

Computer Science Research needed for OSG to be successful in 2020.

OSG Executive Team, April 2012

Today, the Open Science Grid offers a diverse fabric of services that advances scientific discovery through more powerful and effective computing. At OSG's core is the concept of Distributed High-Throughput Computing (DHTC), the shared utilization of large ensembles of autonomous resources toward a common goal, where all the elements are optimized for maximizing computational throughput.

The Open Science Grid is both a virtual facility and an ecosystem. As a facility, it provides deployment and testing opportunities for Computer Science research. As an ecosystem OSG facilitates or even instigates collaborations between Computer Science (CS) researchers and domain scientists, creating interdisciplinary communities that use invocations in computing methodologies and technologies to address domain science challenges.

OSG's existing capabilities are effective but basic and primitive compared to those necessary to support collaborative extreme-scale science in the 21st century. OSG needs to achieve new levels of usability for portable, transparent services across an increasingly diverse and complex set of resource types serving a broader range of science domains with growing diversity of scientific computing skills. Statically federated resources (our classical view of the Grid) need to be integrated with dynamically allocated resources (like clouds) causing new challenges for resource planning, acquisition and provisioning. The ongoing increases in size, complexity and diversity of both the science applications and the computer technologies present significant research and technical challenges.

Moving forward OSG faces foundational changes in technologies, scales of use and collaborative heterogeneity. Computing hardware is undergoing a foundational change in the number of compute cores available on a single node; The emergence and prevalence of mobile computing interfaces causes foundational changes in the interaction of individual researchers with the computing environment; The size and reach of data transport and access are undergoing foundational changes in throughput, connectivity and heterogeneous sharing; More scientists use high level tools to develop complex applications out of community specific building blocks. The nature and diversity of the collaborative research groups are undergoing foundational changes in the breadth and inter-dependence of the research questions being addressed and computational tools employed.

Therefore, for OSG to address these challenges it needs an influx of innovative frameworks and technologies in the areas of data, security, systems, workflows, tools and collaborative environments, working on issues like transparent usage, provisioning and management of resources, maximizing throughput and total benefit, improving robustness, managing identity information and trust, improving

usability and integration. Innovation in these broad but inter-related areas can be only accomplished through a coordinated and collaborative CS research effort.

The CS research needs to be conducted in the context of the domain sciences it will enable, and it needs to be driven by specific challenges. OSG will be considered successful if we not only keep up with the growing LHC needs but also expand to enhance the computing throughput to a broad spectrum of scientists at a variety of scales, from individual users at a single campus to multi-institutional experiments.

Among our scientific user community, we see presently two extremes of such challenges. On one extreme, there are the large collaborations in HEP that need CS innovation to help the transition from petabytes to exabytes of data volumes, and from 10^5 cores to 10^8 cores within the global DHTC systems over the next ten years. On the other extreme, there are an ever increasing number and diversity of small to mid-size collaborations of domain scientists that struggle with crossing order of magnitude boundaries in their computing — transitioning from gigabytes to terabytes of data volumes that are an integral part of scientific DHTC workflows at scales of 1-10 million CPU hours. Over the next 10 years, we expect the number and diversity of groups at this end of the spectrum to increase as data and compute intensive science becomes the norm rather than the exception.

In our assessment, it is crucial to prevent the capability gap between these two extremes from growing. While exascale problems need to be solved for the large collaborations, CS research is equally necessary to ensure solutions are available for the many scientists challenged at the terascale and petascale on the passage to the exascale.

To help our growing spectrum of users, we need to keep up with increasingly dynamic and heterogeneous environments while ensuring that domain scientists with limited computing expertise can use those environments. Collaborative tools are also important to help us bridge the gap.

We list a high-level summary of the research issues that need to be addressed to keep OSG successful for the rest of this decade. The OSG will work out more detailed specification of the different areas working with the CS community and OSG stakeholders as it develops its vision and strategy in the coming months and years.

- 1) Transparent usage and ad-hoc/on-demand scheduling of heterogeneous resources¹
 - a. Discovery and provisioning of resources, including compute, storage, network.
 - b. Software deployment and integration.
 - c. Location independence and transparency.

¹ Resources include compute, storage, network, memory, interconnects. Resource types include sensors, mobile devices, laptops, loosely-coupled clusters, clouds, High-Performance Computing systems etc.

- d. On-demand planning and acquisition of resources.
 - e. On-demand deployment of managed storage.
 - f. Transparent access to dynamically instantiated resources.
- 2) Life-time and end-to-end data management and access across globally distributed and heterogeneous storage resources.
- a. Content-managed, global distribution and low-latency transparent access to vast multi-source data sets.
 - b. Management of dataset namespaces and metadata.
 - c. Implementation of data placement, especially for output from multiple resource types.
 - d. Adaptation of data management, storage, and access mechanisms across all scales from mobile to exascale resources, and from low-bandwidth to extreme-bandwidth connections.
 - e. Efficient management of and access to large aggregations of small, diverse, dynamic, data-types.
 - f. Long-term curation of data and software, for maintenance and reuse.
- 3) Maximization of total benefit obtained from the system
- a. Policy based management of the ensemble of resources.
 - b. Co-scheduling of multiple resources and resource types in complex topologies for heterogeneous, complex workloads.
 - c. Coordination between workflow and data management systems, data driven processing, coordinating the on-demand provisioning of network resource with data access systems.
 - d. Adaptation of resource policies and usage based on measured value, total cost of ownership models².
 - e. Monitoring and accounting.
- 4) Solution Robustness, Quality and Assessment
- a. End-to-end system error propagation and failure handling.
 - b. Automated software and solution testing, validation and assessment.
 - c. Reproducibility and trustworthiness of results.
- 5) Adaptable and flexible use of dynamic many-core computing hardware
- a. Architectural hiding at the application layer.
 - b. Automated adaptation to match number of cores available, and dynamic changes to the core count during runtime.
 - c. Maximize resource usage efficiencies in heterogeneous many-core environments.

² Where costs include hardware, software, maintenance, power, cooling, administration, security, replacements, levels of service, impact of MTBF etc.

- 6) Distributed identity and access control frameworks for sharing and use of information and resources
 - a. Single sign-on, transparency and flexibility in use of heterogeneous identity management systems.
 - b. Management and application of dynamic, short- to long-lived collaboration based access control.
 - c. Models and methods for trust and risk management to predict the response and guide the protection of systems that include diverse levels of trust.
 - d. Models that encompass all the entities in a distributed system, including identity providers (federated and social), users, virtual organizations, central service providers, resource owners and sites.

- 7) User level usability, flexibility and integration
 - a. Provisioning of dynamic, integrated end-to-end environments for shared-use.
 - b. Management and monitoring of environment attributes.
 - c. High level, tailored workflow definitions and engines.