

EGEE

Apel User Guide

| | |
|----------------------|------------------------------------------------------------------|
| Document identifier: | EGEE-JRA1-TEC |
| Date: | September 14, 2007 |
| Activity: | JRA1: Middleware Engineering and Integration (UK Cluster) |
| Document status: | RELEASED |
| Document link: | |

Abstract: This document describes how to setup your site to be part of the Apel accounting system. It describes how to install and configure the Log Parser and Publisher components.

Document Change Log

| Issue | Date | Comment | Author |
|-------|----------|-------------|----------------------|
| 1.0 | 26/04/05 | First Draft | Rob Byrom, Dave Kant |

Document Change Record

| Issue | Item | Reason for Change |
|-------|------|-------------------|
|-------|------|-------------------|

Copyright ©Members of the EGEE Collaboration. 2004. See <http://eu-egee.org/partners> for details on the copyright holders.

EGEE (“Enabling Grids for E-science in Europe”) is a project funded by the European Union. For more information on the project, its partners and contributors please see <http://www.eu-egee.org>.

You are permitted to copy and distribute verbatim copies of this document containing this copyright notice, but modifying this document is not allowed. You are permitted to copy this document in whole or in part into other documents if you attach the following reference to the copied elements: “Copyright ©2004. Members of the EGEE Collaboration. <http://www.eu-egee.org>”

The information contained in this document represents the views of EGEE as of the date they are published. EGEE does not guarantee that any information contained herein is error-free, or up to date.

EGEE MAKES NO WARRANTIES, EXPRESS, IMPLIED, OR STATUTORY, BY PUBLISHING THIS DOCUMENT.

CONTENTS

| | | |
|----------|-----------------------------------|-----------|
| 1 | DISCLAIMER | 5 |
| 2 | INTRODUCTION | 6 |
| 2.1 | WHERE R-GMA FITS IN | 6 |
| 2.2 | R-GMA SERVER GUIDELINES | 6 |
| 3 | APEL CORE | 8 |
| 3.1 | DEPLOYMENT | 8 |
| 3.2 | INSTALLATION | 8 |
| 4 | LOG PARSER | 9 |
| 4.1 | DEPLOYMENT | 9 |
| 4.1.1 | SEPARATE GK/CE | 9 |
| 4.1.2 | COMBINED GK/CE | 9 |
| 4.2 | INSTALLATION | 9 |
| 4.2.1 | LOG PARSER COMPONENTS | 10 |
| 4.3 | CONFIGURATION | 10 |
| 4.3.1 | PREAMBLE | 11 |
| 4.3.2 | DBPROCESSOR | 11 |
| 4.3.3 | DBDELETEPROCESSOR | 12 |
| 4.3.4 | CPUPROCESSOR | 12 |
| 4.3.5 | EVENTLOGPROCESSOR | 13 |
| 4.3.6 | GKLOGPROCESSOR | 13 |
| 4.3.7 | BLAHDLOGPROCESSOR | 14 |
| 5 | LOG PARSER EXAMPLES | 16 |
| 5.1 | SEPARATE GK/CE | 16 |
| 5.1.1 | GK CONFIGURATION | 16 |
| 5.1.2 | CE CONFIGURATION | 16 |
| 5.2 | COMBINED GK/CE | 17 |
| 5.2.1 | GK/CE CONFIGURATION | 17 |
| 6 | PUBLISHER | 19 |
| 6.1 | DEPLOYMENT | 19 |
| 6.2 | INSTALLATION | 19 |
| 6.2.1 | PUBLISHER COMPONENTS | 19 |
| 6.3 | CONFIGURATION | 19 |
| 6.3.1 | JOINPROCESSOR | 19 |

| | | |
|-----------|-----------------------------------------------|-----------|
| 7 | APEL ENCRYPTION SCHEME | 21 |
| 7.1 | RSA KEY GENERATION | 21 |
| 7.2 | DETERMINATION OF KEY SIZE | 21 |
| 7.3 | PSEUDO RANDOM NUMBER GENERATION | 21 |
| 7.4 | ADDING RANDOM INFORMATION TO THE DN | 22 |
| 7.5 | HANDLING LONG DNS | 23 |
| 7.6 | DECRYPTING DATA | 23 |
| 8 | PUBLISHING DGAS DATA TO GOC | 24 |
| 9 | PUBLISHER EXAMPLES | 25 |
| 9.1 | R-GMA SERVER CONFIGURATION | 25 |
| 10 | TROUBLE SHOOTING AND FAQ | 26 |
| 10.1 | REPORTING PROBLEMS | 26 |
| 10.2 | OUTOFMEMORYERROR | 26 |

1 DISCLAIMER

This guide provides an overview of the Apel accounting software, where to install it and how to go about installing it. It does not provide an installation guide for setting up an R-GMA system. Users wishing to find out more about how to do this should refer to documentation available on the [R-GMA web site](http://hepunix.rl.ac.uk/egee/jra1-uk/)¹.

¹<http://hepunix.rl.ac.uk/egee/jra1-uk/>

2 INTRODUCTION

The Apel software is composed of two components: The Log Parser and Publisher.

The Log Parser interprets log files to extract job information and publishes it using R-GMA. Specifically, it processes the LCG gatekeeper logs, the system message logs and PBS or LSF event logs. Extracted data is then stored within a MySQL database. The Apel Log Parser also makes LDAP queries of the Computing Element to obtain the CPU performance figures for the worker node clusters and sub-clusters.

The Publisher is used to generate accounting records derived from the parsed logging data. The accounting records are then published into R-GMA where they are then collected by a central accounting server which, aggregates records from all sites.

2.1 WHERE R-GMA FITS IN

The Apel software uses R-GMA as a transport mechanism for moving accounting records stored per site to a centralised database. This is achieved by using a primary producer to publish accounting records (performed by the Publisher) and a secondary producer which is setup on the centralised store (see Figure 1).

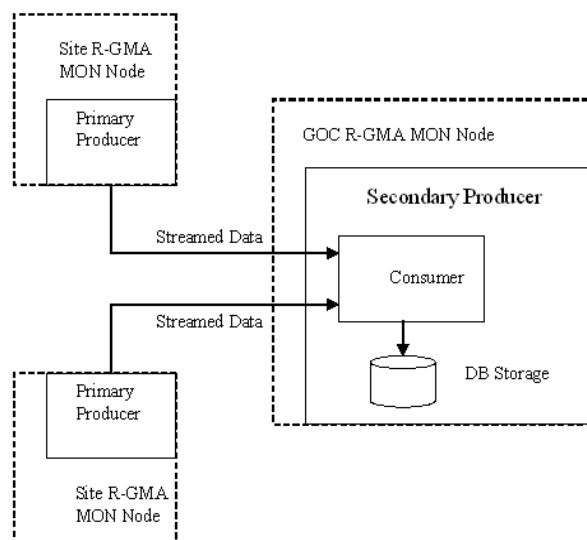


Figure 1: Accounting data is streamed from each site to a secondary producer

For the purpose of this user guide, we assume the secondary producer already exists (i.e. it's been configured by a sys admin located at the central accounting site).

For a more detailed overview of R-GMA, an excellent introduction is covered by the [R-GMA specification](#)²..

2.2 R-GMA SERVER GUIDELINES

The primary producer created by the Publisher is managed by an R-GMA server located locally on the same site. We recommend each site sets up its own R-GMA server, the reasons for this are summarised:

- accounting data may hold sensitive data which is private to each site.

²<https://edms.cern.ch/document/490223/2>

- the site administrator has direct control over the data and can control when locally stored accounting records can be deleted, regenerated and so on.
- the exact number of sites an R-GMA server can support is unclear.

For optimal performance, we recommend installing the R-GMA server on a dedicated host. Ideally this should be separate from the CE and GateKeeper hosts.

3 APEL CORE

The Apel Core contains common functionality used by the Log Parser and Publisher components, it does not provide any direct functionality to the user.

3.1 DEPLOYMENT

Since the Log Parser and Publisher depend upon an Apel Core, this module must be deployed on all hosts running either component.

3.2 INSTALLATION

The Apel Core comes in rpm format available from the [EGEE web site](#)³. The package contains a number of files, the following lists those the sys admin should pay attention to:

- `/opt/glite/share/glite-apel-core/scripts/apel-db-config.py` this script is used to setup a MySQL database for a particular user. It will also install the schema used by both the Log Parser and Publisher. Make sure you run this script where your MySQL server is located.
- `/opt/glite/share/glite-apel-core/scripts/apel-db-update.py` this script is used to update a MySQL database using an existing APEL user account. It is used when configuring the site to account for jobs onto the batch farm via the gLite CE. Make sure you run this script where your MySQL server is located.

On each host where the Apel Core is installed, make sure the following environment variables have been set:

```
export APEL_HOME=/opt/glite
#BouncyCastle Provider
#export BC_PROVIDER
```

If BC_PROVIDER is not set, APEL will assume it is located in

`$APEL_HOME/externals/share/java/bcprov-jdk14-122.jar`

When installing the Apel Core, the `apel-db-config.py` script must be run on the same host as the MySQL server (this will be the same host as the R-GMA server). Keep a note of the name of the database, user and password you use when running the script (you will need these when configuring the Log Parser and Publisher).

³<http://glite.web.cern.ch/glite/packages/>

4 LOG PARSER

The Log Parser is used to parse gate keeper, system message and event logs produced by a site running a batch processing system. Currently, the Log Parser comes in four different flavours: PBS, LSF, SGE and Condor. The Log Parser you choose will depend upon the batch processing system used by the site. However, the deployment and configuration are both equivalent.

4.1 DEPLOYMENT

The Log Parser can be deployed in a variety of ways depending upon the setup of your site. An overview of some of the typical deployments are described as follows.

4.1.1 SEPARATE GK/CE

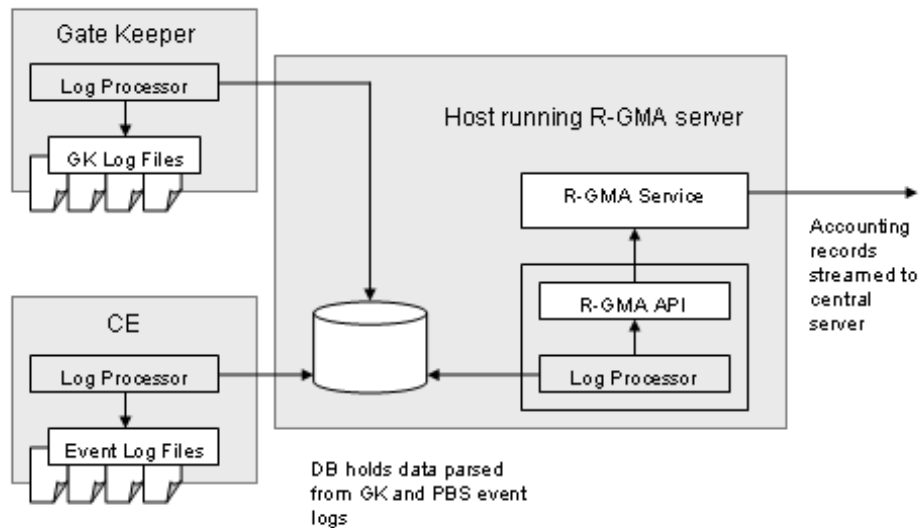


Figure 2: Shows Gate Keeper and CE on different hosts

See Fig 2. In this example, the gate keeper and CE are hosted on separate machines; both require an installation of the Log Parser. On the gate keeper, the Log Parser is configured to process gate keeper logs. Whereas the configuration on the CE is tailored to cater for event logs (conforming to the batch log format).

Fig 2 also shows the R-GMA server. This is where the Publisher is normally installed (described in 6).

4.1.2 COMBINED GK/CE

See Fig 3. In this setup the Log Parser is configured to process both the gate keeper logs and event logs but on the same host.

4.2 INSTALLATION

The Log Parser comes with its own rpm package available from the [EGEE web site](http://glite.web.cern.ch/glite/packages/)⁴. It requires the Apel Core module to be installed (available from the same repository).

⁴<http://glite.web.cern.ch/glite/packages/>

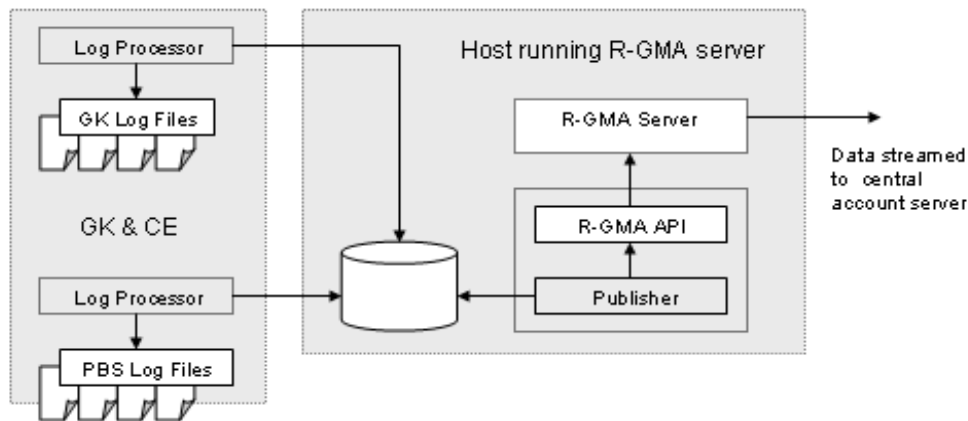


Figure 3: Shows Gate Keeper and CE combined on the same host

4.2.1 LOG PARSER COMPONENTS

The rpm installs the following components (using the PBS Log Parser as an example):

- `/opt/glite/bin/apel-pbs-log-parser` the script used to run the Log Parser.
- `/opt/glite/etc/glite-apel-pbs/parser-config.xml` contains an example config file used by the Log Parser.

The same files can be found when installing any of the LSF, SGE or condor Log Parser by substituting 'pbs' with 'lsf', CONDOR log parser substituting 'pbs' with 'condor' and SGE log parser substituting 'pbs' with 'sge'.

The Log Parser should be run daily. To schedule the Log Parser to run at 2:20 am every day, setup a cron entry as follows:

```
20 2 * * * /opt/glite/bin/apel-pbs-log-parser -f
/opt/glite/etc/glite-apel-pbs/config.xml > /var/log/apel.log 2>&1
```

In addition, it is a good idea to define a log rotation script, an example is show as follows:

```
/var/log/apel.log {
    copy
}
```

The log rotation script in this example is stored under `/etc/logrotate.d/apel`. Note, the cron and log rotate examples may need some tweaking depending upon the flavour of linux you use.

The sys admin should periodically backup data stored in the MySQL database used by the Log Parser. If tables grow large enough to cause disk space shortage, the sys admin is advised to back-up the database (using `mysqldump`) before purging the database contents.

4.3 CONFIGURATION

The Log Parser is setup using a configuration file encoded in an XML format. The configuration file is composed of a list of processors. Each processor carries out a unit of functionality.

When the Log Parser is started, the program will find any processors defined within the config file and will attempt to schedule each one consecutively. The list of configurable processors in order of execution are:

- DBDeleteProcessor
- CPUProcessor
- EventLogProcessor
- GKLogProcessor
- BlahdLogProcessor

An overview of the configuration format is detailed as follows.

4.3.1 PREAMBLE

The config file contains a preamble containing a number of fields that must be filled in by the user. The most important are the database credentials ie the database url, username and password. These credentials define the MySQL database to be used when storing parsed data derived from the gate keeper, system message and event log files.

An example of the preamble is shown as follows:

```
<?xml version="1.0" encoding="UTF-8"?>
<ApelConfiguration enableDebugLogging="yes">
  <DBURL>jdbc:mysql://grid-example.xy.ac.uk:3306/lcg</DBURL>
  <DBUsername>lcg</DBUsername>
  <DBPassword>no default</DBPassword>
  <SiteName>RAL-LCG2</SiteName>
</ApelConfiguration>
```

- enableDebugLogging: enables the user logging.
- DBURL: the URL that points to the MySQL database where the log data is to be recorded. Replace the hostname part with the hostname where your database is located.
- DBUsername/DBPassword: the username and password to be used to access the database on the MySQL database. Since a password is stored in this file you should setup suitable file access privileges.
- SiteName: is the name of the site producing the accounting data. **This must be set to the Sites' public Distinguished Name (DN) and must be consistent in all Apel config files used by the site.**

4.3.2 DBPROCESSOR

The DBProcessor is used to manage the APEL MySQL database schema. This processor examines the MySQL schema for compliance against the current version of APEL. Changes to the schema occur when new middleware releases allow for extra information to be captured in the accounting record.

Examples are any grid-level information that is mapped to local batch resources via the accounting log file (Savannah Patch 898). The UserFQAN derived from the users voms-proxy-init and taken from the LCMAPS mappings, the computing-element resource identifier and the global job identifier.

```
<DBProcessor inspectTables="yes"/>
```

- **inspectTable:** If enabled, the MySQL schema is interrogated. If APEL determines that an internal schema change is required, alterations will be performed. The time taken to complete this operation will depend on the size of the database. Once these changes have been done, subsequent schema interrogations will not modify the table structure.

4.3.3 DBDELETEPROCESSOR

The `DBDeleteProcessor` is used to delete records from the accounting db. An example of the delete processor is show as follows:

```
<DBDeleteProcessor cleanAll="yes"/>
```

Special attention should be taken when using such a configuration. **As a general rule, separate the delete processor within its own config file to avoid any unwanted deletion while parsing new log files.**

- **CleanAll:** If enabled, all table contents in the accounting database will be purged. If disabled, only records that were successfully merged to produce accounting records are removed (this will only delete gate keeper, system message and event log file data stored in the database).

4.3.4 CPUPROCESSOR

If this processor is present, CPU performance information will be generated and inserted into the database to be used when publishing accounting records. Normally this information comes from an LDAP query to a GIIS server, specified by the `GIIS` element.

The default port used is 2170. If GIIS services are available on another port, the port number can be specified in the `GIIS` element using the following format: `host:port`.

An example using a GIIS host with default port 2170

```
<CPUProcessor>
  <!--
    Default: LDAP query to BDII port 2170 with "o=grid"
  //-->
  <GIIS>lcgce.example</GIIS>
</CPUProcessor>
```

An example using a GIIS host with a specified port:

```
<CPUProcessor>
  <!--
    LDAP query to GRIS port 2135 with "mds-vo-name=local,o=grid"
  //-->
  <GIIS>lcgce.example:2135</GIIS>
</CPUProcessor>
```

This processor should be run with the `EventLogProcessor` (see 4.3.5). **It is crucial that spec data is available when accounting records are generated by the Publisher.**

Note, if the GIIS is unavailable, the Log Parser will log the error and will terminate. If this occurs, the admin user should re-run the script using an estimated default spec value (or investigate the health of the GIIS).

4.3.5 EVENTLOGPROCESSOR

The EventLogProcessor is used to parse event logs. The type of event logs parsed depend upon the log parser installed. This processor is usually run on the CE where the event logs are generated. An example config is shown as follows:

```

<EventLogProcessor>
  <SubmitHost>lcgce.example</SubmitHost>
  <Logs searchSubDirs="yes" reprocess="no">
    <Dir>/var/spool/pbs/server_priv/accounting</Dir>
    <ExtraFile>/home/toucan/pbslog00021123.gz</ExtraFile>
    <ExtraFile>/home/toucan/pbslog00021100.log</ExtraFile>
  </Logs>
  <Timezone>UTC</Timezone>
</EventLogProcessor>

```

- Dir: contains the directory where the event logs are to be found.
- ExtraFile: an optional parameter that allows additional filenames to be added to the list of event log files to be processed.
- SearchSubDirs: if enabled, all sub directories will be searched for event logs.
- Reprocess: when an event log file is processed, the log file name is stored in the database to prevent the file from be reprocessed in future. Enabling this option will force all event files to be re-parsed. Using this option is a safe way for re-building data from a batch of log files.
- Timezone: the event logs may be processed on a different machine from where they originated. For example, an admin user may want to run the Log Parser using event logs generated from a different site. Normally, the EventLogProcessor on each site will convert between UTC and local time of the processing host. However, this is not so useful where log files are being processed in a different timezone to that where the were generated.

The behaviour can be modified using the following options:

- derived: looks at the start and end time stamp of each job in the PBS event log (in UTC) to derive the timezone when the record was written.
- default: uses the host's timezone.
- any other value: for example 'GMT+1:30' indicates that 1 hour and 30 minutes has to be added to UTC to obtain the local time where the event log file was generated.

4.3.6 GKLOGPROCESSOR

The GKLogProcessor is used to process to the gate keeper and system message logs generated on the Gate Keeper host. An example with an explanation of each field is shown as follows:

```

<GKLogProcessor>
  <!--
    <SubmitHost>lcgce.example</SubmitHost>
    No Longer used by this component
  -->

```

```
<Logs searchSubDirs="yes" reprocess="no">
  <GKLogs>
    <Dir>/var/log</Dir>
  </GKLogs>
  <MessageLogs>
    <Dir>/var/log</Dir>
  </MessageLogs>
</Logs>
</GKLogProcessor>
```

- **SubmitHost:** Not used by the GKLogProcessor.

This is no longer used in the GKLogProcessor component. However, it may be used by the EventLogProcessor as Condor history and SunGridEngine batch logs do not provide the hostname of the submit node.

- **GKLogs:** contains the directory where the gate keeper log files will be found. gate keeper log file names contain a date suffix appended to globus-gatekeeper and may be of gzip format. File names examples are shown as follows:

```
globus-gatekeeper.log
globus-gatekeeper.log.20040125040202.0
globus-gatekeeper.log.20040201040203.0.gz
```

- **MessageLogs:** contains the directory where the system message log files are found. As with other files, these are recognised by having a particular filename format, examples of which are shown as follows:

```
messages.1.gz
messages.2.gz
messages.200401201123904.gz
messages.3.gz etc
```

- **SearchSubDirs:** if enabled all sub directories will be searched for gate keeper logs.
- **Reprocess:** Same as the attribute found in EventLogProcessor.

4.3.7 BLAHDLOGPROCESSOR

The BlandLogProcessor is used to parse the Blah accounting log file on the gLite CE. This processor is usually run on the gLite CE. An example config is shown as follows:

```
<BlandLogProcessor>
  <!--
    The LRMS submit hostname.
  -->
  <SubmitHost>lbgce01.something.other</SubmitHost>

  <!--
    Please set <BlandLogPrefix> to whatever is defined by the variable BLAHD_ACCOUNTING_I
    in /opt/glite/etc/blah.config
```

Example: If BLAHD_ACCOUNTING_INFO_LOG is set to /var/log/glite/accounting/blahp.log

The output logs are /var/log/glite/accounting/blahp.log-YYYYMMDD and
<BlahdLogPrefix> is set to blahp.log

-->

<BlahdLogPrefix>blahp.log</BlahdLogPrefix>

<Logs searchSubDirs="yes" reprocess="no">

<Dir>/var/log/glite/accounting</Dir>

</Logs>

</BlahdLogProcessor>

- **SubmitHost:** The batch server host name.
- **BlahdLogPrefix:** As explained above. The default log directory for accounting logs on the gLiteCE is /var/log/glite/accounting. All accounting files, which are rotated daily, have the name blahp.log followed by a the string "YYYYMMDD".
- **ExtraFile:** an optional parameter that allows additional filenames to be added to the list of event log files to be processed.
- **SearchSubDirs:** if enabled, all sub directories will be searched for event logs.
- **Reprocess:** when an event log file is processed, the log file name is stored in the database to prevent the file from be reprocessed in future. Enabling this option will force all event files to be re-parsed. Using this option is a safe way for re-building data from a batch of log files.

5 LOG PARSER EXAMPLES

5.1 SEPARATE GK/CE

This example is based on the setup described in 4.1.1. The CE and Gate Keeper are situated on different machines each running the Log Parser. The configuration used for each Log Processor is described as follows:

- GK: will run the GKLogProcessor.
- CE: will run the EventLogProcessor and CPUProcessor

5.1.1 GK CONFIGURATION

```

<?xml version="1.0" encoding="UTF-8"?>
<ApelConfiguration enableDebugLogging="yes">
  <DBURL>jdbc:mysql://grid-example.xy.ac.uk:3306/lcg</DBURL>
  <DBUsername>test</DBUsername>
  <DBPassword>info</DBPassword>
  <SiteName>gridka.de</SiteName>
  <GKLogProcessor>
    <Logs searchSubDirs="yes" reprocess="no">
      <GKLogs>
        <Dir>/var/data/gk</Dir>
      </GKLogs>
      <MessageLogs>
        <Dir>/var/message-data</Dir>
      </MessageLogs>
    </Logs>
  </GKLogProcessor>
</ApelConfiguration>

```

The Gate Keeper configuration only uses the GKLogProcessor. This processor will recursively search all directories below /var/data/gk and /var/message-data in-order to find any relevant log files. Since the 'reprocess' attribute is disabled, only new files will be processed.

Note, the DB credentials given are those used to access the DB located on the R-GMA server (which contains an installation of the MySQL server).

5.1.2 CE CONFIGURATION

```

<?xml version="1.0" encoding="UTF-8"?>
<ApelConfiguration enableDebugLogging="yes">
  <DBURL>jdbc:mysql://grid-example.xy.ac.uk:3306/lcg</DBURL>
  <DBUsername>test</DBUsername>
  <DBPassword>info</DBPassword>
  <SiteName>gridka.de</SiteName>
  <CPUProcessor>
    <GIIS>pbs-server2.gridka.de</GIIS>
  </CPUProcessor>
  <EventLogProcessor>

```



```
<!--
    Add this line for processing Condor and Sun Grid Engine Batch Logs Only
    <SubmitHost>hostname</SubmitHost>
  //-->

  <Logs searchSubDirs="yes" reprocess="no">
    <Dir>/var/data/el</Dir>
  </Logs>

  <Timezone>UTC</Timezone>
</EventLogProcessor>
</ApelConfiguration>
```

When the Log Parser is started, the CPUProcessor will be invoked first. It will try to query the GIIS (pbs-server2.gridka.de) in-order to retrieve a benchmark value of the site's CPU resources. Then all event logs found under /var/data/el will be processed (the search will include all sub directories).

5.2 COMBINED GK/CE

In this example, a site contains a GK and CE that, share the same host (based on the setup described in 4.1.2).

5.2.1 GK/CE CONFIGURATION

```
<?xml version="1.0" encoding="UTF-8"?>
<ApelConfiguration enableDebugLogging="yes">

  <DBURL>jdbc:mysql://grid-example.xy.ac.uk:3306/lcg</DBURL>
  <DBUsername>test</DBUsername>
  <DBPassword>info</DBPassword>
  <SiteName>RAL-LCG2</SiteName>

  <CPUProcessor>
    <GIIS>lcgce.example</GIIS>
  </CPUProcessor>

  <EventLogProcessor>
    <!--
      Add this line for processing Condor and Sun Grid Engine Batch Logs
      <SubmitHost>hostname</SubmitHost>
    //-->

    <Logs searchSubDirs="yes" reprocess="yes">
      <Dir>/var/spool/pbs/server_priv/accounting</Dir>
      <ExtraFile>/home/toucan/pbslog00021123.gz</ExtraFile>
      <ExtraFile>/home/toucan/pbslog00021100.log</ExtraFile>
    </Logs>
    <Timezone>UTC</Timezone>
  </EventLogProcessor>
```

```

    <GKLogProcessor>
      <Logs searchSubDirs="yes" reprocess="yes">
        <GKLogs>
          <Dir>/var/log</Dir>
        </GKLogs>
        <MessageLogs>
          <Dir>/var/log</Dir>
        </MessageLogs>
      </Logs>
    </GKLogProcessor>

  </ApelConfiguration>

```

In this setup, the `CPUProcessor` is run first and is configured to query the GIIS server named `lgce.example`. Then `EventLogProcessor` is will look for all log files located under the `Dir` path, it will also attempt to process the given extra files. The `GKLogProcessor` is the last to be scheduled and will process all GK logs. In this example all log the files found will be re-parsed.

Note, in both the GK and CE processors, the `searchSubDirs` option is enabled. This is the recommended option to ensure all files located under the given `Dir` are processed.

6 PUBLISHER

The Publisher is used to piece together accounting records derived from data parsed from gate keeper, system message and event log files. The generated data is stored locally within a MySQL database and is also published into R-GMA. The published data will then be stored by R-GMA onto the central accounting server.

6.1 DEPLOYMENT

The Publisher is deployed on the same host that runs the MySQL server. By default, the R-GMA server comes equipped with a MySQL server installation so we recommended this as the best place to deploy the Publisher.

6.2 INSTALLATION

The Publisher comes with its own rpm package available from the [EGEE web site](#)⁵. It requires the Apel Core rpm to be installed (also available from the same repository).

The Publisher also requires an R-GMA client installation, please refer to installation notes available on the [R-GMA web site](#)⁶.

6.2.1 PUBLISHER COMPONENTS

The rpm installs the following components:

- `/opt/glite/bin/apel-publisher` Script used to run the Publisher.
- `/opt/glite/etc/glite-apel-publisher/publisher-config.xml` Contains an example config file used by the Publisher.

Similar to the Log Parser, the publisher should be scheduled to run on a daily basis. See notes described in [4.2.1](#).

6.3 CONFIGURATION

The Publisher is setup using a configuration file encoded in an XML format. Compared to the Log Parser, the configuration file is composed of only two processors: the `DBDeleteProcessor` and `JoinProcessor`

The configuration format uses the same header as that used by the Log Parser (see [4.3.1](#)). The `DBDeleteProcessor` is the same as that defined in [4.3.3](#).

6.3.1 JOINPROCESSOR

The `JoinProcessor` is run on the R-GMA server. It attempts to join data stored by the Log Parser to build accounting records. An example configuration is shown as follows:

⁵<http://glite.web.cern.ch/glite/packages/>

⁶<http://hepunix.rl.ac.uk/egEE/jra1-uk/>

```
<!--
  publishGlobalUserName="no" =>
      Do not send User DN information to GOC.
      UserDN information does not leave the site!

  publishGlobalUserName="yes" =>
      Send User DN information to GOC.
      The UserDN is encrypted using a 1024-bit-RSA key-pair.

  disableJoin is optional and may be omitted.
//-->
<JoinProcessor publishGlobalUserName="no" disableJoin="no">
  <Republish>all</Republish>
</JoinProcessor>
```

- `publishGlobalUserName`: if disabled, the Distinguished Name (DN) of the user in each accounting record is set to 'NULL' when published into R-GMA. Some sites may wish to suppress the DN for reasons of *personal privacy*. By default, the DN will always be suppressed from publication.
- `publishGlobalUserName`: if enabled, the publisher will encrypt the UserDN information before publishing to GOC. The data will be decrypted by the GOC and access to this information restricted to the the individual User and VO administrator.
- `disableJoin`: if enabled, APEL will not attempt to join data. This option should be used for sites using the DGAS accounting system in which accounting records are provided by means of the DGAS2APEL conversion program to the APEL LcgRecords format. If omitted or disabled, APEL will try to stitch accounting records together.
- `Republish`: Contains a list of options the user can choose when republishing accounting data. These are listed as follows:
 - `all`: setting the republish option to 'all' causes all accounting data in the database to be regenerated and published. **This option should never be set as a long term configuration. Doing so will impair the performance of the R-GMA server especially if the volume of accounting records is huge (order of tens of thousands).**
 - `missing`: in some cases, the cenral accounting server may not receive the accounting re-courds that were published in a previous join. Rather than having to republish the complete set, the `JoinProcessor` can send a subset of data begining from the last successfull transfer. So when the `JoinProcessor` runs, a timestamp is recorded and the data is published as normal. When the `JoinProcessor` is next invoked, a count of all records that are greater than or equal to the timestamp is derived from the central accounting server. This is compared with the local database and if the accounting server contains fewer records, the `JoinProcessor` will republish all records begining from the timestamp. When the accounting server and local databases are in synch, the timestamp will then be updated as normal. **Note, this is the default behaviour.**
 - `nothing`: disables the republishing mode.
 - `gap`: the gap publisher can be used to send accounting data to the central archiver in a defined period of time. The `recordStart` and `recordEnd` specifies the date when jobs started and finished at the site respectively. Unlike the `missing` option, the gap publisher does not alter timestamps, or attempt to stitch accounting data together. It just sends all accounting records that start and finish in the date range specified.

7 APEL ENCRYPTION SCHEME

This section describes the method used by APEL to encrypt user DN information based on asymmetric private-public key encryption and a randomising function to reduce the likelihood of repeatable patterns arising in the ciphertext.

7.1 RSA KEY GENERATION

Using openssl we generate a 1024 bit private key in PEM format.

```
openssl genrsa -des3 -out private.pem 1024
```

The private key is then used to generate the public key

```
openssl rsa -in private.key -pubout -out public.pem
```

The private key is converted from PEM to PKCS#8 format.

```
openssl pkcs8 -in private_key.pem -topk8 -v2 des3 -nocrypt -outform DER -out private_key.p8c
```

7.2 DETERMINATION OF KEY SIZE

The strength of encryption scheme, the size of input data string that can be encrypted, and the speed of the algorithm are important when encrypting data. 1024 bit keys are regarded as strong and a number of cash prizes on offer by RSA laboratories (key challenges) remain open. Unlike encryption with symmetric keys, which can be performed on plaintext of an arbitrary size, the maximum number of plaintext bytes that can be encrypted using a public key depends on the size of the key (and the type of padding). For example, when using a 1024-bit RSA key with PKCS #1 v.1.5 padding (one of the most common options), it is not possible to encrypt a string, which is longer than 117 bytes (that is 117 ASCII or 58 two-byte Unicode characters). Increasing the size of the RSA key to 2048 bits will allow you to encrypt up to 245 bytes of data, but longer RSA key are expensive: they take more time to generate and operate. Using long public keys (longer than 1024 bits) can seriously degrade application performance. Additionally, the cipher text string generated by a 2048 bit key - after base64 encoding has been performed - is greater than 300 characters, larger than the character limit allowed in a MySQL VARCHAR(255) field.

In EGEE, the longest DNs seen are approximately 117 bytes. After random data has been added to the DN (describe later) this increases to approximately 170 bytes. Thus, a 1024-bit key will not encrypt all DNs. To deal with these rare occurrences (less than 1 % of DNs seen in EGEE), an algorithm is applied to shorten the input string before encryption is applied. This is describe at the end of the section.

7.3 PSEUDO RANDOM NUMBER GENERATION

The SecureRandom class provided by the Bouncy Castle Crypto package is a Java implementation of cryptographic algorithms. This class provides a cryptographically strong pseudo-random number generator (PRNG). A cryptographically strong pseudo-random number minimally complies with the statistical random number generator tests specified in FIPS 140-2, Security Requirements for Cryptographic Modules, section 4.9.1. Additionally, SecureRandom must produce non-deterministic output and therefore it is required that the seed material be unpredictable and that output of SecureRandom be cryptographically strong sequences as described in RFC 1750: Randomness Recommendations for Security. When

you create the class, either you pass it a seed, or it uses a self-generated seed. The seed is cryptographically hashed with the SHA-1 algorithm. Both the result and the internal state of the algorithm are added to the SHA hash to generate the next number.

If you don't pass in a seed, the seed is generated by the "seeder," which is itself an instance of SecureRandom. The seeder itself is seeded using internal thread-timing information that should be fairly hard to predict. The requisite timing information can take several seconds to gather, so the first time you create an instance of SecureRandom, this can be take up to 20 seconds.

The /dev/urandom device provides a reliable source of random output, however the output will not be generated from an equal amount of random input if insufficient input is available. Reads from the /dev/urandom device always return the quantity of output requested without blocking. If insufficient random input is available, alternate input will be processed by the random number generator to provide cryptographically secure output, the strength of which will reflect the strength of the algorithms used by the random number generator. The Output generated without random input is theoretically less secure than output generated from random input, so /dev/random should be used for applications for which a high level of confidence in the security of the output is required.

```
if (new File("/dev/urandom").exists())
{
    System.out.println("Salting SecureRandom (SHA1PRNG) from /dev/urandom");
    byte[] salt = new byte[8192];
    new FileInputStream("/dev/urandom").read(salt);
    secureRandom.setSeed(salt);
    for (int t=0; t<salt.length; t++) salt[t] = 0;
    salt = null;
}
```

7.4 ADDING RANDOM INFORMATION TO THE DN

Since many jobs can belong the same user DN, successive encryptions may lead to repeatable patterns in the ciphertext. To reduce the likelihood of this, two random numbers n1 and n2 are appended to the user DN string before encryption.

Original user DN:

C=UK/O=eScience/OU=QueenMaryLondon/L=Physics/CN=dave kant

After adding random information

956649486681*/C=UK/O=eScience/OU=QueenMaryLondon/L=Physics/CN=dave kant*30956649486681

Following recommendations from RSA laboratories and Bouncy Castle provider, we use an asymmetric block cipher with padding options: "RSA/ECB/PKCS1Padding" together with the GOC public key. The encryption algorithm also makes use of the SecureRandom function. The output cipher text is encoded into BASE64 format and can be written to a file or, in the case of APEL, sent as an encrypted stream via RGMA to the database archiver. The encoded Cipher is 172 characters in length and can be easily accommodated as a database field with limit VARCHAR(255).

Output Cipher text (Base64 String Representation)

```
*APEL V.0.2*Z1R6Lg9CnTr2GZCeJt576Xpe/Dj3EXofwTTGppBzr5nmkt6Vmra7qVqZQXwOcrFi/
yYE8TYkSP1+RRg4kxdifCBBls1gZf4IB8VKf5Tfa+l0s2cRocd0jgrhOmVpt2h4P5lGYaJ+zrZYCl
MmoGjHs2Kx0oZH3Tu+cpfUj5XH7RY=
```

As the encryption scheme may change in future implementations, e.g. with the support of robot certificates from CAs, a plaintext header is added to the cipher string to identify the method used to decrypt the data.

7.5 HANDLING LONG DNS

Known patterns in the UserDN e.g. */OU=personal certificate/* can be reduced to simply strings (*/ou=pc/*). Other techniques involve extraction of the email address from the string. Before any attempt is made to truncate the UserDN, the random numbers n1 and n2 are removed from the string and the strings size is checked before encryption occurs.

7.6 DECRYPTING DATA

Once the data has been sent via RGMA to the central accounting repository, it remains in the RGMA network in its encrypted. To extract the UserDN, the data can be unencrypted using the header information together with the GOC private key. The random information can be easily removed. Once this is done, the unencrypted data **MUST** be removed from the RGMA network and stored in a private database to protect the privacy of the accounting record. Access to the private database will be strictly prohibited and management of accounting data shall follow the guidelines specified by the GDB Accounting Policy document (in preperation).

Message decrypted with file private key:

```
956649486681*/C=UK/O=eScience/OU=QueenMaryLondon/L=Physics/CN=dave kant*30956649486681
```

8 PUBLISHING DGAS DATA TO GOC

An example Publisher to send data to the GOC from a DGAS accounting system. Note that the LcgRecords table is populated by the DGAS2APEL utility. The DGAS enabled site MUST have the LcgRecords and RepublishInfo tables defined. UserFQAN's (e.g. /VO/Role=NULL/Capability=NULL) may be sent via the LcgRecords field "LCGUserVO".

```
--
-- MySQL Table Definition
--
CREATE TABLE RepublishInfo (
    MeasurementDate DATE NOT NULL DEFAULT "000-00-00",
    MeasurementTime TIME NOT NULL DEFAULT "00:00:00",
    VOList TEXT
);

<?xml version="1.0" encoding="UTF-8"?>
<ApelConfiguration enableDebugLogging="yes">
    <DBURL>jdbc:mysql://grid-example.xy.ac.uk:3306/accounting</DBURL>
    <DBUsername>dgas</DBUsername>
    <DBPassword>dgas</DBPassword>
    <SiteName>INFN-site</SiteName>

    <!-- Examine APEL Schema //-->
    <DBProcessor inspectTables="yes"/>

    <JoinProcessor publishGlobalUserName="yes" disableJoin="yes">
        <Republish>missing</Republish>
    </JoinProcessor>
</ApelConfiguration>
```

The DBProcessor will interrogate the APEL schema and modify it if the option is set. Modifications to the schema are required in order to avoid truncation issues as discussed in Savannah bugs 23519,23053,20352.

The JoinProcessor will only look at the LcgRecords and RepublishInfo tables to determine which records should be sent to GOC. The grid users X509 DN certificate is encrypted on-the-fly before its published to the GOC.

9 PUBLISHER EXAMPLES

The Publisher configuration remains unchanged irrespective of where the CE or GK is setup (see 4.1.1 and 4.1.2). This is because the Publisher is deployed on the R-GMA server.

9.1 R-GMA SERVER CONFIGURATION

An example Publisher config file used on the R-GMA server is shown as follows:

```
<?xml version="1.0" encoding="UTF-8"?>
<ApelConfiguration enableDebugLogging="yes">
  <DBURL>jdbc:mysql://grid-example.xy.ac.uk:3306/lcg</DBURL>
  <DBUsername>gridka-mon</DBUsername>
  <DBPassword>mumbojumbo</DBPassword>
  <SiteName>gridka.de</SiteName>
  <!--
    publishGlobalUserName="no" =>
      Do not send User DN information to GOC.
      UserDN information (encrypted or otherwise) does not leave the site.

    publishGlobalUserName="yes" =>
      Send User DN information to GOC according to :-
      Encrypted UserDN sent to GOC.
      Unencrypted UserDNs are set to NULL before sending to GOC
  //-->

  <JoinProcessor publishGlobalUserName="no">
    <Republish>missing</Republish>
  </JoinProcessor>
</ApelConfiguration>
```

In this example, the JoinProcessor is setup to connect to the same MySQL database as the one used by the GK and CE. This is where the log data is held so the credentials of the DB are defined in this config file as well.

An example gap Publisher to re-send all accounting data from last November:-

```
<?xml version="1.0" encoding="UTF-8"?>
<ApelConfiguration enableDebugLogging="yes">
  <DBURL>jdbc:mysql://grid-example.xy.ac.uk:3306/lcg</DBURL>
  <DBUsername>gridka-mon</DBUsername>
  <DBPassword>mumbojumbo</DBPassword>
  <SiteName>gridka.de</SiteName>
  <JoinProcessor publishGlobalUserName="no">
    <Republish recordStart="2005-11-01" recordEnd="2005-11-30">gap</Republish>
  </JoinProcessor>
</ApelConfiguration>
```

The JoinProcessor will not reprocess all log data when generating the accounting records - instead it will concentrate on new data only. However, it will carry out a check to make sure all tuples sent in the last join were successfully received by the central accounting server. If the tuples have not been received it will then attempt to republish the set of records from the last successful synchronization check.

10 TROUBLE SHOOTING AND FAQ

This section describes some common problems encountered with APEL. A version of this FAQ is online at <http://goc.grid-support.ac.uk/accounting>

10.1 REPORTING PROBLEMS

Please report problems using the GGUS ticket service to either the *APEL Support Unit* or the *Accounting Policies Support Unit*. Tickets sent to these support units will be forwarded to apel-support@listserv.rl.ac.uk. Try to include the following information :-

- The siteName
- Description of the problem encountered
- Relevant parts of the APEL log file

10.2 OUTOFMEMORYERROR

The java virtual machine does not have enough memory allocated.

```
Exception in thread "main" java.lang.OutOfMemoryError
```

This is most commonly seen in the publisher logs because Apel must perform a MySQL join of data in order to build accounting records. To fix this problem, increase the JVM memory:-

1. edit the publisher script: `vi /opt/glite/bin/apel-publisher`
2. change the TWO lines "`java -Xmx256m`" to "`java -Xmx1024m`"
3. save changes

REFERENCES