

Proposal to the OSG Council on Open Facility Support - DRAFT

Introduction

Presently the OSG consists of over 100 sites delivering over 700 million wall hours of distributed high throughput computing (DHTC) annually. A number of virtual organizations (VOs) operate primarily on sites not owned by the VO — opportunistically. In recent years the share of opportunistic utilization on the OSG has increased to over 10%, yielding over 100 million wall hours annually. These opportunistic cycles offer DHTC access to a wide variety of US researchers enabling scientific research and innovation. The OSG Project has been actively working to increase: 1) access to DHTC for US Researchers and 2) availability of these opportunistic cycles at all levels of the OSG fabric.

Proposed Principle

The OSG is a sharing ecosystem where members benefit from the contributions of other consortium members. One crucial component of this ecosystem is the compute cycles provided at sites. Compute cycles at any particular site should not remain unused if there is computational work pending for another OSG member; whenever possible, it is our goal to match queued work with available compute capacity at sites.

Recommendations

Site Configuration

We recommend that sites approach the availability and delivery of computing cycles with three tiers of usage in mind; the stated goal being that all available computing cycles are used by one of the three tiers with none remaining idle when technically feasible.

1. The first tier with highest priority consists of VOs that have ownership or partial ownership of the CPU resources at the site as well as local users who submit work directly into the local batch system.
2. The second tier, with the next highest priority, consists of any VOs that the site owners would prefer elevating in priority. Examples of the second tier would include collaborators of site owners or affiliated institutions.
3. The third tier would consist of all other VOs; of special interest are the VOs that implement the OSG Open Facility (currently the osg VO). We recommend that sites delegate management of cycles and related priority in the third tier to the OSG-ET. The OSG-ET, acting on behalf of the OSG Council,

will manage these in an opportunistic mode to address the objectives of the OSG community at-large.

Accounting and Usage Reports

The OSG project will provide each site detailed reports of usage at the level of individual project and wall hours for all opportunistic usage on the Open Facility. This will ensure proper recognition of sites in their role of making DHTC accessible to the researchers who use it and that sites can demonstrate what CPU cycles that might otherwise go idle are used towards. The OSG-ET will also provide regular reports to the OSG Council on access and usage on the Open Facility.