Main Achievements:

An EGI MPI Virtual Team was established. The SA3 MPI (CSIC, TCD, UNIPG) activity participates in this group. The EGI MPI Support Unit is now fully functioning and available in GGUS.

CSIC has reviewed and updated MPI documentation available on the EGI wiki.

In addition, CSIC has also:

- Tested tight integration with SGE and PBS included in newer versions of mpich2 (1.3.1 and 1.5.b1),
- Tested accounting on SGE with different MPI implementations,
- Started to test MVAPICH, an MPI distribution with advanced Infiniband support, which is not currently supported by the default MPI-START,
- Submitted the abstract Parallel Computing workshop at EGI-CF 2012. CSIC are co-ordinating the programme.

CSIC contributed to The SIESTA computation chemistry application software. This was presented at the "Jornada de usuarios de Infraestructuras Grid" Workshop. See:
http://indico.ifca.es/indico/contributionDisplay.py?contribId=21&sessionId=14&confId=399

UNIPG has been working on parallelization of the CHIMERE application. The detailed analysis has identified the extent to which the application will achieve the highest level of parallelization possible, and thus the greatest efficiency. This model is structured around a task-farming model, in which some worker processes share the work by taking charge of part of the domain (subdomain) under the control of the master that sends and receives the relevant data, and that writes the results. The application can read initial input data based on the output of a previous run for a given timeframe. The model has several advantages:

- By chaining the results of many shorter runs, this allows simulations to run over much longer timeframes, and thus overcomes issues with CPU time-bound usage limitations at most resource centres.
- Can also be used as a form of application check-pointing.
- Allows greater granularity/resolution.

The CHIMERE parallelization pattern can be reused in many other applications

TCD has investigated the possibility of supporting generic parallel jobs (non-MPI) using the MPI-START framework. This work aims to exploit the OpenMPI mpiexec as a generic remote node process launcher. This was in investigated in response to a query from the Life Science Grid Community as to whether it was possible to support hybrid DIRAC/MPI workloads.

TCD has also investigated a serious issue in site configurations that may contribute to non-optimal job allocations and MPI job failures. We have developed a test suite to probe all sites supporting MPI-START/(Torque/PBS). Some sites are configured with the job CPU time limit (CPUT) to be set less than the total time for an individual CPU core (PCPUT). In addition, we have been following up on a requirement for the appropriate Information Provider to be fixed and made available as part of the UMD distribution. This issue affects 17 sites. Tickets are being opened against these sites, with recommendations on how to solve the problem.

TCD provided input for GPGPU Information System requirements at the "Towards an Integrated Information System" workshop. See:

https://www.egi.eu/indico/conferenceTimeTable.py?confId=654

In addition, TCD submitted the abstract "Supporting grid-enabled GPU workloads using rCUDA and StratusLab" to EGI-CF 2012.

Issues and Mitigation:

The predominant job scheduler MAUI needs to be upgraded at affected MPI sites. A node allocation bug/problem has been discovered in the standard distribution. EMI does not provide support for MAUI, and the software license may impose significant software distribution restrictions. We are currently working around this problem by recommending that sites install a version from an un-official repository.

MPI/Parallel Computing surveys delayed until month 1 in PQ8.

Plans for QR8:

MPI related tickets in GGUS shall be assigned to the new MPI User Support unit.

PQ8 work shall follow directly on from PQ7.

EGI-CF 2012 Parallel Computing Workshop.