

Table of Contents

The CREAM configuration file for both EMI-2 and EMI-3 releases.....	1
The CREAM configuration file structure.....	1
The Command Executor configuration.....	1
The Delegation Executor.....	2
The BLAH Executor.....	3
The Activity Executor (EMI-ES).....	5
The JDBC datasource configuration.....	6
The security authorization layer configuration.....	7

The CREAM configuration file for both EMI-2 and EMI-3 releases

The CREAM configuration file structure

The default location of the CREAM configuration file is `/opt/glite/etc/glite-ce-cream/cream-config.xml`. It is a XML file composed of a set of elements basically of three different types:

- *commandexecutor* which defines the specific capabilities (e.g. delegation, job management, activity management) provided by CREAM
- *JDBC dataSource* used for setting up the access to CREAM DB (MySQL)
- *argus-pep* or *authzchain* which defines the security authorization layer based on ARGUS service or gJAF

The configuration file looks like the following schema:

```
<service id="CREAM service (core2)">

  <commandexecutor
    id = "the id"
    category = "the category"
    ...

    <parameter name = "name_1" value = "value_1" />
    <parameter name = "name_2" value = "value_2" />
    ...
    <parameter name = "name_x" value = "value_x" />
  </commandexecutor>

  <dataSource name = "datasource_name"
    type = "javax.sql.DataSource"
    factory = "org.apache.commons.dbcp.BasicDataSourceFactory"
    ...
  />

  <authzchain name = "chain-1">
    <plugin name = "the name" classname = "org.glite.ce.xxx">
      <parameter name = "name_1" value = "value_1" />
      <parameter name = "name_2" value = "value_2" />
      ...
      <parameter name = "name_x" value = "value_x" />
    </plugin>
    ...
  </authzchain>

</service>
```

The Command Executor configuration

The *commandexecutor* represents the implementation of specific functionality provided by CREAM (e.g. delegation, job management, activity management). At the current time CREAM provides three different *commandexecutors* (*BLAHExecutor*, *ActivityExecutor* and *DelegationExecutor*) but new ones can be implemented and configured. The configuration of a *commandexecutor* is based on a well defined structure composed of a set of mandatory attributes and parameters which are specific for each executor:

```
<commandexecutor
  id = "the id"
  category = "the category"
```

```

commandqueuesize = "500"
commandqueeshared = "false"
commandworkerpoolsize = "50"
filename = "/usr/share/java/EXECUTOR_NAME.jar">

<parameter name = "name_1" value = "value_1" />
<parameter name = "name_2" value = "value_2" />
...
<parameter name = "name_x" value = "value_x" />
</commandexecutor>

```

The following is the list of the mandatory attributes with their meaning:

- **id**: the unique id (*)
- **category**: the name of the category to which the commandexecutor belongs (e.g. DELEGATION_MANAGEMENT, JOB_MANAGEMENT, ACTIVITY_MANAGEMENT) (*)
- **commandqueuesize**: the size of the in memory prefetched command queue (default: 500) (*)
- **commandqueeshared**: must be set to false (default) if the command queue is used by a single CREAM; must be set to true if there are multiple CREAM services using a single command queue (i.e. multiple CREAM services managing a single farm) (*)
- **commandworkerpoolsize**: defines the internal parallelism degree in terms of the number of commands (e.g. the user requests) processed simultaneously by a pool of threads (i.e. the command workers). Those commands are fetched from the the CREAM's command queue. Please note that the parameter value can influence the overall CREAM performance (default: 50) (*)
- **filename**: the path of the jar implementing the commandexecutor (*)

* please DO NOT change the default value which comes from the YAIM configuration.

The Delegation Executor

The Delegation Executor implements delegation functionality (e.g. make a new delegation request, renew or destroy a delegation, etc). The list of attributes and parameters which define this executor, complete of their meaning and default values is:

attributes

- **id**: *"DelegationExecutor"*
- **category**: *"DELEGATION_MANAGEMENT"* (*)
- **commandqueuesize**: *"0"* (NB: the value MUST be *"0"* because the executor doesn't need the command queue) (*)
- **commandqueeshared**: *"false"* (*)
- **commandworkerpoolsize**: *"0"* (NB: the value MUST be *"0"* because the executor doesn't need the command queue) (*)
- **filename**: *"/usr/share/java/glite-ce-cream-delegationExecutor.jar"*

parameters

- **CREAM_COPY_PROXY_TO_SANDBOX_BIN_PATH**: the path of the script for copying the user delegation proxy into the specified job sandbox (default: *"/usr/bin/glite-cream-copyProxyToSandboxDir.sh"*)
- **CREAM_PURGE_PROXY_FROM_SANDBOX_BIN_PATH**: the path of the script for deleting the user delegation proxy from the specified job sandbox (default: *"/usr/bin/glite-ce-cream-purge-proxy"*)
- **CREAM_SANDBOX_DIR**: the directory where the sandbox files are staged on the CREAM CE node (default: *"/var/cream_sandbox"*)

- **DELEGATION_PURGE_RATE**: specifies how often the delegation purger has to run (default "10" minutes)

* please DO NOT change the default value which comes from the YAIM configuration.

example

```
<commandexecutor id="DelegationExecutor"
  category="DELEGATION_MANAGEMENT"
  commandqueuesize="0"
  commandqueueshared="false"
  commandworkerpoolsize="0"
  filename="/usr/share/java/glite-ce-cream-delegationExecutor.jar">

  <parameter name="CREAM_COPY_PROXY_TO_SANDBOX_BIN_PATH" value="/usr/bin/glite-cream-copyProxyT
  <parameter name="CREAM_PURGE_PROXY_FROM_SANDBOX_BIN_PATH" value="/usr/bin/glite-ce-cream-purg
  <parameter name="CREAM_SANDBOX_DIR" value="/var/cream_es_sandbox" />
  <parameter name="DELEGATION_PURGE_RATE" value="10" /> <!-- minutes -->
</commandexecutor>
```

The BLAH Executor

This executor implements the CREAM's job management capabilities (e.g. jobSubmit, jobCancel, jobInfo, etc). CREAM interacts with the BLAH component, an abstraction layer providing a unified interface to the underlying LRMS. At the time of writing CREAM supports through BLAH the following LRMS: LSF, PBS/Torque, SGE, Condor and SLURM. The list of attributes and parameters which define this executor, complete of their meaning and default values is:

attributes

- id: *"BLAH executor"*
- category: *"JOB_MANAGEMENT"* (*)
- commandqueuesize: *"500"*
- commandqueueshared: *"false"*
- commandworkerpoolsize: *"50"*
- filename: *"/usr/share/java/glite-ce-cream-blahExecutor.jar"*

parameters

- **SANDBOX_TRANSFER_METHOD**: sets the sandbox transfer method (GSIFTP/LRMS) (default: *"GSIFTP"*)
- **CREAM_SANDBOX_DIR**: the directory where the sandbox files are staged on the CREAM CE node (default: *"/var/cream_sandbox"*)
- **CREAM_CREATE_SANDBOX_BIN_PATH**: the path of the executable for creating the job sandbox (default: *"/usr/bin/glite-cream-createsandboxdir"*)
- **CREAM_COPY_PROXY_TO_SANDBOX_BIN_PATH**: the path of the script for copying the user delegation proxy into the specified job sandbox (default: *"/usr/bin/glite-cream-copyProxyToSandboxDir.sh"*)
- **CREAM_PURGE_SANDBOX_BIN_PATH**: the path of the executable for purging the job sandbox (default: *"/usr/bin/glite-ce-cream-purge-sandbox"*)
- **CREAM_PURGE_PROXY_FROM_SANDBOX_BIN_PATH**: the path of the executable for purging the user delegation proxy from the job sandbox (default: *"/usr/bin/glite-ce-cream-purge-proxy"*)
- **CREAM_JOB_SENSOR_HOST**: enable this parameter (and set it to the hostname of the CREAM CE) if the CEMon job sensor plugin is installed (as needed for submissions via WMS/ICE) (default: parameter disabled)

- **CREAM_JOB_SENSOR_PORT**: represents the port used by CREAM to get notifications from the CEMon job sensor plugin. Enable this parameter if the CEMon job sensor plugin is installed (as needed for submissions via WMS/ICE) (default: "49154")
- **MAX_LEASE_TIME**: the maximum allowed lease time (in second). I.e. if a client specifies a lease time too big, this value is used instead (default: "36000")
- **LEASE_EXECUTION_RATE**: specifies how often the job purger has to run (default: "30" minutes)
- **GLEEXEC_BIN_PATH**: the path of the glEXEC executable (default: "/usr/sbin/glEXEC")
- **GLEEXEC_CAT_CMD_PATH**: the pathname of the "cat" command invoked by CREAM via glEXEC (default: "/bin/cat")
- **BLAHP_BIN_PATH**: the path of the blahpd executable (default: "/usr/bin/blahpd")
- **BLAH_JOBID_PREFIX**: represents a prefix for job names for BLAH. If the same BLAH parser serves multiple CREAM CEs, please choose a different value for each of these CREAM services. The value MUST be a string 6 characters long that starts with 'cr' and ends with '_' (default: "cream_")
- **BLAH_COMMAND_TIMEOUT**: represents the maximum time interval accepted by CREAM for the execution of commands by BLAH (default: "300" seconds)
- **BLPARSER_RETRY_COUNT**: represents the number of attempts to contact the BLAH parser (if it is not reachable) before giving up. If -1 is specified, CREAM will never give up (default: "100")
- **BLPARSER_RETRY_DELAY**: the time interval (ms) between two attempts to contact the BLAH parser (default: "60000")
- **LRMS_EVENT_LISTENER_PORT**: the port used by CREAM to receive notifications about job status changes sent by the BLParser/JobWrapper (default: "49152")
- **HOST_SMP_SIZE**: the default value for the *hostsmpsize* attribute (default: "2")
- **JOB_WRAPPER_DELEGATION_TIME_SLOT**: "3600"
- **JOB_WRAPPER_COPY_PROXY_MIN_RETRY_WAIT**: the minimum time interval expressed in seconds, between the first attempt and the second one for retrieving the user delegation proxy (default: "60")
- **JOB_WRAPPER_COPY_RETRY_FIRST_WAIT_ISB**: if during a input sandbox file transfer occurs a failure, the JW retries the operation after a while. The sleep time between the first attempt and the second one is the initial wait time (i.e. the wait time between the first attempt and the second one) expressed in seconds. In every next attempt the sleep time is doubled (default: "60")
- **JOB_WRAPPER_COPY_RETRY_COUNT_ISB**: the maximum number of ISB file transfers that should be tried (default: "2")
- **JOB_WRAPPER_COPY_RETRY_FIRST_WAIT_OSB**: if during a output sandbox file transfer occurs a failure, the JW retries the operation after a while. The sleep time between the first attempt and the second one is the initial wait time (i.e. the wait time between the first attempt and the second one) expressed in seconds. In every next attempt the sleep time is doubled (default: "300")
- **JOB_WRAPPER_COPY_RETRY_COUNT_OSB**: the maximum number of OSB file transfers that should be tried (default: "6")
- **JOB_WRAPPER_TEMPLATE_PATH**: the path where the JobWrapper template is placed (default: "/etc/glite-ce-cream/")
- **JOBS_SCRATCH_DIR**: the tmp directory used by CREAM (default: "/var/tmp")
- **JOB_SUBMISSION_MANAGER_SCRIPT_PATH**: CREAM is able to protect itself if the load (i.e. memory usage, etc) is too high. In this scenario it refuses new job submissions, while the other commands are still allowed. The whole stuff is implemented via a limiter script (i.e. */usr/bin/glite_cream_load_monitor*) very similar to the one used in the WMS. Basically this limiter script checks the values for some system and CREAM specific parameters, and compare them against some thresholds defined in its configuration file (i.e. */etc/glite-ce-cream-utils/glite_cream_load_monitor.conf*). If one or more threshold is exceeded, new job submissions get not allowed and an error message is returned to the client. Moreover the limiter script is run every 10 (default: */usr/bin/glite_cream_load_monitor* */etc/glite-ce-cream-utils/glite_cream_load_monitor.conf*)
- **JOB_SUBMISSION_MANAGER_ENABLE**: in order to disable the limiter, it is needed to set this parameter value to *false* and restart the service (default: "true")
- **JOB_PURGE_POLICY**: the policy concerning which jobs should be deleted by the automatic job

purger and when. For example, if `JOB_PURGE_POLICY` is *"ABORTED 1 days; CANCELLED 2 days; DONE-OK 3 days; DONE-FAILED 4 days; REGISTERED 5 days;"* then the job purger will purge jobs which are:

- ◆ in ABORTED status for more than 1 day
- ◆ in CANCELLED status for more than 2 days
- ◆ in DONE-OK status for more than 3 days
- ◆ in DONE-FAILED status more than 4 days
- ◆ in REGISTERED status for more than 5 days
- **JOB_PURGE_RATE**: specifies how often the job purger has to run (default *"300"* minutes)

* please DO NOT change the default value which comes from the YAIM configuration.

The Activity Executor (EMI-ES)

This Activity Executor implements the functionality (e.g. `createActivity`, `getActivityInfo`, etc) defined in the EMI-ES specification (v1.16). The list of attributes and parameters which define this executor, complete of their meaning and default values is:

attributes

- `id`: *"ActivityExecutor"*
- `category`: *"ACTIVITY_MANAGEMENT"* (*)
- `commandqueuesize`: *"500"*
- `commandqueueshared`: *"false"*
- `commandworkerpoolsize`: *"50"*
- `filename`: *"/usr/share/java/glite-ce-cream-activityExecutor.jar"*

parameters

- **ACTIVITY_STATUS_MONITOR_AGE**: defines how old should be a not terminated activity in order to ask its current status to the LRMS (default: *"24"* hours)
- **ACTIVITY_STATUS_MONITOR_RATE**: specifies how often CREAM checks the status of the not terminated activities (default: *"6"* hours)
- **ACTIVITY_WRAPPER_TEMPLATE_PATH**: the path where the ActivityWrapper template is placed (default: *"/etc/glite-ce-cream-es/"*)
- **BDII_URI**: defines the BDII URI (e.g. *_"ldap://cream.pd.infn.it:2170" _*)
- **BDII_RATE**: specifies how often CREAM has to get information from the BDII (default: *"60"* minutes)
- **BLAH_BIN_PATH**: the path of the blahpd executable (default: *"/usr/bin/blahpd"*)
- **BLAH_PREFIX**: represents a prefix for job names for BLAH. If the same BLAH parser serves multiple CREAM CEs, please choose a different value for each of these CREAM services. The value MUST be a string 6 characters long that starts with 'cr' and ends with '_' (default: *"cr-es_"*)
- **BLAH_COMMAND_TIMEOUT**: represents the maximum time interval accepted by CREAM for the execution of commands by BLAH (default: *"300"* seconds)
- **BLAH_NOTIFICATION_LISTENER_PORT**: the port used by CREAM to receive notifications about job status changes sent by the BLParser/ActivityWrapper (default: *"49155"*)
- **BLAH_NOTIFIER_RETRY_COUNT**: represents the number of attempts to contact the BLAH parser (if it is not reachable) before giving up. If -1 is specified, CREAM will never give up (default: *"100"*)
- **BLAH_NOTIFIER_RETRY_DELAY**: the time interval (ms) between two attempts to contact the BLAH parser (default: *"60000"*)
- **COPY_PROXY_MIN_RETRY_WAIT**: the minimum time interval expressed in seconds, between the first attempt and the second one for retrieving the user delegation proxy (default: *"60"*)

- **COPY_RETRY_FIRST_WAIT_ISB**: if during a input sandbox file transfer occurs a failure, the JW retries the operation after a while. The sleep time between the first attempt and the second one is the initial wait time (i.e. the wait time between the first attempt and the second one) expressed in seconds. In every next attempt the sleep time is doubled (default: "60")
- **COPY_RETRY_COUNT_ISB**: the maximum number of ISB file transfers that should be tried (default: "2")
- **COPY_RETRY_FIRST_WAIT_OSB**: if during a output sandbox file transfer occurs a failure, the JW retries the operation after a while. The sleep time between the first attempt and the second one is the initial wait time (i.e. the wait time between the first attempt and the second one) expressed in seconds. In every next attempt the sleep time is doubled (default: "300")
- **COPY_RETRY_COUNT_OSB**: the maximum number of OSB file transfers that should be tried (default: "6")
- **CREATE_WRAPPER_BIN_PATH**: (default: "/usr/bin/glite-ce-cream-create-wrapper")
- **CREATE_SANDBOX_BIN_PATH**: the path of the executable for creating the job sandbox (default: "/usr/bin/glite-cream-createsandboxdir")
- **DEFAULT_LRMS_NAME**: defines the default LRMS used by CREAM if not specified by the user
- **DEFAULT_QUEUE_NAME**: defines the default LRMS's queue used by CREAM if not specified by the user
- **DELEGATION_TIME_SLOT***: (default "3600" seconds)
- **LIMIT_FOR_LIST_ACTIVITIES**: defines the maximum number of items contained into the response to be sent back to the client (default: "500")
- **PURGE_SANDBOX_BIN_PATH**: the path of the executable for purging the job sandbox (default: "/usr/bin/glite-ce-cream-purge-sandbox")
- **SANDBOX_DIR**: the directory where the sandbox files are staged on the CREAM CE node (default: "/var/cream-es_sandbox")

* please DO NOT change the default value which comes from the YAIM configuration.

The JDBC datasource configuration

CREAM relies on an external relational database to store its internal state (e.g. jobs, activities, delegations, commands queue, etc). This improves fault tolerance as it guarantees that this information is preserved across restarts of CREAM. Moreover, the use of a SQL database improves responsiveness of the service while performing queries which are needed by the usual CREAM operations, such as getting the list of jobs associated with a specific user. The CREAM deployment provides MySQL as preferred database, but any SQL database accessible through JDBC should be well supported although we don't have yet experience on it. Note that the database server can be installed on a dedicated host, or can share the same machine hosting CREAM. The JDBC connection is configured by the following template:

```
<dataSource name="datasource_name"
  type="javax.sql.DataSource"
  factory="org.apache.commons.dbcp.BasicDataSourceFactory"
  driverClassName="org.gjt.mm.mysql.Driver"
  username="the user" password="the password"
  maxActive="200"
  maxIdle="30"
  maxWait="10000"
  url="jdbc:mysql://localhost:3306/DATABASE_NAME?autoReconnect=true"
  validationQuery="SELECT 1"
  testOnBorrow="true"
  testWhileIdle="true"
  timeBetweenEvictionRunsMillis="20000"
  minEvictableIdleTimeMillis="90000"
  logAbandoned="false"
  removeAbandoned="true"
  removeAbandonedTimeout="30"
/>
```

We suggest not to change the default values coming from the YAIM configuration phase. The parameters are all mandatory and their values well defined therefore there is no need to change them. Moreover please check the section 1.5.5 of the *System Administrator Guide* (i.e.

<https://wiki.italiangrid.it/twiki/bin/view/CREAM/SystemAdministratorGuideForEMI2>) if you want to configure the CREAM databases on a host different than the CREAM service (by using YAIM).

The security authorization layer configuration

CREAM supports two different authorization systems based on the ARGUS service or the gJAF (grid Java Authorization Framework). Therefore the CREAM configuration depends on the authZ system selected. In case of ARGUS the XML section looks like:

```
<adminlist filename="/etc/grid-security/admin-list"/>
<argus-pep name="pep-client1"
  resource_id="CREAM_PEP_RESOURCEID"
  cert="TOMCAT_HOSTCERT_LOCATION"
  key="TOMCAT_HOSTKEY_LOCATION"
  passwd=""
  mapping_class="org.glite.ce.cream.authz.argus.ActionMapping">
  <endpoint url="ARGUS_PEPD_ENDPOINTS" />
</argus-pep>
```

while the configuration for the gJAF system is:

```
<authzchain name="chain-1">
  <plugin name="localuserpip" classname="org.glite.ce.commonj.authz.gjaf.LocalUserPIP">
    <parameter name="glexec_bin_path" value="/usr/sbin/glexec" />
    <parameter name="glexec_probe_cmd" value="/usr/bin/id" />
    <parameter name="methods" value="JobRegister, putProxy, getProxyReq, renewProxyReq, getTe
  </plugin>

  <plugin name="bannerpdp" classname="org.glite.ce.commonj.authz.gjaf.BlackListServicePDP">
    <parameter name="blackListFile" value="/etc/lcas/ban_users.db" />
  </plugin>

  <plugin name="admincheckpip" classname="org.glite.ce.commonj.authz.gjaf.AdminCheckerPIP">
    <parameter name="adminList" value="/etc/grid-security/admin-list" />
  </plugin>

  <plugin name="gridmappdp" classname="org.glite.ce.commonj.authz.gjaf.GridMapServicePDP">
    <parameter name="gridMapFile" value="/etc/grid-security/grid-mapfile" />
  </plugin>

  <plugin name="vomspdp" classname="org.glite.ce.commonj.authz.gjaf.VomsServicePDP">
    <parameter name="gridMapFile" value="/etc/grid-security/grid-mapfile" />
  </plugin>
</authzchain>
```

We suggest not to change the default values coming from the YAIM configuration phase. The parameters are all mandatory and their values well defined therefore there is no need to change them.

-- LisaZangrando - 2012-12-04

This topic: CREAM > CREAMConfigurationFile

Topic revision: r11 - 2013-03-26 - LisaZangrando



Copyright © 2008-2021 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.

Ideas, requests, problems regarding TWiki? Send feedback