NASA SBIR 2011 Phase I Solicitation

S6.01 Technologies for Large-Scale Numerical Simulation

Lead Center: ARC

Participating Center(s): GSFC

NASA scientists and engineers are increasingly turning to large-scale numerical simulation on supercomputers to advance understanding of complex Earth and astrophysical systems, and to conduct high-fidelity aerospace engineering analyses. The goal of this subtopic is to increase the mission impact of NASA's investments in supercomputing systems and associated operations and services. Specific objectives are to:

- Decrease the barriers to entry for prospective supercomputing users.

- Minimize the supercomputer user's total time-to-solution (e.g., time to discover, understand, predict, or design).

- Increase the achievable scale and complexity of computational analysis, data ingest, and data communications.

- Reduce the cost of providing a given level of supercomputing performance on NASA applications.

- Enhance the efficiency and effectiveness of NASA's supercomputing operations and services.

Expected outcomes are to improve the productivity of NASA's supercomputing users, broaden NASA's supercomputing user base, accelerate advancement of NASA science and engineering, and benefit the supercomputing community through dissemination of operational best practices.

The approach of this subtopic is to seek novel software and hardware technologies that provide notable benefits to NASA's supercomputing users and facilities, and to infuse these technologies into NASA supercomputing operations. Successful technology development efforts under this subtopic would be considered for follow-on funding by, and infusion into, NASA's high-end computing (HEC) projects: the High End Computing Capability project at Ames and the Scientific Computing project at Goddard. To assure maximum relevance to NASA, funded SBIR contracts under this subtopic should engage in direct interactions with one or both HEC projects, and with key HEC users where appropriate. Research should be conducted to demonstrate technical feasibility and NASA relevance during Phase I and show a path toward a Phase II prototype demonstration.

Offerors should demonstrate awareness of the state-of-the-art of their proposed technology, and should leverage existing commercial capabilities and research efforts where appropriate. Open Source software and open standards are strongly preferred. Note that the NASA supercomputing environment is characterized by: HEC systems operating behind a firewall to meet strict IT security requirements, communication-intensive applications, massive computations requiring high concurrency, complex computational workflows and immense datasets, and the need to support hundreds of complex application codes - many of which are frequently updated by the user/developer. As a result, solutions that involve the following must clearly explain how they would work in the NASA environment: Grid computing, web services, client-server models, embarrassingly parallel computations, and technologies that require significant application re-engineering. Projects need not benefit all NASA HEC users or application codes, but demonstrating applicability to an important NASA discipline, or even a key NASA application code, could provide significant value.

Specific technology areas of interest:

Efficient Computing

In spite of the rapidly increasing capability and efficiency of supercomputers, NASA's HEC facilities cannot purchase, power, and cool sufficient HEC resources to satisfy all user demands. This subtopic element seeks dramatically more efficient and effective supercomputing approaches in terms of their ability to supply increased HEC capability or capacity per dollar and/or per Watt for real NASA applications. Examples include:

- Novel computational accelerators and architectures.

- Enhanced visualization technologies.

- Improved algorithms for key codes.

- Power-aware "Green" computing technologies and techniques.

- Systems (including both hardware and software) for data-intensive computing.

- Approaches to effectively manage and utilize many-core processors including algorithmic changes, compiler techniques and runtime systems.

User Productivity Environments

The user interface to a supercomputer is typically a command line in a text window. This subtopic element seeks more intuitive, intelligent, user-customizable, and integrated interfaces to supercomputing resources, enabling users to more completely leverage the power of HEC to increase their productivity. Such an interface could enhance many essential supercomputing tasks: accessing and managing resources, training, getting services, developing codes (e.g., debugging and performance analysis), running computations, managing files and data, analyzing and visualizing results, transmitting data, collaborating, etc.

Cloud Supercomputing

Cloud computing has made tremendous promises, and demonstrated some success, for business computing. For operations, potential benefits include: resource virtualization, incremental and transparent provisioning, enhanced

resource consolidation and utilization, automated resource management, automated job migration, and increased service availability, and others. For users, potential benefits include: out-sourced operations, on-demand resource availability, increased service reliability, customized software environments, a web user interface, and more. This subtopic element seeks technologies that enable Cloud computing to be used for efficient and effective supercomputing operations and services.