



NASA SBIR 2017 Phase I Solicitation

S5.01 Technologies for Large-Scale Numerical Simulation

Lead Center: ARC

Participating Center(s): GSFC

Technology Area: TA11 Modeling, Simulation, Information Technology and Processing

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 for 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.
 - The need to support hundreds of complex application codes - many of which are frequently updated by the user/developer.

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. For instance, a GPU accelerated (or multi-core) planetary accretion code such as LIPAD (Lagrangian Integrator for Planetary Accretion and Dynamics).

The three main technology areas of S5.01 are aligned with three objectives of NSCI, the National Strategic Computing Initiative, announced by the White House in July 2015. The overarching goal of NSCI is to coordinate and accelerate U.S. activities in HEC, including hardware, software, and workforce development, so that the U.S. remains the world leader in HEC technology and application. NSCI charges every agency that is a significant user of HEC to make a significant contribution to this goal. This SBIR subtopic is an important part of NASA's contribution to NSCI. See <https://www.nitrd.gov/nsci/index.aspx> for more information about NSCI. The three main elements of S5.01 are:

- Many NASA science applications demand much faster supercomputers. This area seeks technologies to accelerate the development of an efficient and practical exascale computing system (10¹⁸ operations per second). Innovative file systems that leverage node memory and a new exascale operating system geared toward NASA applications are two possible technologies for this element. At the same time, this area calls for technology to support co-design (i.e., concurrent design) of NASA applications and exascale supercomputers, enabling application scaling to billion-fold parallelism while dramatically increasing memory access efficiency. This supports NSCI Objective 1 (Accelerating delivery of a capable exascale computing system that integrates hardware and software capability to deliver approximately 100 times the performance of current 10 petaflop systems across a range of applications representing government needs.).
- Data analytics is becoming a bigger part of the supercomputing workload, as computed and measured data expand dramatically, and the need grows to rapidly utilize and understand that data. This area calls for technologies that support convergence of computing systems optimized for modeling & simulation and those optimized for data analytics (e.g., data assimilation, data compression, image analysis, machine learning, visualization, and data mining). In-situ data analytics that can run in-memory side-by-side with the model run is another possible technology for this element. This supports NSCI Objective 2 (Increasing coherence between the technology base used for modeling and simulation and that used for data analytic computing.).
- Presently it is difficult to integrate cyberinfrastructure elements (supercomputing system, data stores, distributed teams, instruments, mobile devices, etc.) into an efficient and productive science environment. This area seeks technologies to make elements of the supercomputing ecosystem much more accessible and composable, while maintaining security. Thus supports NSCI Objective 4 (Increasing the capacity and capability of an enduring national HPC ecosystem by employing a holistic approach that addresses relevant factors such as networking technology, workflow, downward scaling, foundational algorithms and software, accessibility, and workforce development.).