

# Me – My friends – the Grid

Frank Würthwein  
UCSD

We present a basic organizational principle for services within the Open Science Grid that builds on the idea of a Tier of University based computing centers that enables scientists at Universities to benefit from a globally distributed computing grid about which they have little or no control. The ideas expressed here aren't groundbreaking, nor particularly novel. In fact, they are quite obvious. We write them down here primarily to be able to refer to them later.

## **Me and the grid:**

Emerging grid infrastructures like the Open Science Grid are expected to grow by making minimal requirements on the participating sites that provide resources to the grid, and by allowing maximal heterogeneity. Requirements as well as common services are thus kept to an absolute minimum by design.

On the other extreme, the individual scientist who are the end users of the grid expect a feature rich environment that hides the heterogeneity of the underlying grid, and is no more difficult to use than their familiar desktop or laptop environment. In particular, the user of the grid expects to be subjected to a minimal set of interfaces that focus on functionality rather than technicalities, and provide instant gratification. E.g., a workload submission shall not take longer than a minute or so, independent of its size, and the user shall be able to operate the grid with a minimal set of input flags.

For the Open Science Grid to grow in both dimensions, resources and users, a Tier of most likely persistent services need to be developed that mediate between these two seemingly contradictory set of requirements. It would be most natural if those services were operated for the users by a layer of infrastructure that is sociologically close to the University based scientists who are to be the benefactors of the globally distributed and largely anonymous grid.

## **My friends:**

There is an obvious mismatch between the requirements for growth of the Open Science Grid as viewed from the two ends of the spectrum, users and providers of resources. To close this gap we need to deploy a University based infrastructure that provides the individual scientists with services like disk quota, task tracking and bookkeeping, compute quota for fast turn-around on small workload for debugging purposes, user and

group specific data and provenance catalogues, and probably many more. In a nutshell, all the “heavy lifting” required to turn a globally distributed computing grid with heterogeneous and generally minimal interfaces into an easy to use resource at the scientists finger tips needs to be provided by “my friends”.

## **LHC Computing as an example:**

Computing for the LHC in the US includes two large national labs, referred to as Tier-1 centers, where primary data for the experiments is stored, and user level data may be backed up to tape. In addition, a second Tier of smaller more distributed computing resources is deployed at Universities across the country. Each individual LHC scientist is assigned to at least one Tier-2 center, mostly according to their respective research interests. The user community at a Tier-2 center is thus expected to be reasonably homogenous and focused on a set of common goals.

Researchers develop their applications on their desktops or laptops, either at home, their home institutions, or at the Tier-2 center. They submit their applications for large scale processing to the grid by specifying little more than the dataset to analyze as well as the local release directory that contains their application on their laptop. A thin client tar's up the application on the laptop, copies it to the Tier-2 center, all generally in less than a minute.

At the Tier-2 center a set of services translate the processing request into potentially many thousands of jobs that can be submitted to the grid. Co-location of data and CPU as well as complete bookkeeping of these thousands of jobs is responsibility of the set of persistent services at the Tier-2 center. It is the responsibility of the Tier-2 center services to guarantee provenance tracking, best effort towards completion of all jobs that are part of the request, and to allow the users to track the status of their requests by querying the Tier-2 based bookkeeping whenever they see fit. In addition, the Tier-2 provides disk space to store the processing output for the user to within some limits.

## **Beyond the LHC community:**

Most of the LHC Tier-2 centers either already are strongly aligned with campus wide computing infrastructures, or are part of an effort to build such infrastructures. These range from centralized campus facilities to distributed campus grids. The notion of services provided for “me” by “my friends” at the Tier-2 center in order to allow me to easily access globally distributed computing grids is thus likely to be expandable to a much broader audience across the full spectrum of sciences on the Open Science Grid.