**NSF IESP Whitepaper**

**Abani Patra, Rob Pennington, Ed Seidel**

**Office of Cyberinfrastructure**

**National Science Foundation**

Within the context of the "NSF's Vision for Cyberinfrastructure for 21st Century Discovery" document, NSF is developing a comprehensive program for supporting the national cyberinfrastructure (CI) for science and engineering, including major HPC facilities, grids, networks, software, data, and virtual organizations. NSF clearly cannot do this alone, and therefore must pursue global partnerships with other organizations and agencies.

NSF reaches deeply into every campus in the US, covers all the sciences and engineering areas, and in terms of cyberinfrastructure which includes HPC, is very broad. NSF researchers will clearly benefit from a stronger software program, improved support for complex applications and strengthened integration with campuses. Students and postdocs will benefit from training in software engineering, use of advanced CI, and socio-technical activities that are critical to success in many complex research activities.

The computational community is already dealing with several major challenges at petascale, including new hardware using manycore, massive scaling, system software, file systems, applications software, debuggers, applications development, programming environments, machine rooms, cooling and power costs.

Exascale challenges will drive innovation in many CI related areas. Developments in cyberinfrastructure to support scientific and engineering research will need to be integrated across the following major topics:

- Software: A major software grand challenge program responsive to emerging architectures needs to be developed, involving national and international efforts.
- Applications: NSF funded researchers have strength and breadth in the community that will use exascale facilities. New research challenges will further broaden the application coverage.
- Hardware: R&D activities in hardware design that are responsive to the most challenging application needs.

Questions for consideration:

The NSF cyberinfrastructure vision document provides the current high level framework for cyberinstructure strategy. The requirements for cyberinfrastructure are evolving rapidly and, as a result, new questions arise in planning for future cyberinfrastructure. As part of the process of understanding these requirements, we welcome discussion and input on a wide range of questions including the following.

- How will present & emerging applications use exascale systems?

- What are the new applications that are emerging or likely to emerge in the coming decade?
    - Are they new application domains, new modeling modalities, multimodal modeling, dynamic/on-line integration computation and measurements?
    - How will technology advances drive the advancement of applications capabilities (technology-push)?
- How can NSF best stimulate development of exascale software applications?
- How can application needs drive the design of hardware platforms, system software, and applications software development environments?
- How will new architectures aid or impede successful reformulation of problems for parallel solution approaches?
- How can useful software that has been developed as part of the exascale effort be sustained beyond the development period?
- What systems software will be required? Distributed systems support, programming environments, runtime support, data-management user tools?
- In what ways will fault tolerance need to be considered by the applications developers? By the system software developers?
- What application support environments will be needed? Application packages, numeric and non-numeric library packages, problem-solving environments?
- How can NSF aid seamless portability of applications across different hardware and software platforms as they all evolve?
- How can NSF aid or catalyze developments that make it possible to provide the same user experience and where possible use the same tools, including compilers, debuggers and performance tools, on system scales all the way down to the typical researcher's laptop or desktop?
- How can the community of science and engineering researchers who will use exascale systems be best supported in a rapidly changing environment?
- How feasible is the development of generally applicable software that will enable efficient translation of problems to programs? What priority should be given to pursuing this approach?
- What education and training actions should be considered to prepare researchers, students and educators for future cyberinfrastructure?