

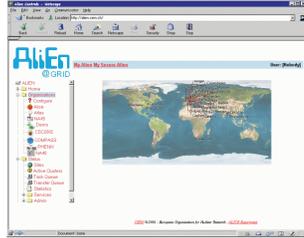
AliEn - EDG Interoperability in Alice

AliEn (Alice Environment) is a grid-like system for job distribution and data management developed by the Alice Collaboration.

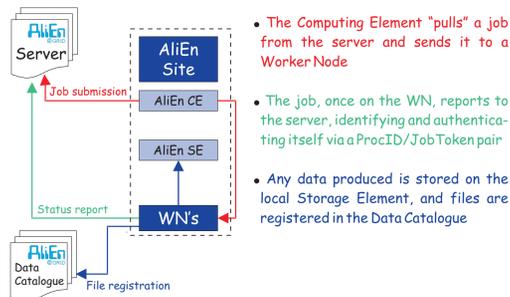
The AliEn system has already been used as a production system for several Alice MonteCarlo simulation productions. Based on existing, open-source modules, it is a simpler, lightweight alternative to full-blown Grid projects; its features include:

- Hierarchical, MySQL-based Data Catalogue with command-line and GUI interfaces
- "Pull-model" for job submission, with built-in support for massive productions
- LDAP-based remote configuration of sites
- Remote package management
- Support for metadata management
- Accepts Globus certificates
- Web-based monitoring and management:

<http://alien.cern.ch>

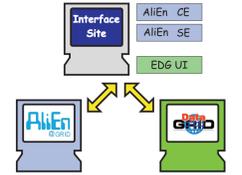


The AliEn Job Cycle



Interfacing to EDG:

The whole Grid (namely, a Resource Broker) is seen by the server as a single AliEn Computing Element, and the whole Grid storage is seen as a single, large AliEn Storage Element.



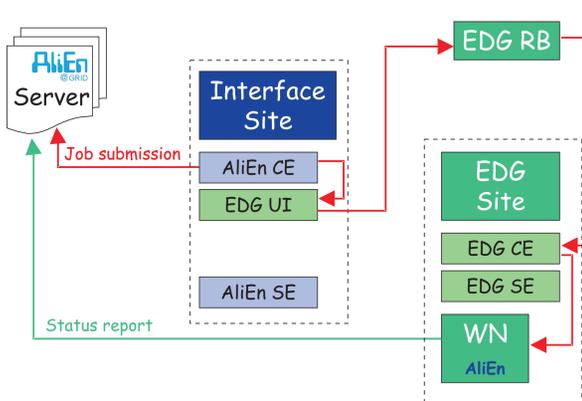
• "Interface sites" are the contact points between the two systems: an Interface Site is an EDG User Interface machine that runs also, as services, the AliEn client suite: Storage Element, Computing Element and Cluster Monitor.

• The Interface Machine takes care of JDL translation, of interfacing to Resource Brokers and Replica Catalogues, and of staging files between the two systems.

• The interface may become a bottleneck, but it is possible to deploy more than one.

• One such interface is in place in Torino, Italy, submitting jobs to the EDG testbed and to the GLUE-enabled DataTAG testbed: grid012.to.infn.it

Job Submission



• The Job Submission interface is implemented as a "Local Queue" object (like PBS or LSF): a Resource Broker is seen as a single queue.

• When the AliEn Computing Element "pulls" a job from the server, a new JDL file is generated from the original AliEn JDL (which is not directly compatible with EDG JDL), adding:

- requirements for a suitable version of AliEn runtime environment;
- requirement for outgoing network connectivity from the Worker Node (needed to contact the server)
- environment variables containing the ProcID and the JobToken (used by the WN to authenticate itself with the server) and some configuration information needed to run AliEn on the WN

• In AliEn JDL files, a SpecialRequirements ClassAd attribute can be used to pass requirements that are not meaningful to AliEn

• The job is then submitted to the Resource Broker

• When the job arrives on the WN, AliEn code is run to contact the server, piping the stdout and stderr streams to the logging system

• A database of EDG/AliEn JobID correspondences is kept by the interface CE to keep track of jobs that did not report yet to the server

```
Requirements = member(other.Packages, "AliRoot");
Packages = "AliRoot";
Arguments = "--round 2002-02 --run 00071 --event 269 --version v3.09.06";
Executable = "/Alice/bin/AliRoot.sh";
InputFile = ("LF:/alice/simulation/2002-02/v3.08.02/00071/Config.C",
"LF:/alice/simulation/2002-02/v3.08.02/00071/grun.C");
Type = "Job";
OutputFile = {"galice.root"};

# JDL automatically generated by AliEn
Executable = "/opt/alien/bin/alien";
Arguments = "RunJob 103742";
StdOutput = "std.out";
StdError = "std.err";
OutputSandbox = {"std.err", "std.out"};
VirtualOrganization = "datatag";
Environment = {"ALIEN_PROC_ID=103742", \
"ALIEN_JOB_TOKEN=mxrC9Roni!F!OpR55Tm9g93gS0yt28", \
"ALIEN_CM_AS_LDAP_PROXY=grid012.to.infn.it:8084", \
"ALIEN_SE_MSS=EDG", \
"ALIEN_SE_FULLNAME=Alice:Torino:EDG", \
"ALIEN_DISABLE_PACKAGES=1"};
Requirements = \
Member(other.GlueHostApplicationSoftwareRunTimeEnvironment, "ALIEN-1.29.10") \
&& Member(other.GlueHostApplicationSoftwareRunTimeEnvironment, "ALICE-3.09.06") \
&& Member(other.GlueHostNetworkAdapterOutboundIP=true);
```

File Registration

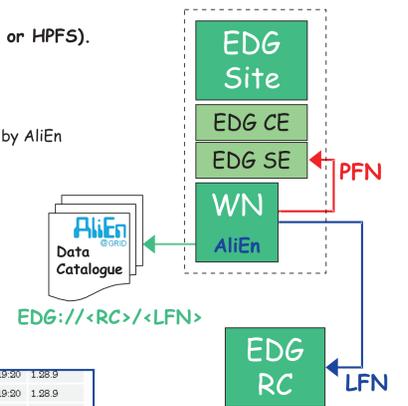
All files generated are doubly-registered in both (AliEn and EDG/Globus) Replica Catalogues, thus ensuring accessibility from both systems.

• The Storage Element/Replica Manager interface is implemented as a "Mass Storage System" object (like CASTOR or HPFS).

- An output file is generated by an AliEn job running on some EDG site
- The CloseSE (in EDG terms) is determined by AliEn code (running on the WN) from the .BrokerInfo file
- The file is stored in the selected SE and registered in EDG Replica Catalogue, while uniqueness of the Logical File Name is enforced by AliEn
- An "intermediate" file name is generated using the EDG LFN and the Virtual Organization name (to allow form multiple VO/RC)
- The file is then registered in AliEn Data Catalogue using this "intermediate FN" as Physical File Name.

```
grid012.to.infn.it> alien
Connecting to database alien.cern.ch:3307 ...
[aliendb.cern.ch:3307] /alice/cern.ch/alice/ > whereis /proc/103742/galice.root
And the file is in Alice:Torino:EDG edg://alice/00001.galice.root.1047921833

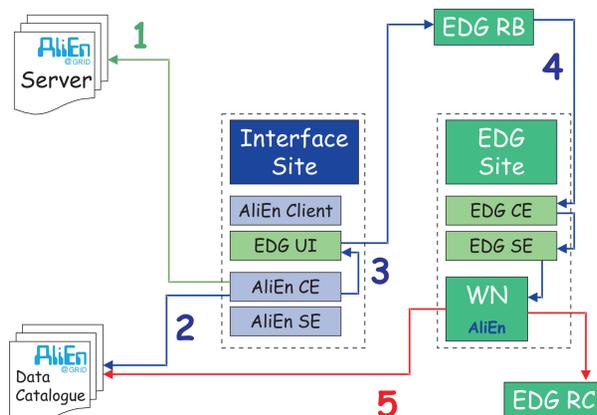
grid012.to.infn.it> edg-replica-manager-listReplicas -c /opt/edg/etc/alice/rc.conf -l 00001.galice.root.1047921833
configuration file: /opt/edg/etc/alice/rc.conf
logical file name: 00001.galice.root.1047921833
protocol: gsiftp
For LFN 00001.galice.root.1047921833 the following replicas have been found:
location 1: grid025.pd.infn.it//shared/alice/00001.galice.root.1047921833
The program was successfully executed.
```



Data Access

The basic case

1. Whenever a job (in its JDL file) requires data stored in an EDG Storage Element, the AliEn job broker will try to assign it to the EDG interface Computing Element
2. The interface queries the AliEn Data Catalogue to translate the AliEn Logical File Name into the corresponding EDG LFN...
3. ...and generates the appropriate InputData ClassAds attributes to be included in the JDL.
4. The generated JDL is then submitted to the Resource Broker, that will take care to send it where the data are available.
5. Once running, the job will query again the AliEn Data Catalogue, and the EDG one subsequently, to actually retrieve the data.



Some exotic cases

- Should the need arise, it is possible for jobs running on either system to access data physically stored in the other.
- For data residing on EDG being accessed from AliEn sites, the Data Catalogue query is a two-step process via the AliEn and EDG Replica Catalogues, while jobs running on EDG sites can access directly the AliEn Data Catalogue.
- The physical transfer of files is always performed by first staging them to the Interface machine(s), that again acts as a contact point between the two systems.

