



Fachgebiet 3-5 – Entwicklung und Einsatz von
Firewallkonzepten in Grid-Umgebungen

Recommendations for Static Firewall Configuration in D-Grid

Autoren

Gian Luca Volpato (RRZN, Leibniz Universität Hannover)

Christian Grimm (RRZN, Leibniz Universität Hannover)

Das diesem Bericht zugrundeliegende Vorhaben wurde mit Mitteln des Bundesministeriums für Bildung und Forschung unter dem Förderkennzeichen 01AK800B gefördert. Die Verantwortung für den Inhalt dieser Veröffentlichung liegt bei den Autoren.

Index

1	INTRODUCTION.....	5
2	GLOBUS GT4	6
2.1	Server.....	6
3	LCG/GLITE.....	8
3.1	Server.....	8
4	UNICORE	10
4.1	Server.....	10
5	SRM / DCACHE	12
5.1	Server.....	12
6	OGSA-DAI	14
6.1	Server.....	14
7	REGISTRATION OF GRID SERVICES	15
	REFERENCES	16
	APPENDIX A: CONTROLLING THE EPHEMERAL PORT RANGE IN GT4 [WEL06]	17
	APPENDIX B: CONTROLLING THE EPHEMERAL PORT RANGE IN LCG/GLITE AND SRM/DCACHE [WEL06]	18
	APPENDIX C: FIREWALL FILTERING RULES CONFIGURATION FOR GLOBUS TOOLKIT 4.....	20
	APPENDIX D: FIREWALL FILTERING RULES CONFIGURATION FOR LCG/GLITE	21
	APPENDIX E: FIREWALL FILTERING RULES CONFIGURATION FOR UNICORE.....	22
	APPENDIX F: FIREWALL FILTERING RULES CONFIGURATION FOR SRM/DCACHE AND OGSA- DAI	23

List of abbreviations

BDII	Berkeley Database Information Index
CE	Computing Element
DAI	Data Access and Integration
DMZ	DeMilitarized Zone
DN	Distinguished Name
GIIS	Grid Information Index Server
GRAM	Grid Resource Acquisition and Management
GRIS	Grid Resource Information Server
GSI	Grid Security Infrastructure
GT4	Globus Toolkit 4
LCG	LHC Computing Grid
LHC	Large Hadron Collider
MDS	Monitoring and Discovery System
MON	Monitoring box
NJS	Network Job Supervisor
OGSA	Open Grid Services Architecture
RFT	Reliable File Transfer
RGMA	Relational Grid Monitoring Architecture
SE	Storage Element
SRM	Storage Resource Management
WS	Web Service
WSRF	Web Services Resource Framework

1 Introduction

The document “Static Firewall Configuration for Grid Middleware” released by DGI FG3-5 in January 2006 analyzed the interaction between Grid middleware and firewalls. Two important results were achieved: a list of network ports and associated transport protocols used by different Grid services and a set of firewall configurations (IP filtering rules) that allow seamless communications among these services.

This document refines the previous work introducing recommendations for middleware deployment and configuration, and defining one single port range to be adopted by all resource providers within D-Grid. The immediate benefit will be a simplified firewall configuration.

Although a minimum number of ports needs to be permanently open, it is possible to restrict the incoming connections to a specific sets of IP addresses (or IP sub-networks). This implies the registration of all service hosts deployed by D-Grid in a common, shared, directory service.

The following sections cover the service releases for:

- Computing resources:
 1. Globus Toolkit 4
 2. LCG/gLite
 3. Unicore
- Data management:
 1. SRM/dCache
 2. OGSA-DAI

2 Globus GT4

Resource providers deploying Globus Toolkit 4 are recommended to use the following ports in their servers' configuration.

Host	Service	Port (TCP)	Modification to default settings
GRAM + MDS + RFT + GSI-SSH	GRAM	2119	No
		Range 20000-25000	Yes
	WS-GRAM	8443	No
		Range 20000-25000	Yes
	WS-MDS	8443	No
	GridFTP	2811	No
Range 20000-25000		Yes	
RFT	8443	No	
GSI-SSH	2222	Yes	

Table 1: GT4 port configuration

With the above recommendations in place the firewall configuration for GT4 becomes as described in Appendix C.

2.1 Server

Resource providers deploying GT4 in D-Grid are advised to configure the middleware services as depicted in Figure 1, i.e.:

1. Install GRAM, MDS, RFT and GSI-SSH on one single host.

Solid lines represent connections between the involved parties. Communication with external entities is displayed on the right side of the figure, where ports to be enabled for incoming connections through the firewall are also indicated.

Resource Provider

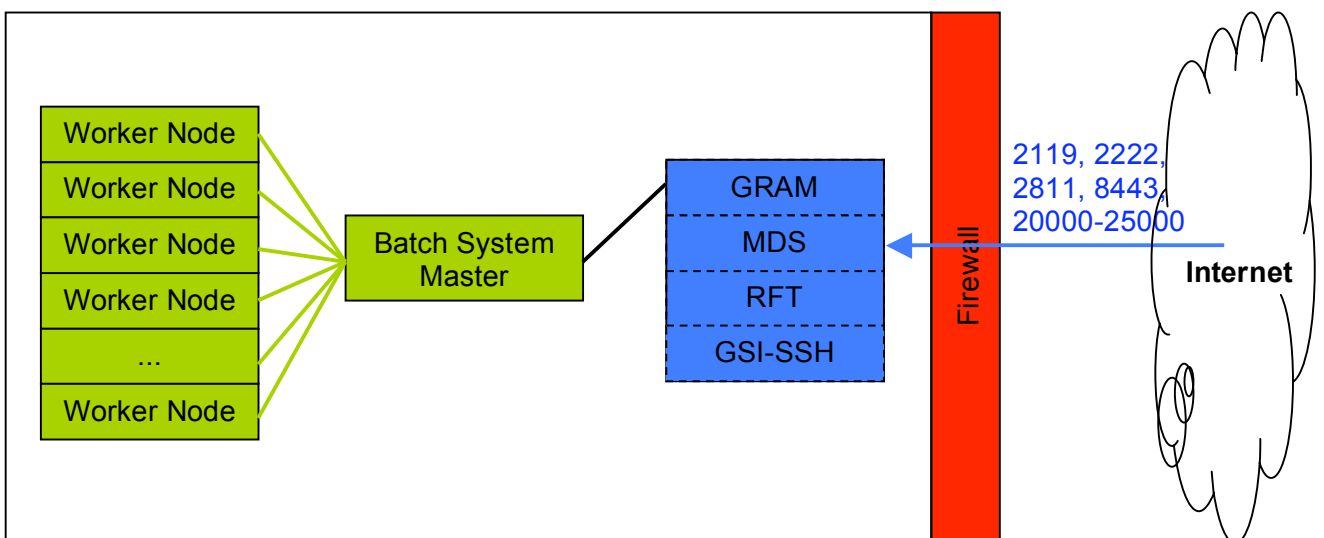


Figure 1: GT4 recommended deployment at resource providers

To adjust the standard Globus installation to the specific D-Grid recommendations apply the following change:

- Set the *controllable ephemeral port range* to the interval 20000 - 25000.
- Set the port used by the GSI-SSH server to 2222/tcp.

Controllable Ephemeral Port Range

Preliminary operational experience with the Globus Toolkit shows that approximately 20 ports per expected simultaneous user on a given host should be provided. Given an average amount of 250 simultaneous active D-Grid users, the controllable ephemeral port range is in the order of 5000 ports.

The environment variable `GLOBUS_TCP_PORT_RANGE` and the Java system property `org.globus.tcp.port.range` restrict the ephemeral ports to a specific range. This range should be formatted as *min,max* (comma separated values) and it should be configured in all Globus Toolkit 4 affected components through the following procedure:

1. Configuration of C libraries, client applications and services.
2. Configuration of pre-WS-GRAM.
3. Configuration of Java libraries, client applications and services.
4. Configuration of GridFTP.

Appendix A describes with details how to set the port range to the interval 20000 – 25000.

GSI-SSH

The GSI-SSH server runs by default on port 22/tcp. The same port is used also for normal SSH connection. In order to separate the traffic of the two services it is recommend using port 2222/tcp for the GSI-SSH server.

Assuming a GSI-SSH server is already installed and running, the following steps change its port to the recommended value 2222/tcp:

1. Edit the file `$GLOBUS_LOCATION/etc/ssh/sshd_config` and set
`Port=2222`
2. Edit the file `/etc/services` and insert a new line
`gsisshd 2222/tcp`
3. Restart the `gsisshd` daemon.

3 LCG/gLite

Resource providers deploying LCG/gLite are recommended to use the following ports in their servers' and clients' configurations.

Host	Service	Port	Modification to default settings
LCG CE + MON + Site-BDII	GRAM	2119	No
		Range 20000-25000	Yes
	GRIS	2135	No
	Extended GRIS	2136	No
	BDII	2170	No
	GridFTP	2811	No
Range 20000-25000		Yes	
MON-RGMA	8080	No	
	8088	No	
	8443	No	
Worker Node	GridFTP	Range 20000-25000	Yes

Table 2: LCG/gLite port configuration

With the above recommendations in place the firewall configuration for LCG/gLite becomes as described in Appendix D.

3.1 Server

Resource providers deploying LCG/gLite in D-Grid are advised to configure the middleware services as depicted in Figure 2, i.e.:

1. Install LCG CE, MON and Site-BDII on one single host.
2. Install the Worker Node package on all batch system nodes.

Solid lines represent connections between the involved parties. Communication with external entities is displayed on the right side of the figure, where ports to be enabled for incoming connections through the firewall are also indicated.

To adjust the standard LCG/gLite installation to the specific D-Grid recommendations apply the following change:

- Set the *controllable ephemeral port range* to the interval 20000 - 25000.

Resource Provider

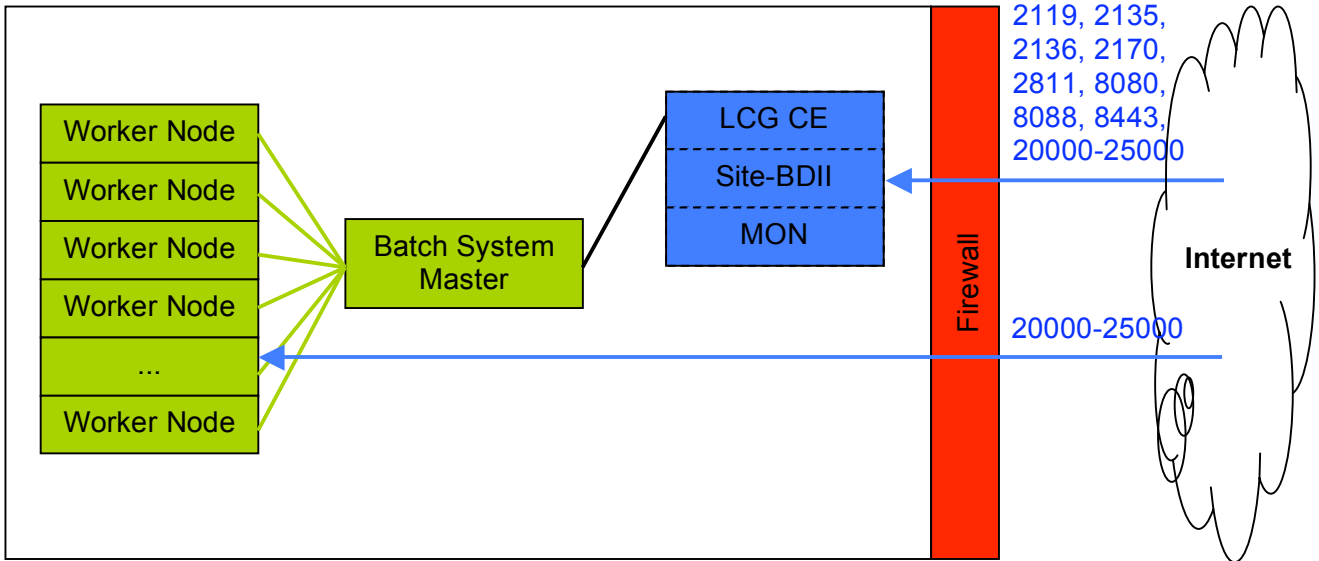


Figure 2: LCG/gLite recommended deployment at resource providers

Controllable Ephemeral Port Range

Preliminary operational experience with the Globus Toolkit shows that approximately 20 ports per expected simultaneous user on a given host should be provided. Given an average amount of 250 simultaneous active D-Grid users, the controllable ephemeral port range is in the order of 5000 ports.

The environment variable `GLOBUS_TCP_PORT_RANGE` restricts the ephemeral ports to a specific range. The value of this variable should be formatted as *min,max* (comma separated values). Setting this environment variable in all LCG/gLite affected components is achieved through the following procedure:

1. Configuration of GT libraries.
2. Configuration of GRAM Gatekeeper and JobManager.
3. Configuration of GridFTP.

Appendix B describes with details how to set the port range to the interval 20000 – 25000.

4 UNICORE

Resource providers deploying Unicore are recommended to use the following ports in their servers' configurations.

Host	Service	Port	Modification to default settings
NJS + TSI + UADB	NJS	1128	Yes

Table 3: Unicore port configuration

4.1 Server

Resource providers deploying Unicore in D-Grid are advised to configure the middleware services as depicted in Figure 3, i.e.:

1. Install NJS, TSI and UADB on one single host.

Solid lines represent connections between the involved parties. Communication with external entities is displayed on the right side of the figure, where ports to be enabled for incoming connections through the firewalls are also indicated.

It can be noticed that incoming connections directed to the NJS should be allowed only when initiated by the official D-Grid Unicore Gateways provided by Forschungszentrum Jülich and Forschungszentrum Karlsruhe. These Gateways are the single point of contacts for all external Unicore connections and act as a filter in front of the NJS.

Resource Provider

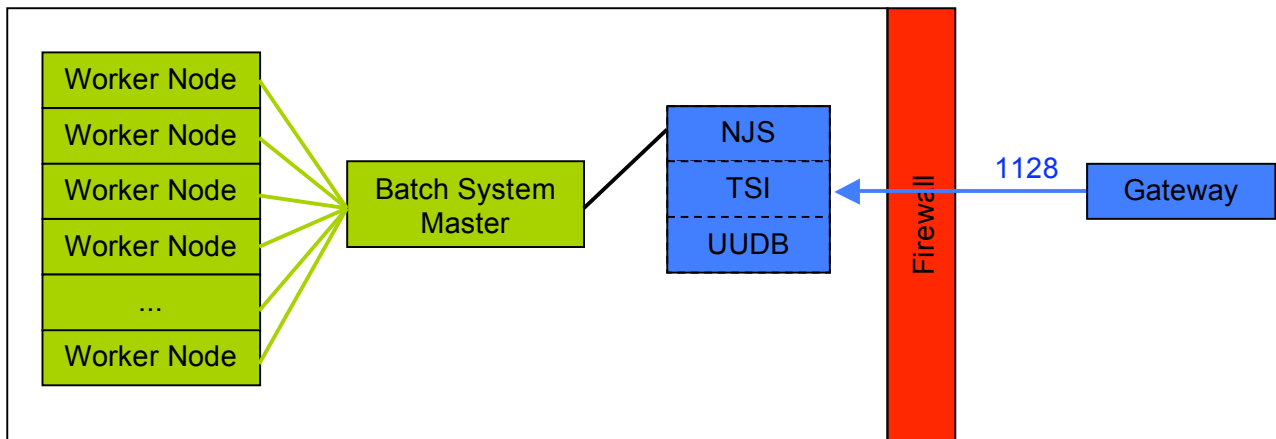


Figure 3: Unicore recommended deployment at a Resource Centre

To adjust the standard Unicore installation to the specific D-Grid recommendations apply the following changes:

- Set the port used by NJS to 1128/tcp.

NJS

Resource providers may deploy one or more NJSs. Each NJS should listen for connection requests from the Gateways on port 1128/tcp. This port value is configured in the file *njs/conf/njs.properties* by the parameter *njs.gateway_port* [Ber06].

The site firewall should allow connections between the Gateways and the NJSs. Since the IP addresses of these hosts are known in advance, the firewall can be easily configured in order to allow network communication only among these entities.

With the above recommendations in place the firewall configuration for Unicore becomes as described in Appendix E.

5 SRM / dCache

Resource providers deploying SRM/dCache are recommended to use the following ports in their servers' configuration.

Host	Service	Port	Modification to default settings
dCache SE	GRIS	2135	No
	Extended GRIS	2136	No
	GridFTP	2811	No
		Range 20000-25000	Yes
SRMv1	8443	No	

Table 4: SRM/dCache port configuration

With the above recommendations in place the firewall configuration for SRM/dCache becomes as described in Appendix F.

5.1 Server

Resource providers deploying SRM/dCache are advised to configure the data management services as depicted in Figure 4, i.e.:

1. Install dCache SE on one single host.

Solid lines represent connections between the involved parties. Communications with external entities is displayed on the right side of the figure, where ports to be enabled for incoming connections through the firewall are also indicated.

Resource Provider

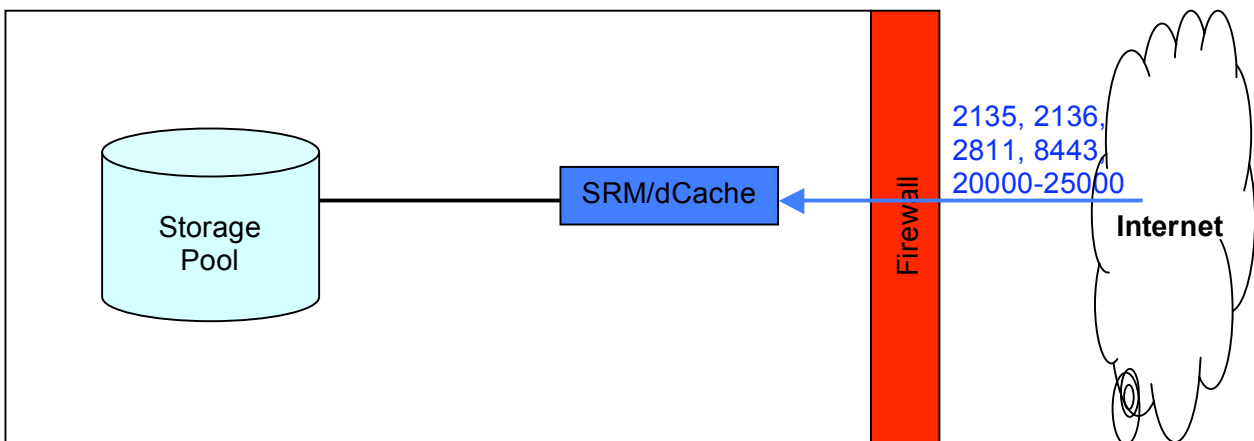


Figure 4: SRM/dCache recommended deployment at resource providers

To adjust the standard SRM/dCache installation to the specific D-Grid recommendations apply the following change:

- Set the *controllable ephemeral port range* to the interval 20000 - 25000.

Controllable Ephemeral Port Range

Preliminary operational experience with the Globus Toolkit shows that approximately 20 ports per expected simultaneous user on a given host should be provided. Given an average amount of 250 simultaneous active D-Grid users, the controllable ephemeral port range is in the order of 5000 ports.

The environment variable `GLOBUS_TCP_PORT_RANGE` restricts the ephemeral ports to a specific range. The value of this variable should be formatted as *min,max* (comma separated values). Setting this environment variable in all SRM/dCache affected components is achieved through the following procedure:

1. Configuration of GT libraries.
2. Configuration of GridFTP.

Appendix B describes with details how to set the port range to the interval 20000 – 25000.

6 OGSA-DAI

Resource providers deploying OGSA-DAI software are recommended to use the following ports in their servers' configuration.

Host	Service	Port (TCP)	Modification to default settings
OGSA-DAI	WSRF - DAI	8443	No

Table 5: OGSA-DAI port configuration

With the above recommendations in place the firewall configuration for SRM/dCache becomes as described in Appendix F.

6.1 Server

Resource providers deploying OGSA-DAI software are advised to configure the data management services as depicted in Figure 5, i.e.:

1. Install OGSA-DAI components on one single host.

Solid lines represent connections between the involved parties. Communications with external entities is displayed on the right side of the figure, where ports to be enabled for incoming connections through the firewall are also indicated.

Resource Provider

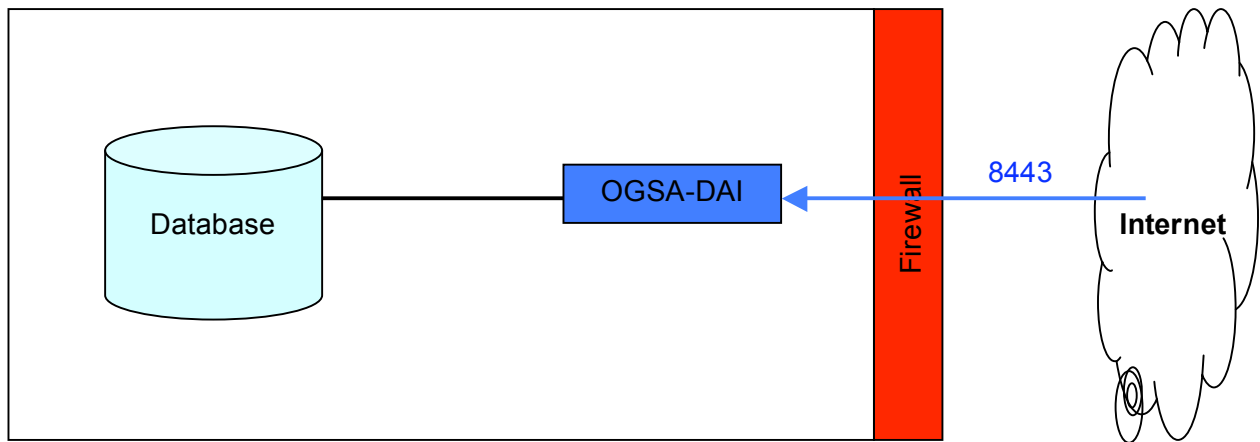


Figure 5: OGSA-DAI recommended deployment at resource providers

No additional configuration is required to adjust the standard OGSA-DAI installation to the specific D-Grid recommendations.

7 REGISTRATION OF GRID SERVICES

To improve the overall security level of Grid resources managed by a provider it is highly beneficial to rigorously know the IP addresses (or IP sub-networks) of all hosts deploying GT4, LCG/gLite or Unicore services in D-Grid. All middleware installations should be registered in a central directory. Such compulsory registration would allow the enforcement of rigorous firewall rules, minimizing the risk of external attacks and intrusions.

The firewall located at the border of the resource provider network should allow incoming connections only when all following criteria are satisfied:

- The source IP address is registered in the central directory.
- The destination IP address matches an internal host running a middleware service.
- The destination port either is a well-known port or it falls within the recommended TCP controllable range 20000 – 25000.

Directory service

The availability of a directory service recording the IP addresses (IP sub-networks) of all hosts deploying GT4, LCG/gLite and Unicore is a key component for the enforcement of optimal, strict firewall rules. The administration and management of this service should be centralized in one single institute. The directory should be constantly accessible online and should offer a notification service to inform interested people about insertions/deletions/modifications of the list.

Entries in the directory should contain at least the following information:

- Fully Qualified Domain Name
- IP address
- DN of host certificate (when available)
- Institute
- Set of Grid services running on the host
- Email address of system administrator
- Telephone number of system administrator
- Email address of security administrator
- Telephone number of security administrator

References

- [Wel06] V.Welch. *Globus Toolkit Firewall Requirements*. Version 9, October 2006
<http://www.globus.org/toolkit/security/firewalls/Globus-Firewall-Requirements-9.pdf>
- [Ber06] S.van den Berghe, *Using the NJS and TSI (V4)*, March 2006
http://www.unicore.eu/documentation/manuals/unicore5/files/NJS_TSI_Manual.pdf

Appendix A: Controlling the ephemeral port range in GT4 [Wei06]

Configuration of C libraries, client applications and services

Edit the file `/etc/profile` and insert a new line:

```
export GLOBUS_TCP_PORT_RANGE=20000,25000
```

Configuration of pre-WS-GRAM

The recommended way to configure the controllable ephemeral port range for pre-WS-GRAM is through modification of the xinet daemon.

Edit the file `/etc/xinetd.d/globus-gatekeeper` and add one “env” line:

```
env += GLOBUS_TCP_PORT_RANGE=20000,25000
```

Configuration of Java libraries, client applications and services

The Java System Property `org.globus.tcp.port.range` controls the ephemeral port range. There are three alternative ways to configure it:

1. Edit the file `~/globus/cog.properties` and insert a new line:

```
tcp.port.range=20000,25000
```

This is the recommended method, since all Java libraries automatically read this value.

2. Pass the value on the command line:

```
% java -Dorg.globus.tcp.port.range=20000,25000
```

3. Specify the value directly in the application:

```
System.setProperty("tcp.port.range","20000,25000")
```

Configuration of GridFTP

The recommended way to configure the controllable ephemeral port range for GridFTP is through modification of the xinet daemon.

Edit the file `/etc/xinetd.d/gridftp` and add one “env” line:

```
env += GLOBUS_TCP_PORT_RANGE=20000,25000
```

Appendix B: Controlling the ephemeral port range in LCG/gLite and SRM/dCache [Wei06]

Configuration of the GT libraries

Edit the file `/etc/profile` and insert a new line:

```
export GLOBUS_TCP_PORT_RANGE=20000,25000
```

Configuration of GRAM Gatekeeper and JobManager

There are four alternative ways to configure GRAM Gatekeeper and Job-Manager:

1. Modify the gLite configuration file.

(replace `$GLITE_LOCATION` with its actual path)

Edit the file `$GLITE_LOCATION/etc/globus-job-manager.conf` and add one line:

```
-globus-tcp-port-range "20000 25000"
```

2. Use a wrapper script.

(replace `$GLOBUS_LOCATION` with its actual path):

```
% mv $GLOBUS_LOCATION/libexec/globus-job-manager \
  $GLOBUS_LOCATION/libexec/globus-job-manager.real
% cat > $GLOBUS_LOCATION/libexec/globus-job-manager
#!/bin/sh
GLOBUS_TCP_PORT_RANGE=20000,25000
export GLOBUS_TCP_PORT_RANGE
exec $GLOBUS_LOCATION/libexec/globus-job-manager.real „$@"
^D
% chmod 755 $GLOBUS_LOCATION/libexec/globus-job-manager
```

3. Modify the xinet daemon.

Edit the file `/etc/xinetd.d/globus-gatekeeper` and add one “env” line:

(replace `$GLOBUS_LOCATION` with its actual path):

```
service globus-gatekeeper
{
    socket_type = stream
    protocol = tcp
    wait = no
    user = root
    server = $GLOBUS_LOCATION/sbin/globus-gatekeeper
    server_args = -conf $GLOBUS_LOCATION/etc/globus-gatekeeper.conf
    disable = no
    env += GLOBUS_TCP_PORT_RANGE=20000,25000
}
```

4. Modify the inet daemon.

Change the line in the file *inetd.conf* that starts the Gatekeeper.

(replace \$GLOBUS_LOCATION with its actual path):

```
gatekeeper stream tcp nowait root \
    /bin/env env GLOBUS_TCP_PORT_RANGE=20000,25000 \
    $GLOBUS_LOCATION/sbin/globus-gatekeeper -conf \
    $GLOBUS_LOCATION/etc/globus-gatekeeper.conf
```

Configuration of GridFTP

There are two alternative ways to configure GridFTP:

1. Modify the xinet daemon.

Edit the file */etc/xinetd.d/gsiftp* and add a "env" line:

(replace \$GLOBUS_LOCATION with its actual path):

```
service gsiftp
{
    socket_type = stream
    protocol = tcp
    wait = no
    user = root
    server = $GLOBUS_LOCATION/sbin/in.ftpd
    server_args = -l -a
    disable = no
    env += GLOBUS_TCP_PORT_RANGE=20000,25000
}
```

2. Modify the inet daemon.

Change the line in the file *inetd.conf* that starts the GridFTP server.

(replace \$GLOBUS_LOCATION with its actual path):

```
gsiftp stream tcp nowait root \
    /bin/env env GLOBUS_TCP_PORT_RANGE=20000,25000 \
    $GLOBUS_LOCATION/sbin/in.ftpd -l -a
```

Appendix C: Firewall filtering rules configuration for Globus Toolkit 4

This table describes the incoming network connections for service nodes running **Globus Toolkit 4**, when all resource providers implement the recommended configuration.

C = Controllable Ephemeral Port Range: 20000 - 25000.

This port range is locally configurable using the GLOBUS_TCP_PORT_RANGE environment variable and the Java System Property `org.globus.tcp.port.range`.

Node	Service	Source		Destination	
		Host	Port (TCP)	Host	Port (TCP)
GRAM + MDS + RFT + GSI-SSH	GRAM Gatekeeper	any UI	*	Localhost	2119
	GRAM JobManager	any UI	*	Localhost	C
	WS-GRAM (job startup)	any UI	*	Localhost	8443
	WS-GRAM (job control)	any UI	*	Localhost	C
	WS-MDS	External clients	*	Localhost	8443
	GridFTP control	External clients	C	Localhost	2811
	GridFTP data (single channel)	External clients	C	Localhost	C
	GridFTP data (multiple channel)	External clients	C	Localhost	C
	RFT	any UI, WN	*	Localhost	8443
GSI-SSH	any UI	*	Localhost	2222	

Appendix D: Firewall filtering rules configuration for LCG/gLite

This table describes the incoming network connections for service nodes running **LCG/gLite**, when all resource providers implement the recommended configuration.

C = Controllable Ephemeral Port Range: 20000 - 25000.

This port range is locally configurable using the GLOBUS_TCP_PORT_RANGE environment variable.

Node	Service	Source		Destination	
		Host	Port (TCP)	Host	Port (TCP)
LCG CE + MON + Site-BDII	GRAM Gatekeeper	any RB	C	Localhost	2119
	GRAM JobManager	any RB	C	Localhost	C
	GRIS (LDAP)	any BDII, RB, UI, CE, SE, WN	*	Localhost	2135
	Extended GRIS (LDAP)	any GridICE server	*	Localhost	2136
	BDII (LDAP)	any BDII, RB, UI, WN	*	Localhost	2170
	GridFTP control	any UI, CE, SE, WN	C	Localhost	2811
	GridFTP data (single channel)	any UI, CE, SE, WN	C	Localhost	C
	GridFTP data (multiple channel)	any UI, CE, SE, WN	C	Localhost	C
	RGMA – MON	Information Catalogue, any MON	*	Localhost	8080, 8443
RGMA – MON	any host in the local site	*	Localhost	8080, 8443	
RGMA – MON	any MON	*	Localhost	8088	
Worker Node	GridFTP data (multiple channel)	any SE	C	Localhost	C

Appendix E: Firewall filtering rules configuration for Unicore

This table describes the incoming network connections for service nodes running **Unicore**, when all resource providers implement the recommended configuration.

Node	Service	Source		Destination	
		Host	Port (TCP)	Host	Port (TCP)
NJS + TSI + UUDB	NJS	any Gateway	*	Localhost	1128

Appendix F: Firewall filtering rules configuration for SRM/dCache and OGSA-DAI

This table describes the incoming and outgoing network connections for service nodes running **SRM/dCache** and **OGSA-DAI**, when all resource providers implement the recommended configuration.

C = Controllable Ephemeral Port Range: 20000 - 25000.

This port range is locally configurable using the GLOBUS_TCP_PORT_RANGE environment variable.

Node	Service	Source		Destination	
		Host	Port (TCP)	Host	Port (TCP)
dCache	GRIS (LDAP)	any BDII, RB, UI, CE, SE, WN	*	Localhost	2135
	Extended GRIS (LDAP)	any GridICE server	*	Localhost	2136
	GridFTP control	any UI, CE, SE, WN	C	Localhost	2811
	GridFTP data (single channel)	any UI, CE, SE, WN	C	Localhost	C
	GridFTP data (multiple channel)	any UI, CE, SE, WN	C	Localhost	C
	SRMv1	any UI, WN, SE	*	Localhost	8443
OGSA-DAI	WSRF - DAI	any UI, WN	*	Localhost	8443