



# Grid Support

**Andrew Cormack**



**Technical Guide**



## UKERNA Technical Guides

UKERNA Technical Guides are a series of user guides available to JANET customers. Technical Guides contain detailed technical information and are intended for technical support staff at JANET sites, network specialists, or those with a particular interest in the specialist area.

If you have any queries or comments about the Guides or would like to obtain copies, please contact:

JANET Customer Service

UKERNA

Atlas Centre, Chilton, Didcot  
Oxfordshire, OX11 0QS

Tel: 0870 850 2212

Fax: 0870 850 2213

E-mail: [service@janet.ac.uk](mailto:service@janet.ac.uk)

Further details of the documents in this series are available at:

<http://www.ja.net/services/publications/technical-guides/>

# Table of Contents

<b>1. Introduction .....</b>	<b>3</b>
<b>2. Types of Grid Service.....</b>	<b>5</b>
<b>3. Defining Grid Services.....</b>	<b>7</b>
<b>3.1 Technical Specification .....</b>	<b>7</b>
<b>3.2 Users, Authorisation and Virtual Organisations .....</b>	<b>7</b>
<b>3.3 Policies .....</b>	<b>8</b>
<b>4. Supporting Networks for Grids.....</b>	<b>9</b>
<b>4.1 Network Design.....</b>	<b>9</b>
<b>4.2 Network Performance .....</b>	<b>9</b>
<b>5. Supporting Systems for Grids .....</b>	<b>11</b>
<b>5.1 Supporting Servers.....</b>	<b>11</b>
<b>5.2 Supporting Workstations .....</b>	<b>12</b>
<b>6. Supporting Grid Users and Applications.....</b>	<b>13</b>
<b>Appendix A: Checklist of Support Activities .....</b>	<b>15</b>
<b>Appendix B: Good Practice Ideas and Examples .....</b>	<b>17</b>
<b>B.1 Service Definitions .....</b>	<b>17</b>
<b>B.2 Network Support .....</b>	<b>17</b>
B.2.1 Network Design .....	17
B.2.2 Network Performance.....	17
<b>B.3 System Support.....</b>	<b>18</b>
<b>B.4 User and Application Support.....</b>	<b>18</b>

---

# 1. Introduction

Grids represent a new approach to computing, using high-speed networks to let researchers collaborate and gain access to the computing and data resources they require. Grids should reduce two of the constraints on current research: the requirement for researchers to be physically close and the need for researchers to have either physical possession or prior allocation of all the computing resources they will need. Grids should therefore improve the efficiency of both researchers and computers.

This will only be achieved once researchers can simply use a Grid as a research tool, without having to think about the computers, networks and software that comprise it. In other words Grids and their users need to move from an ad hoc approach, where Grids are constructed as technology experiments, to a service approach where the behaviour of a Grid is predictable and defined, and its availability can be relied upon. This change is much more about people and processes than about technologies.

This document therefore identifies the major areas that need to be considered by any organisation that wishes to provide or use Grid technologies as a service, and gives some examples of how these challenges have been addressed. A companion document, *Deploying Grids*, covers the technology issues that are likely to be encountered in using Grid technologies.

Whether an organisation wishes to provide its own Grid services, or just to enable its own users to use Grid services provided by others, it is likely to have to develop policies and processes in four main areas:

- Defining Grid services, or understanding and complying with definitions of Grid services provided by others.
- Developing processes to ensure that network resources required by Grids are available, without disrupting other uses of the organisation's networks.
- Developing processes to ensure that computing, storage and workstation systems are available for Grid use, while retaining appropriate security for these and other applications.
- Supporting users of Grid services and their applications.

These discussions should include both Grid users and IT service providers, and should start as early as possible to give the greatest likelihood of success.

This guide reviews the types of Grid service that exist or are in development. It then considers each of the support areas in turn. A checklist of the areas likely to require action can be found in Appendix A: Checklist of Support Activities. Examples of how organisations have approached these issues are in Appendix B: Good Practice Ideas and Examples.



---

## 2. Types of Grid Service

Work is in progress to create Grid services at the levels of individual campuses, both nationally for the UK and internationally. Each level is likely to be based on components contributed by the levels below it. This section therefore gives a brief overview of the characteristics of each type of service and the agreements that are needed to build them into a consistent whole.

Campus Grids are normally intended primarily as a service for members of organisations where they are located; it is usual for users to have to obtain permission to use the Campus Grid rather than the service automatically being open to everyone in the organisation. Users may authenticate to a Campus Grid using their existing login credentials or using special credentials for the Campus Grid. The aim of most Campus Grids is to make better use of the computing resources that the organisation already owns, though it should not be thought that the additional capacity comes for free as there will be an increased cost – at least in electricity and air conditioning – when computers are active rather than idle. Campus Grids may use any mixture of spare processing cycles on computers bought for other purposes, such as classroom workstations, and dedicated computing hardware, such as a High Performance Compute cluster. Spare cycle grids are usually constrained to run the operating system required for the systems' primary use, so Grid users may have to adapt their jobs to a new working environment. It is also important to ensure that the background use of these systems for Grid processing does not unduly affect their primary purpose. A number of Campus Grids have been created and many of these are moving from experimental to service status.

The UK National Grid Service (NGS) was created to provide computer and data services to researchers in the UK. The aim is to provide a national service, more powerful than is commonly available on a single campus, to let researchers run jobs that are too large for their local campus resources to complete in the available time. The NGS is currently available for scientific and academic research projects, which may be subject to peer review. Users are authenticated using digital certificates from the UK Grid Certification Agency or other recognised certification agencies. The first four nodes on the NGS were funded as dedicated services at four locations with good connections to JANET. These nodes have been used to define standard services and interfaces. The NGS is now working with partner organisations who undertake to provide those same services and interfaces, with service levels chosen by each partner, to allow their Campus Grids to be accessed, managed and monitored as part of the centrally managed NGS.

A number of international efforts are in progress to allow researchers to collaborate using computers and data from different countries. Possibly the most advanced, in terms of providing a Grid service, is the EGEE (Enabling Grids for ESciencE) project. This requires that all computers run the same package of software and requires their managers to commit to providing defined levels of service. Researchers submit batch jobs for processing somewhere on the EGEE pool. This pool includes national grid services (which may themselves contain campus grid services) as well as services provided by national and international computing and data centres. Once available resources are identified, the jobs run on them and results are returned to the user. EGEE has the same goal as the UK NGS, to make large computing resources available for relatively short periods of time, but also focuses on the need to support collaborations by making large data sets available over a wide area.



---

## 3. Defining Grid Services

Grid services are distinguished from ad hoc Grids by the fact that the service they provide is clearly defined. It is obvious that an organisation that wishes to provide a Grid service must consider how that service will be defined, but organisations that plan to be users, rather than providers, of Grid services also need to review the service definitions to confirm that the service is appropriate and that the organisation (and the user) can satisfy any technical and procedural requirements that may be placed on them.

### 3.1 Technical Specification

The service definition should include details of what technologies (software, interfaces, etc.) are supported by the service. At present there is a wide variety of incompatible software that is able to run a Grid, and it is always likely that there will be restrictions on the programming and application tools that can be used on any particular Grid service. Different types of Grid software will make different demands on the networks used for communications (see later) so the prospective user and their network providers must ensure that they can support these before spending time preparing to use a Grid service that may, in practice, be unusable. Anything claiming to be a 'service' should provide some technical support for its users, but thought must be given to what level of support can be provided by the service and whether this matches the requirements of its intended users. For all of the defined services – both technical and support – it should be clear whether the level of service is guaranteed or provided on a best efforts basis. Even a best efforts service will allow some users to do things they could not otherwise achieve, but for others a lower-performance guaranteed service will be preferable to a service that may complete their job much faster but offers no guarantee that it will not be slower.

### 3.2 Users, Authorisation and Virtual Organisations

Grid service providers must also consider who will be permitted to use the Grid service. Restricting the service to local users avoids many problems, from technology through authentication and policy enforcement to support and communications, but loses many of the potential benefits of Grids as a collaborative tool for a widely distributed research community. Whatever user group is chosen, the grid provider must consider whether the service will be automatically available to all users in that class (staff, students, etc.) or whether users will be accepted individually. Again there is a trade-off between the additional administration required for individual authorisation and the uncontrolled demand for service and support if the service is available to all. As discussed later, not all applications will be appropriate for a Grid service, so it may be best to start working with selected research groups to allow both the users and the service provider to develop their understanding of what the technology can and cannot achieve.

Some Grid services are made available to particular groups, such as research projects or communities of interest. Where these groups include individuals from a number of different organisations, it is common to view them as a Virtual Organisation and to pass to them the responsibility for authenticating and accounting for the activities of individual users. Some Virtual Organisations will be formally constituted organisations in their own right, but others may have no existence other than through the Grid. This may raise issues for the service provider who may no longer be able to identify individual users, for the Virtual Organisation which must ensure it can meet the requirements of the services its members access, and for the members' home organisations which may be expected to act as the ultimate enforcers of policy for services they were not aware their users were using. The responsibilities of these various parties need to be clearly stated and agreed in advance to reduce potential problems.

Service providers may also wish to remove access permissions, either temporarily or permanently, from certain users or groups, for example because their period of authorisation

has ended or because problems have arisen. Definitions of when and how this will be done must be clear.

### **3.3 Policies**

Any computing or network resource is likely to have some policies that regulate its use and may impose obligations on its users and their organisations. These policies will aim, among other things, to ensure that the resource is available on a fair basis to all those who are authorised to use it. Typically, policies will state what the acceptable uses of the resource are (including who is allowed to use it), and the steps that are taken by the service and required of its users to ensure the security of the service. Security is important both to ensure that the service is available and reliable and that it does not present an undue risk to its users or to others. By their nature, Grids tend to be relatively open to authorised users, so are likely to depend more on users behaving responsibly and abiding by policies than on controls enforced by technical means. There is therefore the possibility that users may accidentally or deliberately breach the Policies of a Grid and that their home organisations will be expected by the Grid provider to enforce any sanctions that apply for these breaches.

Grids also present a Policy problem in that Grid providers may not know who their users are and users may not know which Grid resources they are using. Determining what Policies apply to a particular user, and informing the user of those Policies in advance, may therefore be difficult. Grid providers can help by ensuring that their terms of use are similar to those generally used by a particular type of service or group of users, and by publicising their Policies to users wherever possible; users also need to be sensitive and cease any activity that they are informed is a breach of a particular local policy. Organisations must support the Policies of others, even where a particular activity would not have breached the Policy of the home organisation.

---

## 4. Supporting Networks for Grids

Grid technologies often make novel demands of networks, both in terms of performance and by their use of new protocols. It is therefore essential to have processes both to assess what proposed Grid activities will require of site and external networks, and to ensure that any changes to provide for these do not undermine existing measures that have been implemented on networks and servers to reduce security risks. Many technical and design solutions exist that allow Grids to be used on existing production networks (these are discussed in the *Deploying Grids* Technical Guide), but unless an appropriate solution is chosen then serious damage may be done both to the Grid and to the network service.

### 4.1 Network Design

It is unlikely that the use of Grid technologies will be successful unless the design of the network has taken them into account. Providing connections that allow the required protocols to be transmitted with the necessary characteristics of bandwidth, delay, packet loss and jitter may require adjustments at any level of the network, from physical connections to transport and application layers. New control points, such as application gateways, may be required both to ensure the Grid performs adequately and that its use does not damage other applications with which it shares the network. Many different technical solutions exist that will allow most requirements to be met; however, this will only be achieved if there is a clear understanding of the requirements of both the Grid and the network service, a reasonable assessment of the risks to both, and adequate resources provided to satisfy both the performance and security requirements. Unless Grid and network providers and operators work together to create and maintain a Grid service, it is very unlikely that it will deliver the expected results.

As Grid technology – and the organisation’s use of it – develops, it is likely that the demands on the network will change. This may involve greater demands on the network, or may allow simplification. A continuing process must therefore be established to review the service, to identify problems and opportunities, and to evaluate, implement and maintain solutions.

### 4.2 Network Performance

The intensive use of networks involved in some Grid applications means that performance issues are particularly important. For traditional network applications, available bandwidth and packet loss have been the most visible performance metrics for end-users; however, some Grid applications may also be adversely affected by latency and jitter, which are less commonly measured and may change more frequently. The use of Grids is therefore likely to require new performance measurement and diagnostic tools for network operators, as well as processes for users to characterise and report problems as they are occurring. For some applications, network behaviour all the way to the workstation may have an impact on perceived performance. Measuring and diagnosing problems in this area may well require the assistance of the user, aided by appropriate user-friendly tools. Some Grid systems attempt to tune their behaviour to the current network performance: this can be helpful but such behaviour must be controlled so that it does not affect other uses of the network.

Assessments of network performance will be needed at three different stages of the use of Grid technologies. Users will wish to be assured of an adequate level of performance before they commit effort to using new technology, so it will be necessary to validate the network’s capabilities in advance; routine monitoring both of the performance of Grid networking activities and their interaction with other uses of the network will be required to manage and develop the service; and effective and prompt data gathering and analysis will be needed to resolve faults and incidents. Each of these will require its own processes and may require some specific tools and expertise.



## 5. Supporting Systems for Grids

Grid systems consist of large combined CPU resources, connected by fast networks and with a considerable degree of mutual trust. These are essential features for their intended purpose but they also make Grids a significant operational and security problem. Any failure or misuse of a Grid is likely to spread rapidly, because of the trust and inter-connections, and to have a large impact because of the power of the computing resources involved. To prevent failures as much as possible, and to reduce the impact of those that do occur, it is essential to have effective processes to configure and maintain Grid systems, and to detect problems quickly and respond rapidly to contain them.

### 5.1 Supporting Servers

Computers that provide Grid services need to run a large amount of complex software, some of which may be experimental rather than of full production quality. Software configuration and maintenance are therefore likely to require significant planning and continuing resources, to cover all layers from the operating system and network services to Grid services and application software. A failure at any of these layers can wreck the security and operational stability of all the others.

Some potential problems can be reduced by careful choice and initial configuration of software. Packages that are familiar and are known to be well-maintained are to be preferred where possible. All software, including the operating system and network services, will need to be updated regularly and the ease of doing this will be particularly important if there is a large number of systems to be managed. Operating systems and packages that provide updates as self-installing patches are likely to require much less effort than those where code needs to be edited and re-compiled. In particular, any package that limits the ability to install patches to other parts of the Grid system must be treated with great care. Any system for which a patch is available but has not yet been installed is at very high risk and may need to be isolated from other systems and shared networks. Software developers must avoid creating such dependencies.

It must be clear who is responsible for installing and maintaining all the software components of each Grid server, and those people must have sufficient time available and the necessary authority to do this vital job. They must subscribe to relevant mailing lists provided by software vendors and others; when a new problem or patch is discovered, remedial action is likely to take priority over all other work.

Where a number of similar computers make up a Grid, for example in a Campus Grid or CPU farm, it will often be simplest to automate the process of installing and updating software on all the computers. A number of sites include Grid server software as part of their standard managed workstation image, though this may require individual agreements and configurations if the computers that make up the Grid are not all managed by the central computing service.

Like any other computer, Grid systems should require each user to authenticate themselves to prove their entitlement to use the resource and their personal or group files and programs. Grids may use a user authentication system that already exists in the organisation, in which case additional configuration is likely to be required to give the Grid software and systems access to the central authentication database, or they may have their own stand-alone authentication methods. Authenticated users may be mapped to individual local user accounts, or all jobs may run as a single Grid user with the Grid software ensuring that users cannot interfere with each others' files or jobs. Whatever approach is used, there should be some reliable way to identify the person who owns each job and file, and to contact them speedily if the job appears to be causing problems either accidentally or deliberately. Some means also needs to be provided for users to get data on and off the Grid, either by transferring files manually or by giving the Grid systems access to a central networked filestore.

Despite care in configuring, maintaining and using the Grid system it is likely that there will be operational or security incidents from time to time. Processes for detecting and responding to these incidents are essential to limit the impact when they do occur. The processes for responding to incidents on networked systems are relatively well understood in the incident response community (for example in documents produced by the CERT Co-ordination Center – <http://www.cert.org/csirts/>) but inter-organisational Grids pose new challenges because they cross traditional organisational boundaries and areas of responsibility. Speed of response is also essential because the mutual trust within the Grid will allow the impact of any incident to spread very rapidly. Depending on the likely extent of the problem, it may be appropriate to contain it by disabling a single user or identity or group, a single computer, a site or an entire Grid. An agreed process for making, authorising and implementing these decisions whenever an incident occurs is essential.

## 5.2 Supporting Workstations

Although some Grids provide access through standard applications such as web browsers, others require additional software to be installed on the user's workstation. As with the server software, this can be a complex process and is probably best done as part of a central configuration management system.

Authentication can also present a challenge for the workstation. Some Grids use the username and password combination with which users are familiar, and these do not normally require any special arrangements on the workstation. However, other Grids use hardware tokens, software tokens or digital certificates, and therefore require additions to the client workstation. If a physical card reader or other interface is required then this will obviously limit the workstations from which the user can access the Grid. Software tokens and certificates that have to be saved on the local disk also restrict which workstations can be used: if the user has their own personal workstation on which the certificate is installed then the workstation effectively becomes an expensive hardware authentication token and other users must not be allowed to use it. If the user wishes to use multiple, shared workstations then there are problems both of moving certificates from one workstation to another and of ensuring that the certificate is not left behind on disk or in memory when the user leaves the workstation, potentially allowing others to gain unauthorised access to the Grid. A number of solutions are being developed to store certificates in a central store where users can access them by other authentication methods, and these seem the most promising solutions to these issues.

## 6. Supporting Grid Users 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.



---

# Appendix A: Checklist of Support Activities

## Service Definitions

- Grid Service Providers
  - ensure that technical and support services offered are defined and available to users, along with any service level definitions
  - define who may use the Grid service, and document the process for gaining authorisation (and authentication credentials if required)
  - define technical and network requirements to use the Grid service
  - define acceptable use and other policies for the Grid and make these available, as far as possible, to all authorised users
  - ensure that appropriate records of use are kept (subject to the Data Protection Act 1998) to support investigations of misuse
  - define responsibilities of all parties to prevent misuse and to assist in investigations (particularly if virtual organisations are allowed to authorise their members)
  - define circumstances where authorisation to use the service may be withdrawn, and ensure that organisational and technical processes exist to enable this to be done promptly.
- Organisations with Grid users
  - ensure that Policies of the Grid services used are known and supported
  - ensure that local Policies and processes allow responsibilities required of Grid users' organisations to be satisfied.

## Network Support

- Network Design
  - establish processes to identify and monitor network requirements of Grid applications
  - based on risk assessment, identify appropriate technical solutions to support Grid use while maintaining appropriate performance and security for Grid and other applications
  - incorporate Grid requirements into network design and provision.
- Network Performance
  - review network provision to Grid workstations and servers, and improve components or design if necessary
  - deploy processes, tools and expertise to diagnose performance and functional problems.

## System Support

- Server Support
  - assign responsibilities and authorities and provide necessary effort to operate servers securely
  - configure and maintain Grid servers: operating system, network services, Grid services, applications

- consider software and configuration management systems
- plan and provide for user authentication and identification
- provide file management systems
- assign responsibilities and document process for incident response.
- Workstation Support
  - if necessary, implement software management practices as for servers
  - implement secure authentication methods that allow users to access grids from necessary locations without the risk of unauthorised access
  - educate users on safe use of authentication tokens.

## **User and Application Support**

- Promote Grid Use
  - work with IT Service to ensure Grid and network requirements are both satisfied
  - assign liaison staff to work with potential users
  - identify appropriate technologies/applications for Grid services.
- Support Early Adopters
  - discuss requirements of early adopter projects
  - help early adopters to transfer their applications to the Grid service (e.g. software modifications, client requirements).
- Support Development of Grid Service
  - identify standard application services of benefit to multiple users
  - monitor service use and adapt service to suit user requirements.

---

## Appendix B: Good Practice Ideas and Examples

This appendix contains ideas and examples of good practice in Grid support gathered from various conferences and meetings with organisations that are making effective use of Grid technologies. It is hoped to expand this collection in future editions: please send any suggestions for ideas to include to [A.Cormack@ukerna.ac.uk](mailto:A.Cormack@ukerna.ac.uk).

Where possible, references to the source of the idea are included, though some of the case studies were given by personal communication or on condition of anonymity.

### B.1 Service Definitions

- Documentation of the services provided by the UK's National Grid Service is at: <http://www.grid-support.ac.uk/content/category/7/29/49/>
- A Security Policy for UK e-Science projects is at: [http://www.pparc.ac.uk/Rs/Fs/Es/New\\_Security\\_Policy.asp](http://www.pparc.ac.uk/Rs/Fs/Es/New_Security_Policy.asp)
- The LHC Computing Grid project is carrying out a series of 'Security Challenges' to ensure that security measures implemented on components of the Grids are working as intended. The security challenges are designed to test both the technical and human implementation of security measures – for example, running a job on the Grid and then asking sites to trace which users, processes and files were related to the job. This activity should help to ensure that when a truly malicious job appears on the Grid, service managers both have the tools to identify and deal with it and are familiar with how to use them: <https://twiki.cern.ch/twiki/bin/view/LCG/LCGSecurityChallenge>

### B.2 Network Support

#### B.2.1 Network Design

A wide range of technical measures for adapting networks to the requirements of Grids can be found in the companion UKERNA Technical Guide *Deploying Grids* (<http://www.ja.net/services/publications/technical-guides/tg-grid-deployment.pdf>). In some cases there will be a number of possible measures to choose from, working at different levels of the transmission system. If lack of bandwidth or congestion is a problem then it is likely that solutions will need to re-arrange or improve the physical infrastructure. Security requirements may be addressed at the physical layer or at the transport layer, for example by using VLANs, or at the application layer by using tunnels. Application-layer gateways can often simplify other security measures: for example, a number of sites are planning to use Condor-G to link their Campus Grids together using Globus protocols rather than running the more complex Condor protocols over wide area networks.

In some cases it may be necessary to connect a demanding Grid resource directly to the external network if its traffic would be too disruptive to existing security measures. If this is done then the Grid resource must ensure that it has its own security measures as it will have no other protection from the hostile 'background noise' on the Internet. Any internal systems that need to connect to the exposed Grid server may also need special security measures to prevent them acting as bridges for any hostile activity.

#### B.2.2 Network Performance

The GN2 project included a Performance Enhancement and Response Team that investigated network performance problems and built up a knowledge base of common problems and solutions as well as strategies for investigating problem reports. The results

are published as a Wiki (<http://pace.geant2.net/cgi-bin/twiki/view/PERTKB/WebHome>) and a report (<http://www.geant2.net/upload/pdf/GN2-05-176v4.1.pdf>).

### B.3 System Support

Campus grids are often constructed where there is an existing managed workstation service – for example, covering open access terminal rooms. Since the software images run on these systems are managed centrally and installed automatically on all workstations it is relatively straightforward to install grid software as part of this standard image. This also allows upgrades and patches to be rolled out relatively easily:

**<http://www.nesc.ac.uk/action/esi/download.cfm?index=2457>**

Care may be needed where the workstations forming a Campus Grid are (or are perceived to be) intended for use by a particular faculty or department, rather than as a general university facility. In these circumstances there may be objections to running jobs from other departments or organisations unless local users and system managers are consulted. In some cases it may be appropriate to have separate Grids incorporating central and departmental machines, or at least to apply different access rules or priorities (if the software allows this).

### B.4 User and Application Support

The computing service at one organisation has regular meetings with departmental users. During these meetings the opportunity of using the Campus Grid was discussed and promising projects or activities identified. Computing service staff then worked with these projects to help them port their applications to the Grid. Although it will not be possible to offer this type of one-to-one assistance in the future, it is hoped to gather sufficient experience and good practice information that future projects will need less intensive support.

Successful early adopters should be used to promote the use of the Grid. Where researchers have had a good experience and achieved things that would not have been possible without the use of Grids, they are likely to report this to colleagues both within and outside the organisation so that use can grow by word of mouth. Internal newsletters and published papers can be effective ways to promote Grid use.

Plan for your Grid to be successful! Know in advance what you will do when demand fills the available resources, and don't make promises to early adopters that you will no longer be able to deliver when they are not the only users of the Grid.

A common problem experienced by Campus Grids is that the available workstations run Windows® operating systems whereas large applications are written for various flavours of UNIX®. Rebooting machines into 'Grid mode' overnight is a possible solution but needs to be scheduled carefully to avoid disrupting the primary use. Operating systems virtualisation may be preferable, where a UNIX® virtual machine can run on a Windows® platform:

**<http://www.nesc.ac.uk/action/esi/download.cfm?index=2462>**

[Note that since the presentation in the above URL was given, the VMWare 'player' software has been made available free of charge.]



## Tell us what you think

Technical Guides are produced and published by UKERNA for use within the JANET Community.

We welcome your comments on all aspects of this document and on any other UKERNA publication.

Please direct feedback to JANET Customer Service, at the address below, or e-mail us directly at:

documentation@janet.ac.uk

UKERNA manages the networking programme on behalf of the higher and further education and research community in the United Kingdom. JANET, the United Kingdom's academic and research network, is funded by the Joint Information Systems Committee (JISC).

For further information please contact:

JANET Customer Service  
UKERNA  
Atlas Centre, Chilton, Didcot  
Oxfordshire, OX11 0QS

Tel: 0870 850 2212  
Fax: 0870 850 2213  
E-mail: service@janet.ac.uk

### Copyright:

This document is copyright The JNT Association trading as UKERNA. Parts of it, as appropriate, may be freely copied and incorporated unaltered into another document unless produced for commercial gain, subject to the source being appropriately acknowledged and the copyright preserved. The reproduction of logos without permission is expressly forbidden. Permission should be sought from JANET Customer Service.

### Trademarks:

JANET®, SuperJANET® and UKERNA® are registered trademarks of the Higher Education Funding Councils for England, Scotland and Wales. The JNT Association is the registered user of these trademarks.

Microsoft® and Windows® are registered trademarks of Microsoft Corporation in the United States and/or other countries.

UNIX® is a registered trademark licensed to X/OPEN.

### Disclaimer:

The information contained herein is believed to be correct at the time of issue, but no liability can be accepted for any inaccuracies.

The reader is reminded that changes may have taken place since issue, particularly in rapidly changing areas such as internet addressing, and consequently URLs and e-mail addresses should be used with caution.

The JNT Association cannot accept any responsibility for any loss or damage resulting from the use of the material contained herein.

### Availability:

Further copies of this document may be obtained from JANET Customer Service at the above address.

This document is also available electronically from: <http://www.ja.net/services/publications/technical-guides>



© The JNT Association 2006

The logo for JISC, consisting of the letters 'JISC' in a large, bold, orange, sans-serif font.



