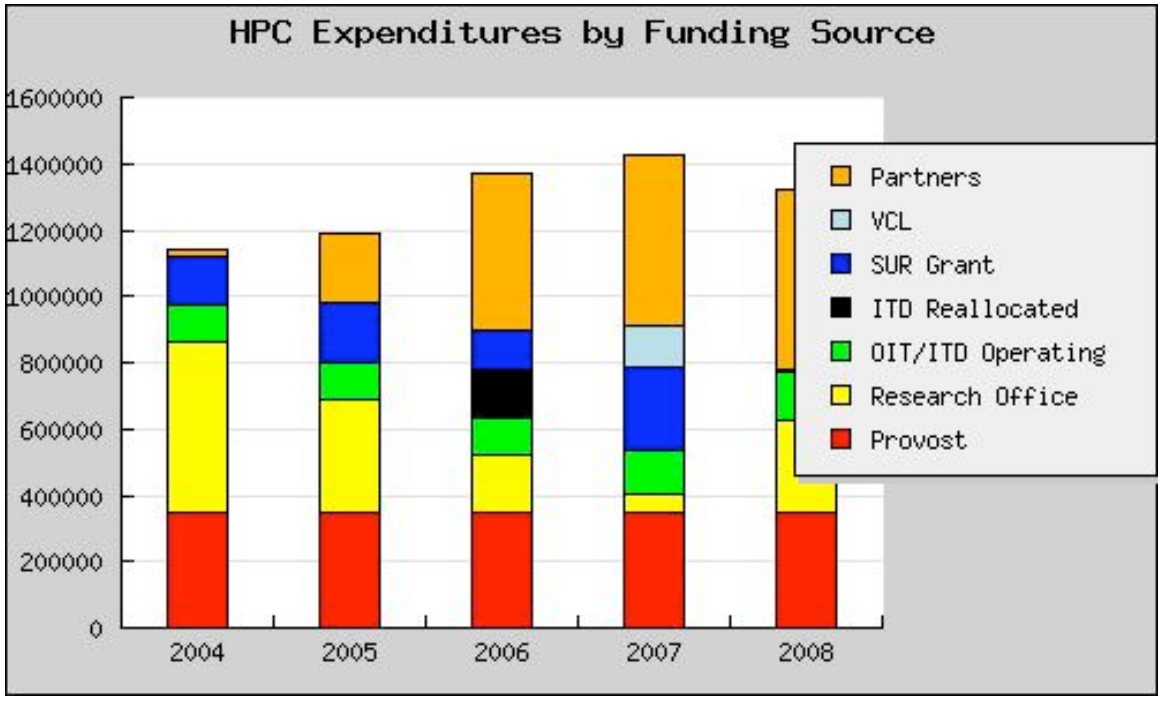
HPC Expenditures by Funding Source
1600000
1400000
1200000
1000000
800000
600000
400000
200000
0
2004
2005
2006
2007
2008
Partners
VCL
SUR Grant
ITD Reallocated
OIT/ITD Operating
Research Office
Provost

Distribution by College of HPC Compute Hours Utilized

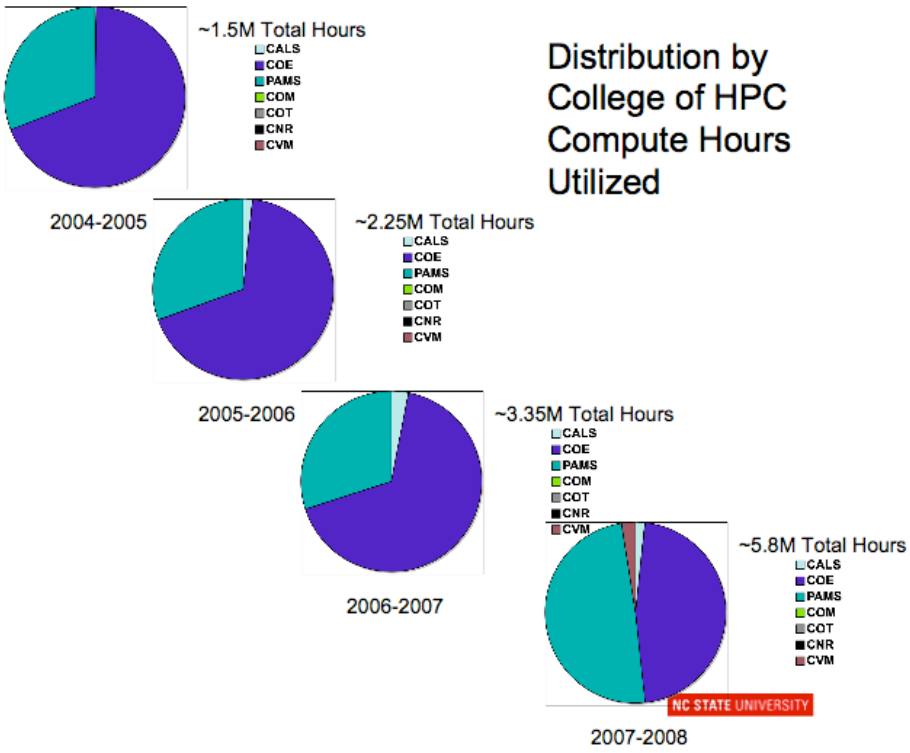

Distribution by College of HPC Mass Storage Utilization

~14 Total TB

June 2006

~16 Total TB

May 2007

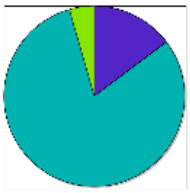

~26 Total TB

June 2008

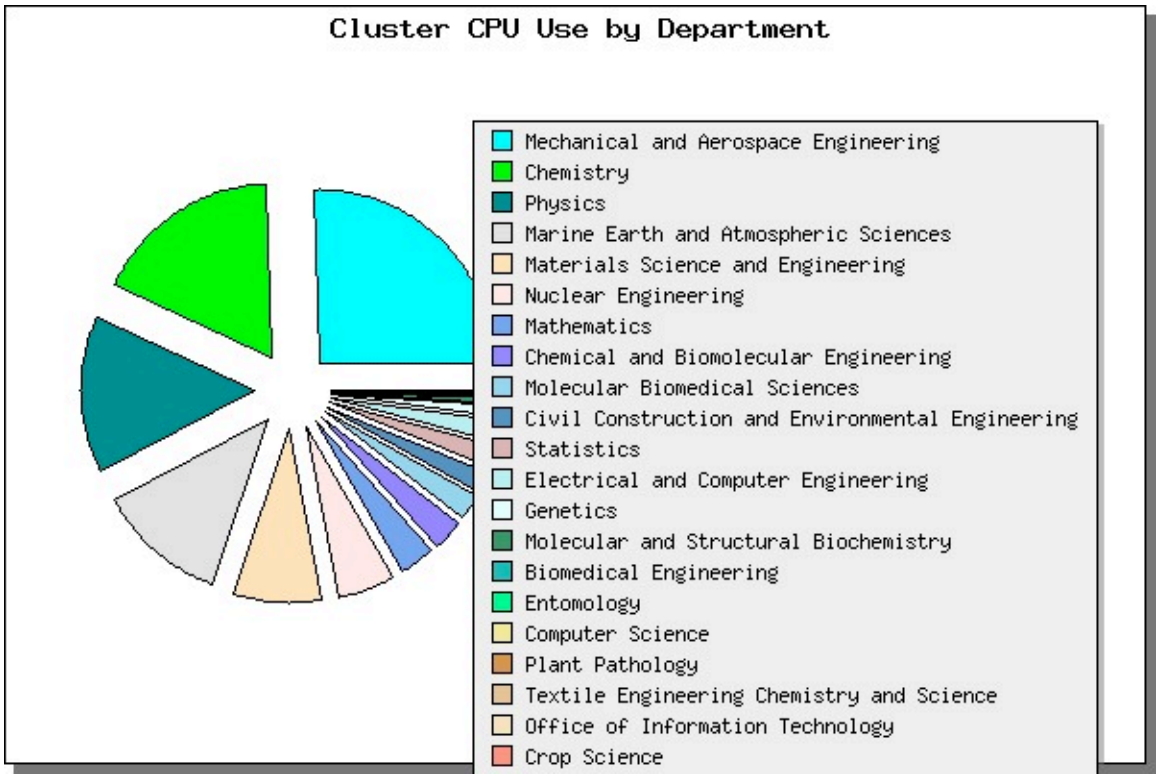
Cluster CPU Use by Department
Mechanical and Aerospace Engineering
Chemistry
Physics
Marine Earth and Atmospheric Sciences
Materials Science and Engineering
Nuclear Engineering
Mathematics
Chemical and Biomolecular Engineering
Molecular Biomedical Sciences
Civil Construction and Environmental Engineering
Statistics
Electrical and Computer Engineering
Genetics
Molecular and Structural Biochemistry
Biomedical Engineering
Entomology
Computer Science
Plant Pathology
Textile Engineering Chemistry and Science
Office of Information Technology
Crop Science