General issues

Aspects of Grid Protocols

The network protocols used by Grid systems are often complex. They may use ephemeral source and destination ports selected randomly each time a process starts, mixtures of TCP (Transmission Control Protocol) and UDP (User Datagram Protocol) and, in some cases, connections opened from server to client rather than in the normal direction from client to server. Routers and firewalls designed to handle conventional TCP connections from a client to a static, well-known port on a server are unlikely to manage these new protocols correctly. Firewall vendors have written software to interpret other complex protocols, such as FTP (File Transfer Protocol) and streaming media, but Grid protocols are not yet sufficiently mature or widespread to be supported in standard firewall devices.

Grid protocols may also be required to transmit very large quantities of data at very high speeds, so routers and firewalls must not require excessive amounts of CPU power or other resources to interpret them. There is always a risk for any protocol that an under-specified router will become a performance bottleneck, but the design of networks for Grid systems presents a particular challenge.

Aspects of Firewalls

Firewalls and router controls are deployed on production networks for three reasons:

- to reduce the exposure of vulnerable systems to attack
- to limit the impact and contain the spread of successful attacks
- to manage network traffic to ensure that important applications have the network resources they need.

For example most universities will have controls:

- separating their administrative computers from the general network, to protect the data on those systems and ensure that they are available for running the university
- between the general network and any connections (e.g. wireless networks) used by short term visitors whose computers may be infected by Internet worms
- between their LAN (Local Area Network) and the wider internet to ensure that e-mail traffic is routed via the campus mail hub where it can be checked for viruses.

All of these purposes – protection, containment and traffic management – are also very relevant to Grid systems. As noted by Hillier [Hillier, p4], Grid software is mostly experimental and under continuous development and it is therefore likely that it will contain bugs that make it vulnerable to attack. Grid machines are high profile targets of a kind particularly attractive to the intruder community – having large resources of CPU, disk and bandwidth – so must

expect to receive more than the average number of attacks. Grid systems are also designed to facilitate movement from one system to another, so the consequences of a successful attack are likely to be widespread and dramatic unless it is contained. The large network flows involved in many Grid applications require uncongested networks where they are not competing with other traffic that is not necessary. Grids on shared networks have as much need for traffic control measures as any other system. The trust model used by Grids, where an authenticated user can gain access to many server systems, means that both workstation and server machines need to be protected. A successful attack on a Grid workstation is likely to spread rapidly, using the Grid itself, to all the Grid servers that the system and its users have access to. In fact most of the major security incidents affecting Grid systems have started with the compromise of a client workstation, using a vulnerability in a non-Grid service, which gave the intruders access to Grid identity certificates. Once an intruder has a certificate, he can gain access for malicious purposes to any Grid resource the original owner can use, and can often use this to steal the credentials of other users of the same clients, resources and networks. In most cases the service that was originally attacked could have been protected if well understood network controls had been in place. A review of this kind of incident is published in [Skow].

The security measures for a network must be considered as a whole so that measures are consistent. In many cases the presence of network controls may reduce the urgency of managing individual systems. If a particular protocol is blocked by a router then it may be possible to treat patches for vulnerabilities as part of planned maintenance rather than requiring emergency action; a firewall may create a perimeter within which more experimental software can be used. Conversely, any reduction in the protection provided by a firewall will require correspondingly greater efforts to keep the systems within it secure. The deployment of Grid systems and software must therefore be considered alongside the network control measures to be used, both to protect the Grid systems and data and to ensure that changes to existing measures do not expose other systems to greater risk.
