

The use of firewalls in the U.K. e-Science grid:
ETF Level 2 and beyond.

DRAFT VERSION 0.1 Oct 24th 2002. Author Jon Hillier (OeSC).

Introduction

This document discusses the necessity, configuration and maintenance of firewalls with the expected growth and evolution of the U.K. e-Science community. The problems outlined, and solutions presented, are based in part upon previous work performed by member institutions of the U.K. Engineering Task Force (ETF) [Baker01, Booth], in response to a survey of current and anticipated grid use and network topology in said institutions and through discussion with members of the ETF.

ETF Level 2

The U.K. e-science grid is in the process of evolving from its level one state (a prototype grid, with the emphasis on testing new procedures and achieving a consistent level of expertise and performance throughout the nationwide e-science centres) to its level two state.

Whilst the exact definition of the level two grid is currently purely a working definition, with particular reference to the flux of members it can be defined as follows: a heterogeneous collection of Globus-enabled clients and servers with a well-defined collection of tools and procedures for day-to-day user management, accounting, resource information sharing and infrastructure monitoring. In addition to this, the collection of resources available through the grid (gatekeepers and attached resources) is expected to function efficiently as a single entity. The number of servers within the level two grid is expected to remain relatively constant, and the flux of new clients to be relatively low.

The level three grid, which is expected to replace the level two grid after ** 200*, may be more dynamic, with clients, servers and resources connecting to the grid backbone as and when necessary. It may also be more OGSA/web service based.

The problem with Globus

For those unfamiliar with the Globus project and toolkit, the Globus website and the U.K. Grid Support web sites are the best starting points¹. The Globus toolkit is best described as a “bag of services” – tools and components for developing a grid. The Globus tools act not only as a form of black-box for access to existing job management and scheduling systems via secure PKI (x509) certificate authentication, but also provide a secure method of resource status querying and secure file transfer facilities. Globus has recently released a document with detailed information on firewall requirements for successful operation of the Globus Toolkit [VonWelch02] – this document is essential reading for anyone considering setting up or modifying a firewall for use with Globus!

¹ www.globus.org and www.grid-support.ac.uk respectively.

Globus uses three primary ports in the IANA registered ports range, these are:

2119 *Globus Gatekeeper* (tcp)

2135 *Globus MDS* (tcp/udp)

2811 *GridFTP* (tcp)

The majority of the Globus tools utilise a passive ftp-like port usage², with tcp returns on the IANA registered/dynamic/private port ranges, 1024-65535. By default Globus will use any ports within this range, starting with the lowest available port, but it is possible to restrict the port usage to a server-specified range using the GLOBUS_TCP_PORT_RANGE environment variable. This port range must be large enough to cope with many simultaneous connections, whilst not presenting a serious security risk if these ports are open on a firewall. This problem, together with recommendations on suitable port ranges, is discussed in more detail in a forthcoming document by A. Richards and D. Hanlon at the CLRC Daresbury Laboratory [Richards02].

Whilst the three primary ports should not present any particular difficulties with firewall configuration – as most sites are reasonably comfortable with allowing similar services (https, ftp, ssh), the tcp return ports could present a major security risk.

Finally, the construction of the Globus toolkit – as wrapped and modified versions of existing open source software (openssl, wuftp etc.), leads to a lag in the application of patches should a security flaw be found in the original open source software.

Web Services

The Grid community, particularly those currently using the Globus software are expected to migrate to the use of web service based grid computing with the release of the next generation Globus Toolkit (GT3). A discussion of GT3 is given at <http://www.globus.org/toolkit/gt3-factsheet.html>. Use of web services considerably reduces the difficulty of firewalling grid resources and gatekeepers as it involves the use of far fewer ports (with Universal Resource Identifiers, URIs, effectively replacing ports), and importantly, no dynamic tcp return ports.

The ETF Firewall Questionnaire

- 1) Where is(are) your gatekeeper(s) with respect to your firewall?: Inside a domain firewall, in a DMZ, not firewalled?
- 2) Where would you ideally like to have a firewall in your system? Not on the gatekeeper itself?
- 3) What is the degree of firewalling between the gatekeeper and the resources it uses? Does this differ from the domain firewall? Does the degree of firewall access to any Grid/Institution shared facility differ between grid users and local non-grid users?

² Detailed information on port usage during a typical Globus interaction is presented in Von Welch's Globus document [VonWelch02], and also in Baker et al. [Baker01].

- 4) What is the expected degree of client use of a gatekeeper versus VO gatekeeper-gatekeeper connections? Is client access important at your site?
- 5) What control do you have over your local firewall settings? Can you make changes quickly?
- 6) Would you be prepared to allow access to all UK e-Science participants through your firewall or would you insist on selective VO-based access only?

Possible firewall configurations

The following list (taken from the ETF Level 2 document [Allan02]), summarises the possible firewall configurations currently under consideration which will allow Globus-based communication between computers at different institutions.

- a) Do not use a firewall.
- b) Use a separate subnet with no firewall.
- c) Use a firewall but open all ports to a matrix of machines.
- d) Use a firewall but open a range of ports to either all machines or a matrix of machines (the Globus recommendation).
- e) Use the Nexus proxy.
- f) Use some form of Dynamic firewall.
- g) Use a multi-institution Virtual Private Network and selective firewall.

Suitability of firewall configurations for Level 2 grid

Current firewall configurations at the various ETF institutions range from no firewalls at all to more complicated systems with site-wide firewalls, different security domains within the site firewall and even software firewalls running on the gatekeepers themselves.

The common theme for all sites, whether they currently have a relaxed site firewall policy or a more restrictive one, was that their firewalls were expected to become considerably more restrictive in the near future.

So, why do we need a firewall at all? Why not simply put the gatekeepers and resources on a separate subnet with no firewall, or have no site firewall at all? The firewall is the first line of defence in applying access controls to services, programs or files on networks both internal to the firewall and external. Aside from the need to restrict port usage by certain file sharing programs (such as Napster and its many clones) and other restricted services, most institutions rely on their firewall for protection against malicious crackers and the plethora of viruses and Trojans ever present on the internet. Apart from attacks such as Denial of Service (DoS), the main system vulnerabilities are the result of "poor" coding.

Theoretically, if a system contained only 100% secure, tested programs, with access restrictions applied where required, then a firewall would only be necessary to protect against DoS-type attacks. Unfortunately, this is not the case in the real world. In addition to software most often not being 100% secure, users have a nasty habit of expecting email, web access, ftp access, remote X windows and all manner of other potentially insecure software to be available on their system. As and when a problem arises with a particular piece of software or service then action must be taken quickly

to close any security hole. It is far easier and quicker to be able to shut off access to a particular service, for example incoming http, at a site firewall than trust local users and system administrators to solve the problem on each individual machine quickly – patches and software upgrades can then be applied before re-opening/starting affected services.

The use of dynamically assigned tcp return ports, together with the potential lag in software security updates makes Globus particularly vulnerable unless protected by a well configured firewall. Unfortunately the first reason, one of the very reasons that Globus *should* be firewalled, also makes it extremely difficult to firewall well.

It should be noted at this stage that security is of particular importance in a grid environment – the resources accessible via the grid are often world class supercomputing facilities, and, as use of the grid increases, the sensitivity of data stored and used in a grid environment will increase considerably.

As well as the potential security risks to computers within a particular institution, if the institution is involved in collaboration with other institutions then the firewalls of the member institutions may be more relaxed towards the member machines. This could lead to a house of cards effect with a single compromise at an insecure site leading to compromises at more secure sites within a trusting framework. If any form of grid, containing firewalls restricting access on the basis of IP address is implemented then *all* the member institutions must have effective firewalls – the system will be as weak as its weakest member. This will be discussed further later in this document with respect to points c) and d) above.

The answers from the ETF members regarding the current and future state of the majority of the firewalls already in use, together with the security reasons already discussed, effectively discount the possibility of choices a) and b) from the list above.

The Nexus proxy [Nexus] (option e) was developed to overcome the problem of communicating through a firewall to nodes within the firewalled domain possessing only locally-recognisable IP addresses. Baker et al. [Baker01] discuss the use of this proxy in considerable detail, and whilst concluding that it could form a valuable method of firewall tunnelling in the future, were unable to get it to work successfully and consequently recommended the method of opening specific ports on the firewall. The Nexus proxy also requires changes to the Globus source code, was developed to function with version one of the Globus toolkit and, importantly, requires a computer external to the site firewall with which to form the tunnel. For these reasons use of the Nexus proxy is unlikely in the future U.K. grid and option e) is discounted from the list.

All four remaining options (c, d, f, g) allow the use of firewalls, albeit with differing levels of flexibility.

Two options (c, d), those of firewalls which have a default deny policy to all computers apart from a specified subset, are closely related and will be considered together. For a system such as the level two grid, a viable, although slightly inflexible and possibly not scaleable, solution is to *hard-wire* the IP addresses of the participating machines into a firewall. Whilst in principle all ports (option c) could be opened to these specified machines (or networks, sites etc.) for reasons of damage limitation and further security, quite apart from the fact that not all ports actually need to be open, a more suitable solution is to open a subset of ports to a selection of

computers (option d). The exact firewall rules required (for a software firewall – IPChains in this case) are described in [Baker01] and will not be mentioned here. This method would allow the satisfactory operation of the level two grid. There are (of course!) problems still remaining: specifically entering the IP address (or, in some cases a range of IP addresses) into each institutional firewall may introduce a degree of inertia to the grid, as the number of participants in the grid increases the firewall rules will grow more complex/lengthy and the firewalls may eventually form bottlenecks, and the “roaming” use of IP addresses would be ruled out. Possible solutions to many of these problems, together with a more detailed plan for the implementation of such a system will be discussed later in this document.

As a solution to the problem of “roaming” users, the same solution applied in areas of networking other than specifically the grid can be used – namely VPNs. VPNs using systems such as IPsec are becoming ever more popular as internet use becomes less office-based and more mobile. If a U.K. wide e-Science VPN, hosted at a particular institution (such as the National e-Science Centre for example), was implemented then local firewalls could be configured to pass a specific range of ports from a specific range of IP addresses. As IPsec is effectively invisible to the end-user, and can use x509 encryption and hence Globus-compatible certificates, it would make the ideal choice for an e-Science VPN. Additionally, VPNs are commonly used as a way “through” restrictive firewalls. Problems with use of a VPN exist however: the implementation of a multi-institutional e-Science VPN is non-trivial, the security of many sites would depend on the security of the site hosting the VPN, there may be network bandwidth problems and there *may* be scaling problems (with address assignment) as the grid grows. However, a multi-institutional VPN is a possible solution to the firewall problems faced by the level two grid.

The dynamic firewall method (f) works in a similar way to a VPN, although technically it is considerably simpler to implement. The dynamic firewall works by first assuming that access to a gatekeeper should only be granted upon presentation of a valid certificate to a gatekeeper. The presence of a valid certificate does not necessarily mean that jobs can be staged to the gatekeeper – for this to happen certain certificate details (the Distinguished Name, DN) have to be mapped to a local user account on the gatekeeper. The prototype dynamic firewall script checks for both a valid certificate and the ability to actually run Globus jobs on the gatekeeper. Initially only a single port, that of the Globus gatekeeper, has to be open on the firewall. A client wishing to connect sends a Globus “ping”,³ to the gatekeeper. The dynamic firewall script monitors the success of this ping and, if the ping is successful, further ports (MDS, GridFTP and any others – including the return ports) are then opened to the IP address from which the ping originated. To add further security, the access times out after an admin-specified period and the firewall then returns to its original state. This method also avoids the need for manual intervention in adding or removing new firewall rules. Unfortunately, not only is the system currently (and foreseeably) only available for the Linux operating system, but it also requires a daemon with root privileges making direct changes to the firewall as a result of a remote request. This, together with the need for the firewall to be a software firewall such as IPTables running *directly* on the gatekeeper itself, made the dynamic firewall unacceptable to the majority of the ETF member institutions. The possible use of such a system as an

³ globusrun -a -r <gatekeeper>

“emergency” method of access to a gatekeeper will be discussed later in this document.

The chosen solutions, which are reliable and easily implemented are:

- d) Open a range of ports to a trusted selection of IP addresses.
- g) Use a multi-institutional VPN, together with a selective firewall.

There now follows a more detailed discussion of the implementation of these solutions.

The *clique* Grid – a trust based solution

If considering restricting access to Globus gatekeepers and resources based on IP address then the obvious choice is to maintain a centralised list of allowed, trusted, IP addresses. This leads to the following:

- 1) A secure site must exist on which to store the IP addresses – this site is potentially responsible for the security of *all* sites participating in U.K. e-Science.
- 2) A secure method of adding new IP addresses must be implemented.
- 3) A secure method of modifying or removing IP addresses must be implemented.
- 4) A fast and secure method of propagating additions or removals to the database to firewall administrators has to be found.
- 5) Any interfaces created for use in the level two grid must be compatible with future grid developments.

Responses from the ETF also indicated that many institutions would prefer to restrict access through their firewalls to other members of Virtual Organisations (VOs - projects etc.) that they are specifically involved in. This means adding further layers of complexity to the trusted host database.

The proposed solution is to create a web-based database accessible using the GridSite⁴ [McNab02] software. This is an existing secure web server developed for use in the GridPP community and already in use by other areas of the U.K. e-Science community (The ETF web pages for example). Importantly, GridSite can use modified Globus certificates to allow read/write access (via Access Control Lists – ACLs) to specific parts of a site. This will allow VO granularity – VO members will only be able to access IP address information pertinent to them.

IP address information would be fed into the GridSite accessible pages either straight from VO management software itself (such as LeSC’s VOM) or via an XML-based grid user/resource database (currently under development at Daresbury). The web pages should be XML-based and fed via a GSI-based web service from the database. For security reasons it is suggested that all additions to the database are moderated by a nominated individual – a VO project leader for example. Regardless of the form that this back-end takes, the important factor is that the GridSite method for accessing the data is set up in such a way as to be able to cope with future grid developments as transparently as possible.

⁴ www.gridsite.org

Changes to the population of a VO would have to be propagated to the local firewall administrators at the member sites as quickly as possible – perhaps via email, or, if VO-management software is closely linked to the relevant databases, a warning could appear within the software itself. For some sites quick changes to their current firewalls are difficult to make – this is a problem which will have to be resolved within the institutions themselves, especially those institutions which foresee their firewalls becoming more restrictive in the near future. It may well be that as new members of a VO apply for their certificates, the appropriate information for firewall changes is also submitted – changes to any firewalls could then be made in preparation for the successful issue of a certificate.

This trusted host (clique) database solution is a good solution for the static level two grid, but it cannot address the problem of dynamic, roaming IP addresses. To allow for this the use of VPNs is required.

The multi-institutional single VPN solution

By setting up a repository of VPN available IP addresses at a single institution (NeSC for example), roaming users could then authenticate using their Globus certificates from anywhere in the country. Firewalls at institutions around the country could then be configured to allow access from any of these addresses. This method could be used in conjunction with the trusted host database solution, computers requiring constant connectivity to the grid (gatekeepers, resources and some clients) would have their IP address details entered into the database, and consequently be hard-wired into the relevant firewalls. More dynamic machines would make use of the VPN to obtain IP addresses which were allowed through the firewalls.

There are problems with using a single institution to host the VPN. The maxim of never putting all your eggs in one basket applies – a single institution may be a weak spot – should the network fail at that institution then all roaming access to the grid would be stopped. Also, the security of the VPN servers themselves is of critical importance – they would also have to be behind a firewall. Finally, a VPN at a single institution may be limited in terms of available IP addresses and also bandwidth. A partial solution to the problems is to host a VPN at each participating institution.

Per-institution VPN and host database hybrid solution

A combination of an per-institution VPN (*not* a multi-institutional VPN) and the host database method would appear to form a suitable solution for the near to mid-term growth of the grid.

Hosting a VPN service at each individual institution increases the reliability of the service, places the onus on VPN security with each institution and can potentially cope with many more users than a single U.K. wide e-Science VPN.

The host database method would work as described previously, with VPN acting as a pool of pre-registered IP addresses at each site. New users with static IP addresses not yet registered and updated through the host database system could use the VPN as a temporary solution whilst more mobile users could make continual use of the VPN.

Totally static machines – gatekeepers, resources and some clients, would have their IP addresses hard-wired into any firewall necessary as before. It is difficult to find fault with this system!

Emergency access!

Although in no way a suitable large scale solution, systems such as the dynamic firewall may be of use for granting short-term access to gatekeepers and resources at a particular site – during demonstrations at a conference for example. If use of software firewalls on a primary gatekeeper is out of the question then it may be possible to implement a low usage emergency gatekeeper with access to the same resources as the primary gatekeeper.

Recommendations

For a short term solution to firewall configuration and maintenance in the level two grid the recommended solution is the trusted host database solution. For the mid to long term evolution of the grid, complete with roaming users, the per-institutional VPN and trusted host database hybrid solution is recommended. Should sites insist on maintaining selective firewalls when web services are in use then the recommended GT2 based solutions are still perfectly valid for use with both GT3 and other web services.

Acknowledgements

The author wishes to thank all the members of the ETF for their patience and indulgence in responding to the firewall questionnaire. In particular thanks to Steven Newhouse, Rob Allan and Stephen Booth for valuable discussion and input. Thanks also to Matthew Dovey for patiently explaining VPNs to me..

Notes:

ETF member institutions participating in the questionnaire:
Welsh e-Science Centre, Cardiff.
Cambridge e-Science Centre, Cambridge University.
London e-Science Centre, Imperial College.
CLRC, Rutherford Appleton Laboratory.
CLRC, Daresbury Laboratory.
Oxford e-Science Centre, Oxford University.
Southampton e-Science Centre, University of Southampton.
e-Science North-West, Manchester.

References

Baker01: Mark Baker, Hong Ong and Garry Smith, “A report on experiences operating the Globus toolkit through a firewall”, September 2001. Available from <http://esc.dl.ac.uk/Papers/firewalls/globus-firewall-experiences.pdf> .

Nexus: The Nexus proxy documentation, available from <http://www.apgrid.org/download.html> .

McNab02: Andrew McNab, “User, Admin and Installation Guide for GridSite version 0.2.1”, February 2002, Available from <http://www.gridpp.ac.uk/gridsite/guide.html> .

Booth: Stephen Booth, “ETF Grid Firewall Recommendations”, Available from <http://www.grid-support.ac.uk/etf/firewalls/Firewalls.html>.

Welch02: Von Welch, “Globus Toolkit Firewall requirements”, August 02, Available from <http://www.globus.org/security/v2.0/Globus%20Firewall%20Requirements-0.3.pdf>.

Allan02: Rob Allan and David Boyd, “The U.K. e-Science Grid, Level 2 Deployment Plan v0.4”, October 02, Not generally released – possibly available from the authors on request?

Richards02: Andrew Richards, Rob Allan and Daniel Hanlon, “Globus Toolkit Firewall Port Selection”, October 02, available from the grid support web site- <http://www.grid-support.ac.uk/etf/firewalls/index.html>