

Managing Expectations and Understanding Needs for Campus Level Shared Research Computing

Jan Odegard, Ken Kennedy Institute
for Information Technology, Rice

Lee Panetta, Department of
Atmospheric Sciences, TAMU

Escalating Concerns

- Administration & faculty recognize that “having” is a competitive advantage, but nature and needs of computational science not always recognized at **highest** levels at our institutions
 - (not even a familiar domain for most CIOs / VPs for IT / ...)
- Campus Information Technology Divisions are struggling with what to do or how to support research computing
- Where does campus IT leadership get **advice?**
- Trend challenges traditional mission of campus IT organizations
- Supporting research computing (staff and infrastructure) is a bottomless pit
 - Where and how to draw the line and how to expense cost?
 - Escalating demand on data centers (power/cooling cost)
- How do we support and best leverage national resources?

Regional/National Discussions

- EDUCAUSE
 - Campus Cyberinfrastructure (CCI) Working Group
- TeraGrid Future (<http://www.teragridfuture.org>)
- An Alternative Approach to NSF Funding of HPC
<http://www.hpcwire.com/hpc/2059714.html>
- LEARN
 - Technology Advisory Group (TAG)
 - Research Advisory Group (RAG)
- CASC
 - 50+ academic research institutions discussion
 - CASC is dedicated to advocating the use of the most advanced computing technology to accelerate scientific discovery for national competitiveness, global security, and economic success, as well as develop a diverse and well prepared 21st century workforce
- What is the role of campus level research computing?
- Should this be the responsibility of the campus IT organization?
 - Recall NSF used to fund mainframes and networking

Understanding Users and their Needs

- Understanding a diverse workload is both time consuming and difficult but necessary
- Allocation issues
 - “Users have unbounded needs, centers have bounded budgets”
 - funded vs unfunded, judgment of relative scientific value in general
 - large scale tightly coupled versus large volume loosely coupled (clouds)
- Tradeoffs:
 - “response time” vs. “efficient” use of machines
 - “sophisticated” computational projects vs less sophisticated but valuable science/engineering
- Users often lack understanding of
 - total cost of ownership
 - competing needs generated by supporting diverse work loads (wall time versus wait time)
- Users tend to advocate the path of least resistance
- Does scarcity encourage innovation and more attention to code efficiency?
- Cluster huggers

HiPCAT Community Opportunity

- Each institution are working hard to support their respective faculty and researchers
- Each institution have unique needs and goals
- Each have lessons learned and best practices
- It would be beneficial to all to share these lessons and learn from each other even if it does not mean sharing resources
- HiPCAT is the perfect community for these discussions
- Topics for discussion:
 - Users profiling and support
 - Tools and support
 - Policies and support
 - Management
 - ...

The Money Pit?