Published on *Jisc community* (https://community-stg.jisc.ac.uk)

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.

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 standalone 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/ [1]) 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.

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.

Source URL: https://community-stg.jisc.ac.uk/library/janet-services-documentation/supporting-systems-

grids

Links
[1] http://www.cert.org/csirts/