

Making HPC accessible for SMEs

Andreas Wierse¹

¹SICOS BW GmbH

1 Motivation

Numerical computation, often also named simulation, plays nowadays in many enterprises an important role in the development process. Especially in large companies in the automotive or aerospace industry it is actually impossible to develop a new product without simulation technology. In recent decades the necessary know-how and personnel has been built up, but there have also been and will be significant investments into the infra structure.

Investments of a size, that can easily be handled by large companies, can be a real challenge for an SME. Especially a Return on Investment (RoI) calculation is much more critical for smaller companies. At least as long as it is not clear that the RoI will come in an acceptably short time frame, given that in comparison to the company's size the investment can be quite large. A powerful computer that has to be depreciated over several years must be loaded strongly in order to be worthy. In addition in smaller companies the usage of the infrastructure depends stronger on the sequence of projects. For them it is much more difficult to keep the load well balanced, which in turn makes the technology more expensive. This effect is visible even in larger companies.

A third issue comes up especially at smaller enterprises in the scope of a regular usage, but for some users due to the kind of the problems they have to solve already at the beginning: what can be done if a numerical simulation is too big for the available hardware or would block it for days or weeks and make it thus unavailable for other users. This can for example happen, when a very fine mesh is needed; or for parameter studies where a certain computation has to be carried out very often with slightly varying parameters.

2 Usage of a Supercomputer

A solution for the problems mentioned in the introduction could be to buy compute time at an external resource. This could take place via offers that are available in the internet: even Amazon provides resources today that could be used for this kind of numerical simulations; however compared to computers that are really optimized for numerical simulations these systems do not deliver the performance that specialized computers can provide; access via virtual machines for example brings significant performance disadvantages with it compared to the optimized architectures of dedicate high performance computers, especially for communication intensive applications like parallelized computation fluid dynamics simulations.

However normally the most powerful supercomputers are not available for industrial users. Typically they are installed at research centres and universities and allow only research use, implying a publication of the results. This is not acceptable for enterprises in most cases (except perhaps for pre-competitive issues), especially not for regular productive use.

An exception offers the High Performance Computing Centre Stuttgart (Höchstleistungsrechenzentrum Stuttgart, HLRS): the Cray XC40 (see image below; currently with more than 2.7 PetaFLOPS



on position 16 in the Top500 list of the worlds fastest computers) is open for industrial users for the simulations. In a long tradition of cooperation between research and industry in the state of Baden-Württemberg high performance computing plays an important role. Porsche for example performs a large part of its computations on HLRS systems since a long time, not the least due the reasons mentioned above, especially the flexibility.



Another research system that is open to industrial users is the bwUniCluster (see image above), installed at the Steinbuch Center for Computing (SCC) at the Karlsruher Institute of Technology (KIT). Being smaller than the large Cray system it nevertheless provides access to significant computing resources.

3 Access and Security

The access to the HLRS systems is for security reasons of course restricted; however this is realized in a way that still allows the integration of the system into the end users work flow. First there is only one way to login onto the system: via ssh; this allows fully encrypted access using Secure Shell mechanisms. The encrypted connection in combination with encrypted authentication provides a high level of security; there are no other means to reach the systems from the outside. Ssh however allows very easily to use scripts from the users local workstation to start and control scripts to run the simulation jobs properly. Naturally data transfer is also encrypted, based on scp, a close relative of the ssh tool.

Once the user is logged on, he is on a so-called login-node; here data can be transferred, but simulation runs cannot be started here. All numerical simulations take place on the compute-nodes, that will be assigned to a single user at any time only. All this is controlled by a batch system, typical for large and of course very large (about 100.000 cores) systems: the users submits his job into a waiting queue and depending on the load of the machine, the different queues (priorities) and the needs of this job, the batch system will start the jobs in a reasonable way to make good use of the system and minimize the waiting time for the users.

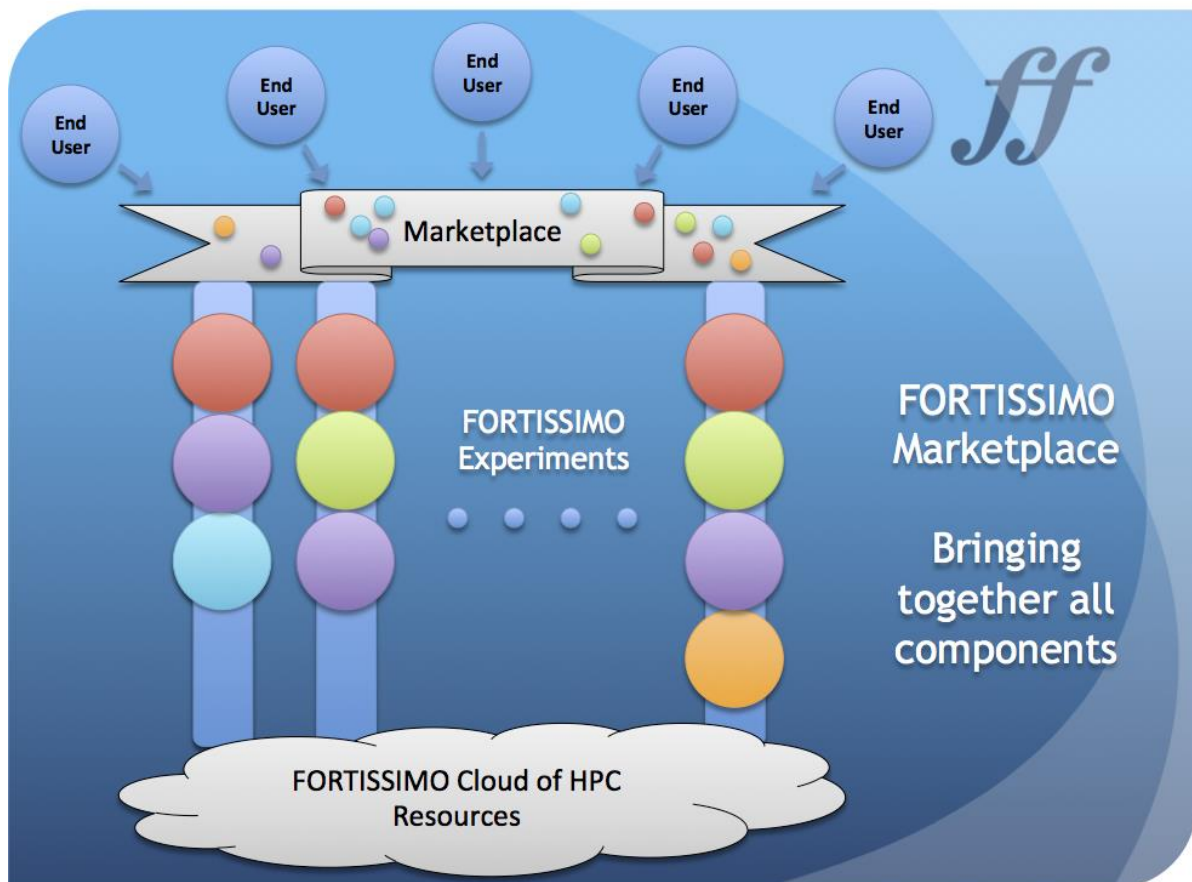
The required software licenses have to be provided by the end users, since the available licenses are only open for research and academic usage; however SICOS BW supports the users when it comes to getting the licenses and integrate them into the HPC environment (e.g. access to remote license servers).

The bandwidth to transfer data to and from the supercomputer is always an issue, especially in the context of large computations. In most cases the bottle neck is on the users side, since HLRS and SCC are connected to the state wide research net (BeIWue) with a bandwidth of several 10s of Gbit/s, soon 100 Gbit/s. Normally the start data (computational mesh plus boundary conditions) are significantly smaller than the result data (especially for time dependent computations). Bottle necks therefore are more likely to occur during the transmission of the results after the computation. But also for smaller bandwidths technology is available to support a reasonable access: remote visualisation with tools like VNC allows the users to leave their data sets on the system and just stream the video of

a visualisation session making use of dynamic compression; this works well, even with bandwidths in the order of 10 Mbit/s.

4 Fortissimo

Fortissimo is an EU-funded project to support especially small and medium sized enterprise to use HPC technology supported by cloud technology. Initially 45 partners, now close to hundred institutions are working together to develop the infra structure requirements that allow SMEs to quickly find the right partner for the solution of their simulation problem. To support this, a Fortissimo Marketplace will be established, where potential HPC users, HPC experts, Independent Software Vendors (ISVs) and HPC resource providers/computing centres can easily find each other. In addition tools will be developed to make the usage as easy as possible, focussing on the needs of small and medium size enterprises.



The Fortissimo partners are spread all over Europe; in addition to the centres in Edinburgh, Paris, Barcelona and Bologna the Höchstleistungsrechenzentrum Stuttgart (HLRS) is the German HPC provider partner in the project. When the project started twenty application projects have been selected and funded, that were used to systematically analyse the different application scenarios in the HPC context. The second call, after which additional projects were selected, has started already (again with about 20 experiments) and the experiments of the third call will start later this year. The focus is on understanding the requirements and needs of small and medium size companies when they use high performance computing resources.

Based on the experience gained in the experiments existing tools will be integrated and improved (if needed), to make it easier for SMEs to:

- find suitable partners
- find the right supercomputer
- find suitable applications/software
- use HPC resources in an easy-to-use environment

- perform the accounting with a full service accounting package
- with need-based accounting (pay-per-use)

The goal is, to provide the small and medium size enterprises with a tool, that makes access to the potential of high performance computing as easy as possible and allow them to improve their competitiveness independent from the size of their company.

5 Summary

Even if Supercomputing sounds expensive, it is well within the reach of even small and medium sized enterprises. Simulation tools are available as well as fairly easy to use access methods and the costs are very reasonable. SICOS BW and its shareholders in Karlsruhe and Stuttgart are working continuously to improve the accessibility even further; we are open to support SMEs to learn how to use this important technology to improve their competitiveness.