Why is the Grid different from other applications?

Grid computing uses a distributed architecture to let teams of researchers in different physical locations and time zones collaborate on large scale projects in many disciplines. It is a novel application that involves new models of computing and new network protocols. Grids require complex and dynamic patterns of trust to be established and implemented, with computers relying on one another to make decisions on the identity and rights of individual users. As noted in the Gridwelten report [Lindner, p49], the most popular Grid software has often been designed for closed environments where there is a high level of trust between systems and the network infrastructure is very open. Inside such an environment, the major concern is ensuring that legitimate users only obtain the resources to which they are entitled: those who are not entitled to use the system at all are assumed (often implicitly) to be excluded entirely by some external perimeter surrounding the closed, trusted environment. Security tends therefore to concentrate on allocating resources to authenticated users within the trusted environment, not on dealing with attacks that may come from outside.

However, universities and university networks have moved away from a simple security model where everything inside a perimeter is trusted and everything outside is not. Pressure for change has come from two directions:

- the increased need to collaborate with people and systems outside the perimeter
- the increased threat from rapidly spreading worms and viruses that may appear either outside or inside the perimeter.

Production networks in education now tend to have security measures, such as routers and firewalls, controlling the flow of traffic both between and within organisations. These measures are designed to manage traditional computing models: where client and server functions can be distinguished and are normally performed by different systems; where network protocols are clearly defined and recognisable; and where trust relations are simple, generally static, and reflect existing organisational hierarchies.

As a result, deploying a Grid system on a production network may present a number of challenges. Within organisations, Grids may not find the open networks and trust models that they are designed for. Between organisations, Grids are unlikely to be protected by a complete security perimeter. Indeed, cross-institutional Grids themselves require a breach in the perimeter, so need to be concerned about direct or indirect threats from outside the Grid environment. Discussion is likely to be needed between Grid managers, designers and users, and production network managers on how the openness of the Grid can be accommodated without increasing the risks to both Grid and existing systems. This document aims to facilitate those discussions, by describing the requirements of networks and Grids, using in particular good practice that has been established by sites with successful Grid deployments. It aims to show how Grids and security measures can be adapted to work together on the same network while preserving both security and functionality.

The document considers general aspects of how Grids should affect network design, then looks at tools that may contribute to a successful Grid deployment. A series of appendices covers deployment aspects of particular Grid software packages and it is hoped to increase the number of packages covered in future editions as information becomes available.
