

Highlights of your project (5-8 lines)

The main objective of the ComPat project is to develop generic and reusable High Performance Multiscale Computing algorithms (so called HPMC Patterns) that will allow to efficiently run multiscale applications on heterogeneous architectures and lift the simulations performance up to the exascale level. The project deals with various problems related to large scale multiscale applications development, deployment and execution with extreme data requirements, scalability, robustness, resiliency and energy efficiency.

What is anticipated technology (hw/sw/methodology) suggested for inclusion in an EsD project and describe the current maturity?

The ComPat project develops software and methodology that will allow to efficiently design, create and execute multiscale simulations on top of existing peta-scale and hopefully future exa-scale e-Infrastructures. It is expected to produce an innovative distributed computing system, which is able to spread computations on many HPC/cluster resources to boost the summary computational capacity intelligently and smoothly. For this reason, ComPat focuses on development of new pattern-based algorithms and their smart integration with middleware systems. The special attention is devoted to advanced scheduling, including energy-aware aspects, based on availability and characteristic of computing resources and profiles of applications. The project deals also with the optimisation of load balancing, data handling and fault tolerance. It is assumed to produce ready-to-use generic software components that may be useful for researchers and engineers from a broad spectrum of science domains. The ComPat outcomes possibly interesting from the EsD perspective includes experience, methodology as well as a set of software components. One of the most stable and mature developments employed in the ComPat is the QCG middleware software stack. Since QCG is available in UMD repositories and productively used in many HPC centers, it can be easily said that the middleware is on TRL 9. Currently the ComPat project is in the middle of the development and the first version of the system was already released. Thanks to the developed technology a few selected grand challenge multiscale applications can be already efficiently executed in a distributed fashion on several HPC clusters at the European level.

How should this technology be used / integrated (I/F, APIs)

The ComPat project is user- or in other words application-oriented. The aim of the project is to build a comprehensive environment for creation and execution of highly-demanding multiscale applications. The ComPat system exposes a consistent interface and APIs to its users. It is based on MML (Multiscale Modelling Language) description of multiscale scenarios and QCG middleware I/F. The MML language was designed and developed in the several previous EU projects, particularly in the FP7 project called MAPPER. The computing jobs in ComPat are managed by the QCG middleware, which constitutes an abstraction layer over local LRMS systems. QCG exposes its own Web Service interfaces and uses its own job description dialects (JobProfile, QCG-Simple). Integration with QCG may be done manifold: on a basis of integration with extensible access tools (e.g. QCG-Portal, QCG-Now, QCG-CommandLineClient) or using available Java and Python APIs. It is worth to note that some of the QCG's capabilities themselves (e.g. logical integration of many clusters, advance

reservation support) may be used not only for multiscale applications in ComPat, but also in variety of other, not only multiscale scenarios.

Are there any pre- or co-requisite items?

The advanced computations addressed in ComPat require flexible and efficient access to distributed computing infrastructure. The generic high performance multiscale patterns developed in the project assume that the simultaneous execution of many simulations' models will be easy to achieve. In order to provide necessary capabilities on this field, the ComPat project has selected and adapted the QCG middleware. For the ComPat it is therefore necessary to have QCG services deployed and configured on HPC/cluster resources.

Any extra work/interaction (on top of current project roadmap) needed to make them ready?

The ComPat project is implemented in accordance to the planned schedule. The current release of the developed system provides basic functionality needed to execute selected ComPat's applications. In the next phase of the project it is expected to extend the system in several ways. The QCG-Broker service's scheduling capabilities will be improved with support for energy and queue time metrics. Since currently only 2 of 3 HPMC patterns are supported, there will be also performed activities aiming to adjust the ComPat system to handle the remaining HPMC pattern. In order to prepare algorithms and functionality able to support multiscale applications in reaching the performance of the exa-scale class, the ComPat project assumes the intelligent integration of many tera- and peta-scale application modules that will be executed on a few large pata-scale machines. Unfortunately, it is currently not possible to easily integrate several production HPC/clusters to simultaneous execution of the application. The major problem is of procedural nature: it is a lack of advance reservation capability on majority of HPC/cluster resources. The only solution here is to have static reservations set up in advance at least on some dedicated HPC/cluster resources.

What information / actions are needed to best prepare for EsD projects?

The ComPat project joins together several groups of researchers that develop grand challenge applications from diverse domains: astrophysics, material-science, fusion and biomedicine. It also associates specialists in systems design, application optimisation and infrastructure technologies. On a basis of intensive cooperation between all project participants, ComPat tries to create generic and flexible system. It is expected that the final release of the ComPat tools and services will be exposed to the wide publicity. For majority of the scenarios the system should be ready to use out-of-the-box. Only for very specific use-cases small adaptations may be required. Therefore, it is foreseen that there are no particular actions needed to prepare ComPat for the participation in EsD. The feedback returned by current and further users should be the most valuable property.