

Network design for Grids

Networks are a key part of the Grid vision, so network design must be considered as part of a Grid deployment. The physical networks to which Grid systems are connected, the allocation of IP addresses and the use of appropriate network controls can all have significant benefits for the performance and security of a Grid. Conversely, if these issues are not included in the early planning they can cause endless problems.


Physical Design

As discussed above, Grid systems will often give rise to high bandwidth flows using complex network protocols. For such systems performance, security and deployment requirements all point to simple network designs being the most effective. Simple network paths should reduce delays and bottlenecks, reduce the opportunities for security vulnerabilities to occur, and reduce the number of network devices that have to be specified or configured to handle Grid protocols.

High bandwidth flows are likely to see the greatest benefit from simplification, and any deployment should try to predict where these are most likely to occur and design its network to suit. On some Grids the largest flows may be the transfer of data across the WAN (Wide Area Network) between replicated servers at different sites; on others they may be the capture of raw information from an instrument connected to the same LAN. Grid services may have their largest aggregate flows coming from multiple clients at other sites. In each case the network segments that carry the largest flows should be reviewed and if any are likely to act as a bottleneck, either because of their low capacity or high existing load, then alternatives should be considered. This may involve rearranging existing network connections or providing new links. At the same time the active network devices – switches, routers, firewalls, etc. – in the path of large traffic flows should be identified to ensure that their performance is sufficient to handle the traffic. If there are a large number of devices on the path then it may be better to change the network topology to reduce the number of devices that need to be traversed. However, a certain number of control devices should be used to exclude unwanted traffic, both to remove competition for bandwidth and to reduce the exposure of Grid computers to attack across the network.

Where the largest flows between clients and servers occur within the site, they will often be naturally located close together and it will be straightforward to connect them to the same or nearby network segments. Where a Grid service is offered to external and internal clients, or is to communicate with other external servers, a common design with security and performance benefits is to provide an additional interface to the site router (or switch), with the Grid and site networks each connected to their own interface via their own firewall (see Figure 1 overleaf). Providing separate firewalls simplifies the configuration of security, performance and policy controls. If cost is an issue then the functions of router and firewalls can be

and site networks




[1]

The process of simplification is likely to result in Grid servers and clients naturally being clustered together on local networks, with high-performance paths across the site or national networks between the clusters. Although clusters of high-traffic systems might appear to run the risk of creating network hotspots, it will usually be easier to provide and manage high performance, end-to-end networks to a few groups of systems, rather than checking the performance of every network component on all the paths to systems scattered randomly across the network. Clustering Grid systems also increases the opportunities for separating Grid traffic from other uses of the network to reduce the risk of performance and security problems arising from interference between the two types of traffic.

Where pairs or groups of Grid systems will be working intensively together, it may be possible to take this separation to its extreme by providing dedicated network links between the Grid systems rather than sharing the general-purpose network. This may apply within an organisation as well as within or between countries. For example, a number of sites in the high energy physics community have obtained dedicated network links to CERN. However it is important to ensure that such links do not become an alternative route for general-purpose traffic – the normal behaviour of Internet routing protocols in seeking out uncongested routes will need to be disabled – and that they do not provide a way for hostile traffic to circumvent security or policy controls installed on the general-purpose network. Wherever a system is connected both to a dedicated link and a general purpose network there is a risk that it will act as a bridge for the spread of security problems.

As shown in Figure 2 below, Grid systems with dedicated wide area links may need to implement controls between themselves and their local network to protect the local network



[2]

These controls must have the same net effect as those implemented by the site firewall, but splitting the controls across two points, one on the dedicated link and one on the internal link between the Grid system and its local area network, should allow the Grid system to achieve the appropriate balance between connectivity and protection.

Address Allocation

Simplification at the network level should be accompanied by simplification of IP addresses, where it has further benefits.

If possible, clusters of Grid systems with similar functions should be given network addresses from the same address range, and no other systems should be allocated addresses within

that range. For example, if the range 192.168.1.16 to 192.168.1.31 contains only Grid servers, then this range can be referred to in the CIDR (Classless Internet Domain Routing) form 192.168.1.16/28. Many routers and firewalls allow CIDR addresses to be used in configuration files, and may even convert them to hardware routing tables, so controls expressed using this format can have both management and performance benefits. Grid services that configure their services within CIDR ranges make it much easier for user and peer sites to configure their network and system controls. The use of CIDR ranges can also make it easier to implement virtual organisations in network terms. Exchanging well-defined address ranges is the key to the idea of the trusted 'clique grid' discussed in [Hillier].

It may also be possible to use CIDR ranges to implement VLANs (Virtual Local Area Networks) in hardware. At least this will allow Grid traffic to be restricted to only those network segments that it needs to flow along. With some network switches and routers it may also be possible to reserve a fixed amount of bandwidth for the Grid traffic.

Control Points

As discussed above, Grid systems are likely to benefit from simple network paths, but this does not mean that all routers and firewalls should be eliminated. Wherever networks are shared with other non-Grid systems or users, or are connected to such shared networks, it is strongly recommended that some network devices (such as routers or firewalls) be retained to act as control points for traffic flows. These are essential to protect Grid traffic flows from contention as well as to protect Grid systems (both clients and servers) from hostile or accidental traffic. The complete removal of network level controls should be considered only where a network and all its connected systems are dedicated solely to Grid work. Even then, to do so means that a single failure of security on any system is likely to compromise the integrity of the entire network. In particular, Grid systems should never have direct, uncontrolled access from the public Internet, as this will seriously harm their reliability.

Ideally, Grid systems should be protected by firewalls or routers that only allow legitimate Grid traffic, but the performance requirements and complexity of protocols may mean that this cannot be achieved. Even if full protocol level filtering cannot be achieved, simply limiting the IP addresses, or ranges, between which traffic can flow will greatly reduce the exposure to threats. High performance routers are likely to process lists of IP addresses in hardware so there should be little or no performance reduction in blocking traffic from unknown addresses. Grid servers will usually occupy well known IP addresses or ranges so server to server traffic can usually be protected in this way. Grid clients are likely to be more numerous and have less predictable IP addresses, but their requirement and ability to handle high-volume flows is also likely to be less. This means there are a number of possibilities to give them controlled access to Grid servers. Users within a site may be required to authenticate themselves before being allocated an IP address from a trusted range. At present this is likely to require a double log-on, one to prove entitlement to an address and one to gain access to the Grid servers, but work is in progress on the investigation of single sign-on systems that would use the same Grid credentials to gain access both to the network and to servers. Alternatively, local and remote users can be supported by using a VPN to bring traffic from a client to a central VPN concentrator, from where it can be routed onto a trusted Grid network or address range. Further details and references to systems of these types can be found in Section 4.2 on Tunnelling.

Control points are needed to protect Grid systems from the rest of the network, but protection

in the reverse direction may also be required. Grid systems typically have access to large computing resources, so any compromise here represents a considerable threat to other systems on the network. Network level controls should therefore be available to contain the spread or consequences of any security incidents within the Grid, and between the Grid and the rest of the network. Since Grid systems often have remote users from other organisations, it may also be necessary for licensing or privacy reasons to prevent those users from accessing other data, software or systems on the local network. Since resources often rely on IP address restrictions to distinguish local (licensed) and remote (unlicensed) users, careful allocation of IP addresses will be needed to avoid Grid systems being a gateway for unintended leaks of information.

Other Network Issues

Performance

Even on a simplified network there are a large number of technical components that contribute to the user's perception of end to end network performance. These range from congestion on backbone networks or routers to poor quality wiring or incorrect hardware or software configuration at the desktop. Any of these may be critical to a high-performance application while being imperceptible in normal operations. The PERT (Performance Enhancement and Response Team) is a project to develop techniques for analysing and resolving performance problems, and aims to produce both debugging strategies and recommended configurations for systems involved in high performance network applications [PERT]. The majority of problems reported to the PERT in early 2004 were traced to inappropriate settings of network cards or TCP buffer sizes.

Source URL: <https://community-stg.jisc.ac.uk/library/janet-services-documentation/network-design-grids>

Links

[1] <http://community.ja.net/system/files/images/tg-deployinggrids-01.jpg>

[2] <http://community.ja.net/system/files/images/tg-deployinggrids-02.jpg>