

Brainstorming Thoughts

Guy Almes <galmes@nsf.gov>
Office of Cyberinfrastructure



CCRW Workshop
Arlington, Virginia
April 2006



What does this mean for the Campus?

- Attention to elements:
 - excellent networking
 - access to high-performance computing
 - facilitation of use of national resources
- Raise awareness of the opportunities and challenges presented to science and engineering users --
Cyberinfrastructure is holistic



Excellent Networking

- Network performance should be good everywhere, and possible to make excellent as needed
- Encourage excellent wide-area resources
- But most performance problems are due to flaws in the last 100 feet



When it's done badly ...

- Inadequate provisioning within campus backbones or department networks
- Inadequate or ill-configured switches
- Weak host hardware or software
- Lack of connection between CIO shop and motivating users



An Emerging Model

- As an element of network planning, be able to extend 'excellent performance' where needed
- gigE is becoming ordinary and 10gigE is no longer heroic
- Know the users and their needs



Access to Advanced Computing

- Capability systems are national
 - facilitate use
- Capacity systems can be local
 - ordinary, but large, clusters of commodity components, integrated with appropriate software and operated professionally



When it's done badly ...

- Small clusters funded across campus
- PI procures, configures, operates (with enthusiasm tapering off)
- Sysadmin often by grad students
- Low utilization due to bursty needs
- Security weaknesses
- Strain on departmental power, air conditioning, and space

An Emerging Model

- Clusters are funded from many sources
- Condominium clusters hosted by computer center
 - power, air conditioning, space
 - professional sysadmin support
 - PI assured of 'his/her' nodes
 - but usually gets much more when needed
 - limited, but valued, access provided to others on campus
- Requires technical and nontechnical planning





Oh, and security too

- Understand the tension between security and the needs of an open research community
- Flexibility, security, and performance need to be balanced



Software is key

- Operating systems for servers and clusters
 - TCP stacks and performant file systems
 - Effective authentication, authorization, and accounting support
 - Grid software stacks to encourage sharing and flexible access to non-local resources
- Applications support



And take a holistic approach

- The key to Cyberinfrastructure is the combining of advanced computing, networking, middleware, etc., in support of the science/engineering user!
- Enabling science is the goal



Relating to NSF

- Consider those top-ten users
- Foster mutual accountability among:
 - PIs
 - That person's NSF POs
 - Campus network organization
 - (Maybe) Sponsored Research Office