

Supporting Grid user and applications

Grid technology is not yet at the stage where an average researcher can simply start to use Grids without help. Early adopters are therefore likely to need support, both to identify which applications can benefit from Grids and to enable them to use the technology effectively. Moving existing applications to a Grid is likely to involve effort in modifying software and data, perhaps to run on a different operating system, as well as the time to learn a new interface and style of working.

Other than in a small number of subject areas, where research simply cannot be done in any other way, organisations are likely to need to persuade their users of the benefits of Grids. These efforts should be concentrated on users who are likely to see the greatest benefit; if satisfied, these users can then promote the service to others and form the basis of a self-supporting user community. A gradual build-up will also allow users, organisations and the Grid service to develop together, without overloading any part of the system through a sudden increase in demand. Organisations should plan to work individually with these early users to help them transfer their applications to the Grid: this may include anything from porting software to a different operating system to obtaining the necessary credentials to use the Grid. It may be worth designating one or more members of staff to liaise between users and the service provider: these individuals should aim to become familiar both with Grid technologies (including their implications for the network and services) and the types of research done in the organisation. They will then be well-placed to identify those research areas that are likely to benefit most from Grids, as well as working with Grid service provider(s) to find the most appropriate technologies to support research and other users. With a wide view of the potential applications it may be possible to identify standard application-level services that a Grid can provide that will be useful to a range of users.

It is important to match the Grid to appropriate user requirements. Shared Grids can be used either by users who have occasional requirements to run large, intensive jobs, so leave the Grid free for others when they are reviewing the results, or by users who are prepared to have their jobs run in the background, using whatever spare CPU cycles are available. Applications that demand continuous use of large resources are unlikely to be comfortable sharing a service with others. The computers used in a Grid also offer different characteristics: some are designed to support tightly-coupled simulations, where a large number of processing units work simultaneously on the same data; others provide a large number of independent processing units under central management (sometimes referred to as CPU farms) and are best suited to tasks that involve running a large number of independent jobs whose results are combined afterwards. Running an application on a mismatched system is at best wasteful of an expensive resource and may even result in the work taking longer than on a single system. Organisations that own tightly-coupled systems may find it worthwhile also to implement a CPU farm Grid so that the expensive tightly-coupled system can be dedicated to jobs that require its capabilities. Simply adding more CPUs to a farm may not be the most effective way to increase performance: Grids that let users submit their own programs will need compilers and the efficiency of compilation can be a significant factor in the time taken to complete the

resulting jobs. Money spent on a compiler that produces fast executable code may well be well spent.

Source URL: <https://community.jisc.ac.uk/library/janet-services-documentation/supporting-grid-user-and-applications>