

# Topic 6

## Grid and Cluster Computing: Models, Middleware and Architectures

Craig A. Lee, Thilo Kielmann, Laurent Lefèvre, and João Gabriel Silva

Topic Chairs

Grid computing represents the common vision of truly general distributed computing across a ubiquitous, open-ended infrastructure supporting a wide range of different application areas. Realizing this vision will require a long-term collaboration of fundamental and applied computer science, industry, commercial infrastructure providers, and many, many application domains. Each meeting like Euro-Par is an important contribution to this long-term effort.

Grid computing has already made significant progress in the design, deployment and use of grids, but many challenges still remain before the vision can be realized. Many large scientific and engineering projects are adopting grids to support their project goals. Examples include the EU Data Grid project, NEESGrid, GriPhyN, and the NaReGI project, to name just a few. It is also possible to set-up small-scale grids to accomplish flexible scheduling and workflow management. Industrial organizations, such as Sun Microsystems, IBM, Microsoft, and Oracle already have or are starting to develop grid products and solutions. Standards bodies are pushing to bring stability to the marketplace with the development of standards such as the WS-\* set of specifications.

Even with this degree of progress, grids still suffer from issues of ease of deployment and use, fault tolerance and reliability, and the need for a true dominant best practice to emerge. Even though Globus is considered the de facto standard (and rightly so), there is still a wide “definition” of what constitutes a grid. Infrastructure developers, application domain experts and industry are all truly looking at an elephant. This is hampering the emergence of commonly accepted tools and best practices that can be codified in standards.

To help address these issues, we present these papers at Euro-Par that document the best work on many aspects of grid computing. This year 40 papers were submitted to Topic 6. The chairs assembled a team of 78 reviewers that produced 158 reviews. All accepted papers had at least four reviews. Many of the papers “in the middle” in effect received more than four reviews as the chairs read, reread and debated their merits. In the end, ten papers were accepted – an acceptance rate of 10%. These papers represent work in the areas of peer-to-peer, grid brokering, grid scheduling, load balancing in grids, processor farms for grids, virtual grid workspaces, grid information services, grid modeling, data replication, fault tolerance, and grid applications. We invite you now to browse these papers and closely study the ones that spark your interest and will help you in your work as much as possible.