



Enabling Grids for E-science

CREAM-CE

Massimo Sgaravatto

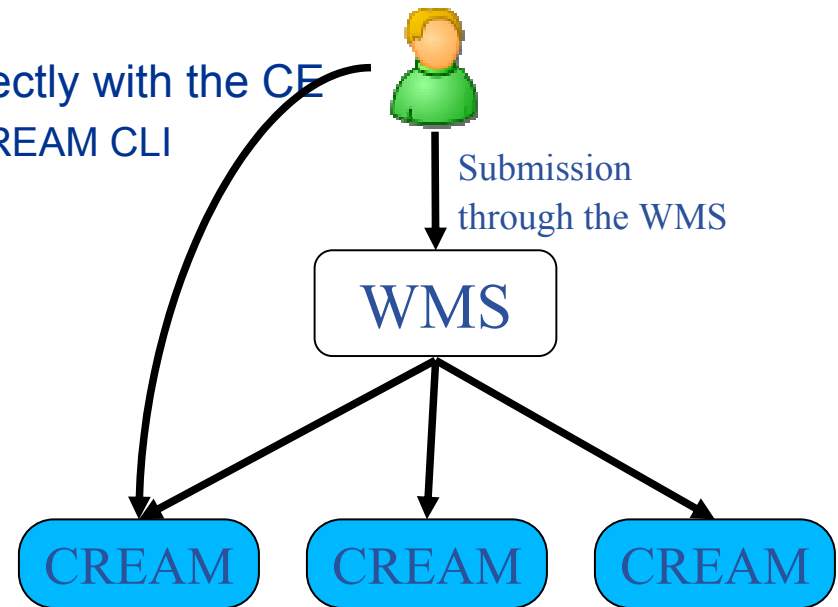
INFN Padova

On behalf of the CREAM product team

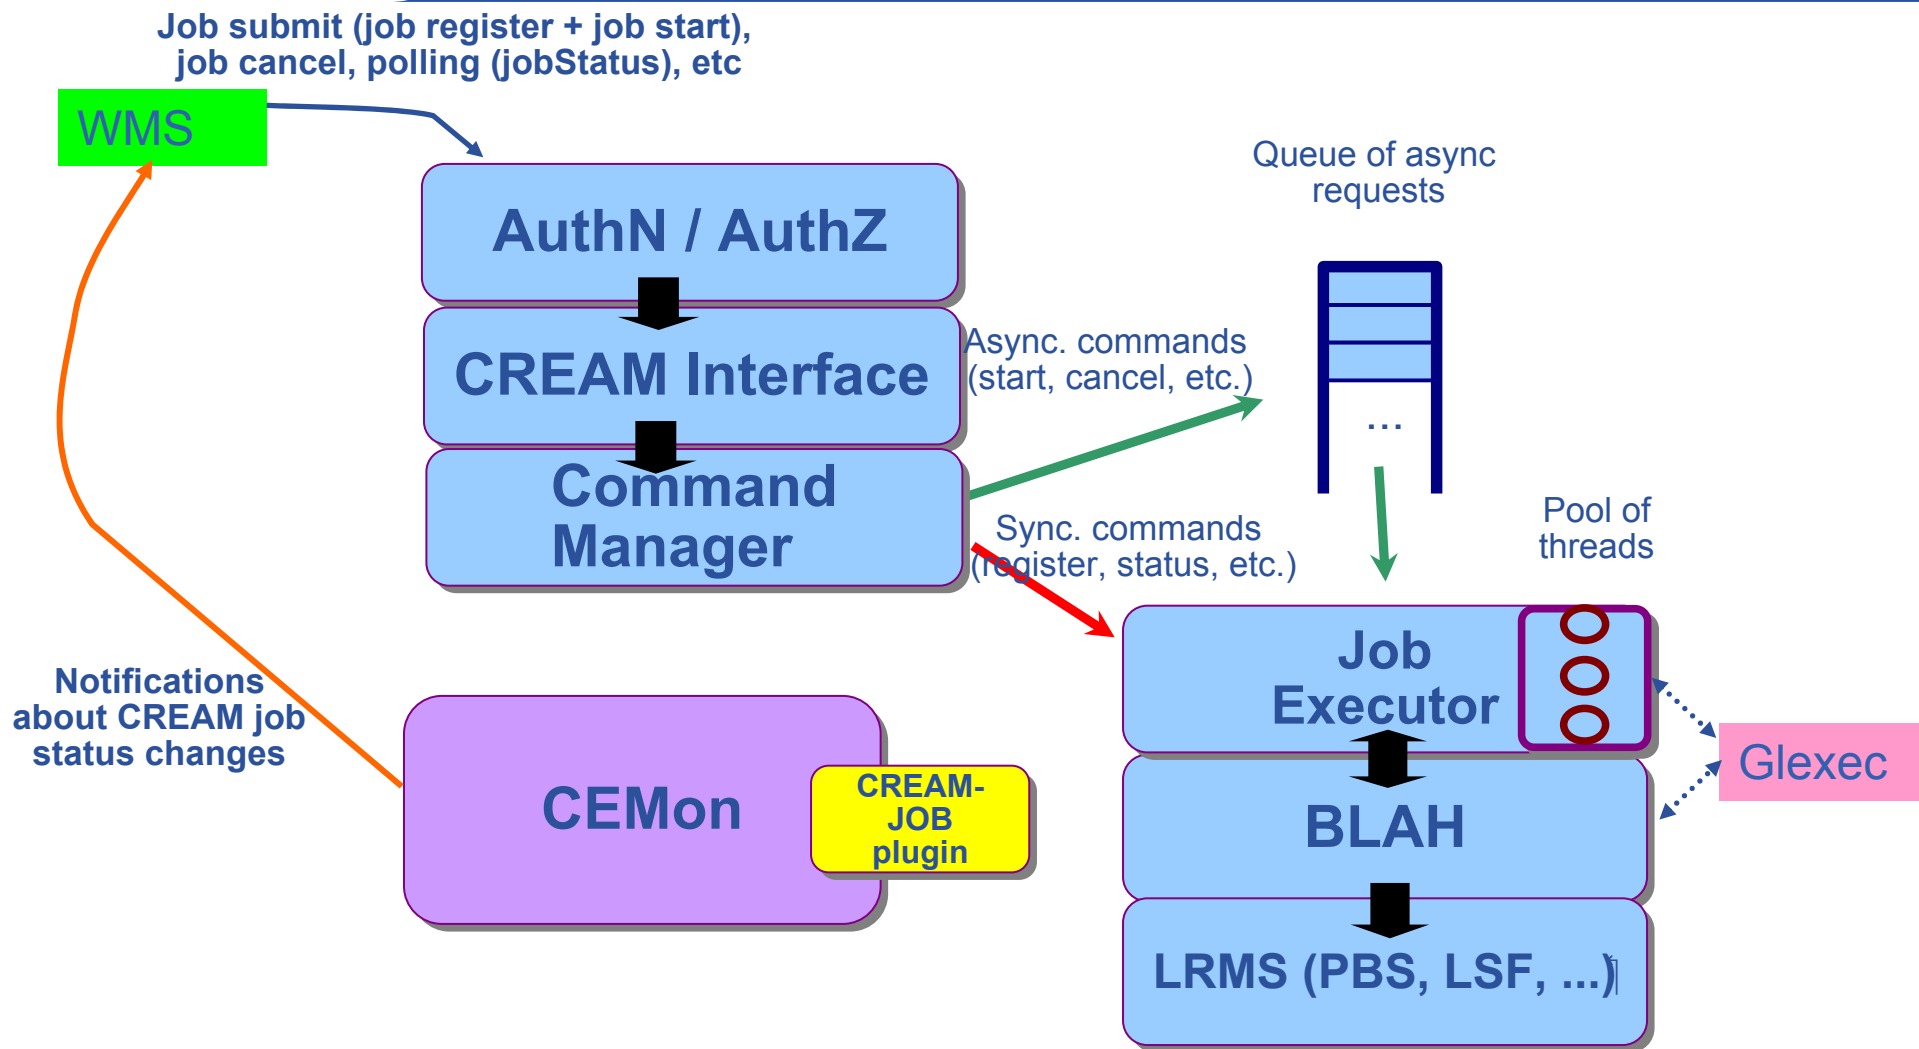
www.eu-egee.org



- **CREAM service: Computing Resource Execution And Management service**
- **Service for job management operations at the Computing Element (CE) level**
- **Allows to submit, cancel, monitor, ... jobs**
- **Web service interface**
- **CREAM can be used:**
 - through the gLite WMS
 - by a generic client willing to interact directly with the CE
 - We provide and maintain an “official” CREAM CLI
 - Users can build their own clients using a Web Service framework



“Simplified” view of CREAM-CE architecture



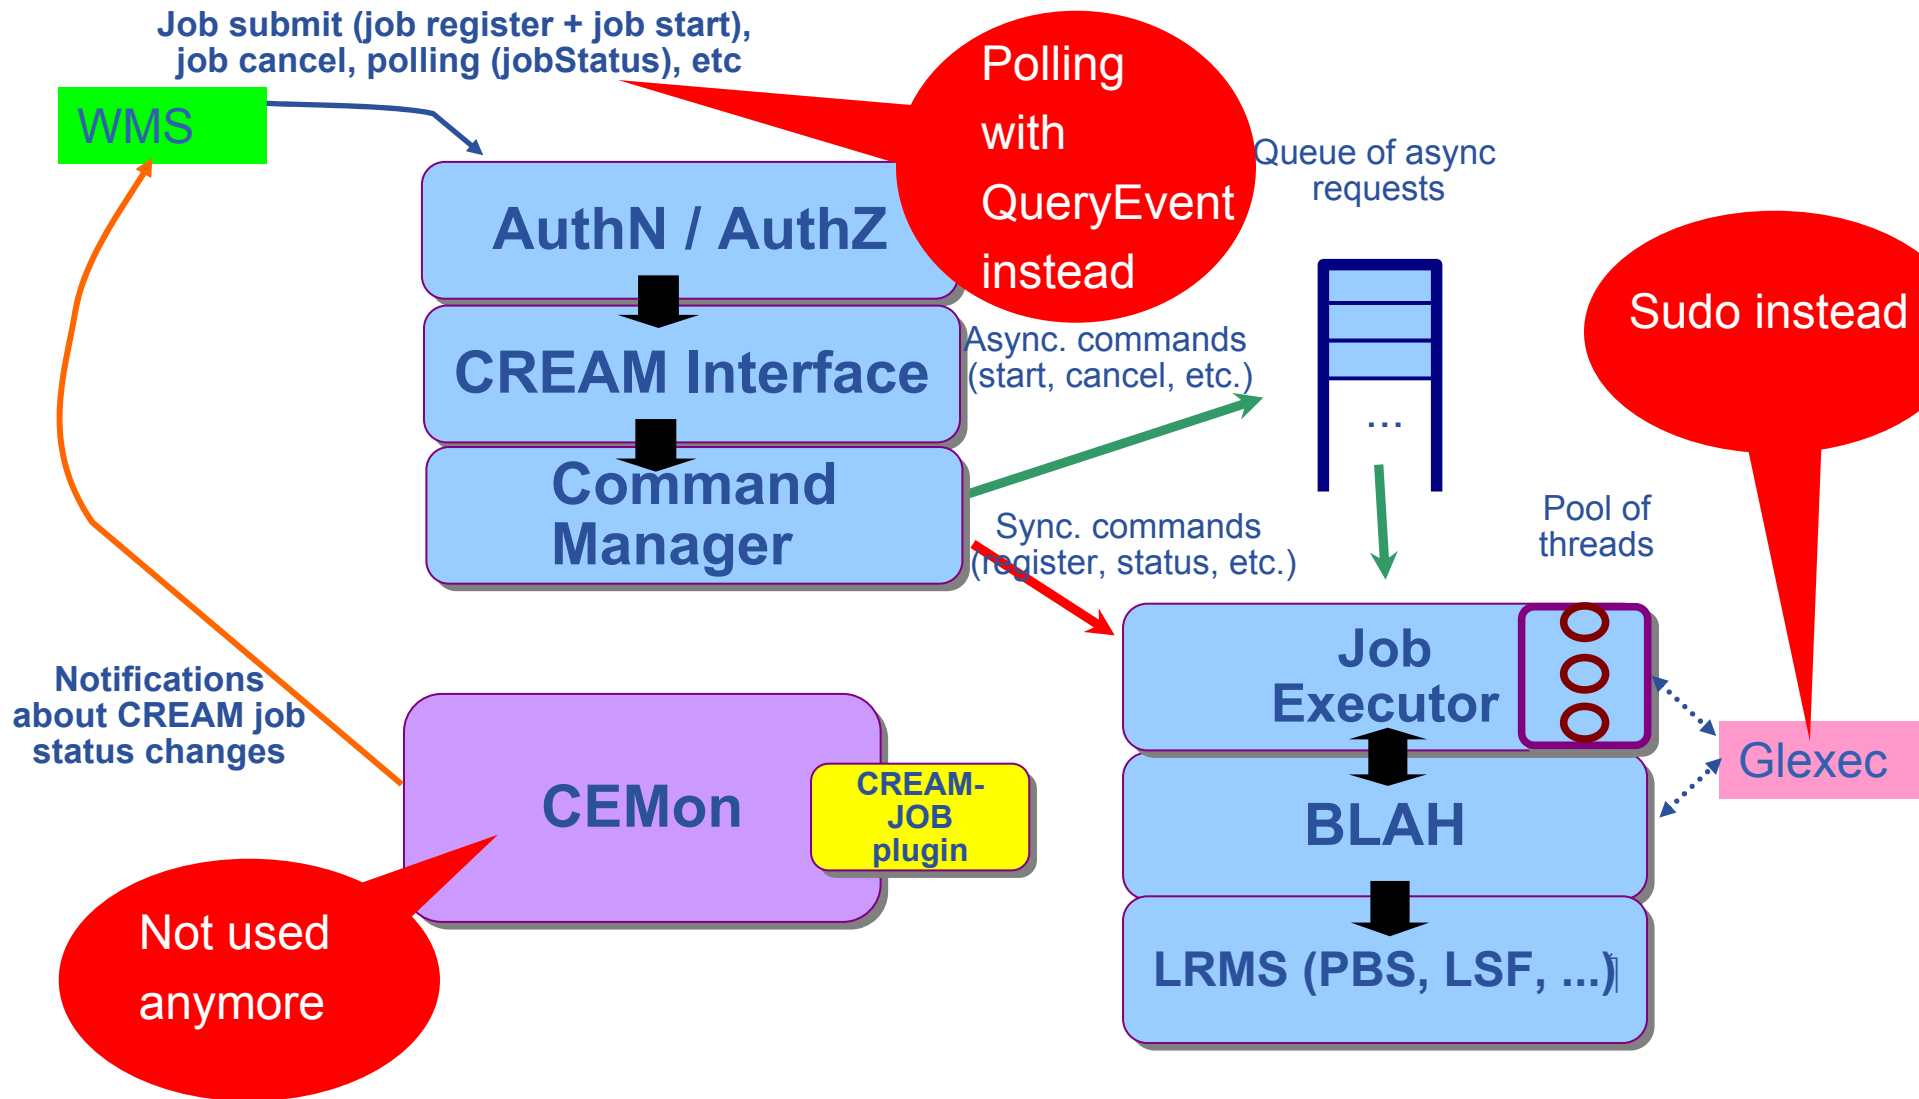
- CREAM and CREAM CLI in production since Oct. 2008
- Support in WMS (ICE) came later
- Initially publishing *Special* as *GlueCeStateStatus* to not get matched by default by the WMS
- SA1/WLCG invited sites to deploy CREAM CEs in parallel to their LCG-CEs
- On Oct. 2009 the WLCG MB decided that the CREAM CEs should publish *Production* as *GlueCeStateStatus*
- As of Jan 25, ~55 CREAM CEs published in the EGEE production BDII
- OSG is evaluating CREAM (Igor Sfiligoi, UCSD)
- Now in production: CREAM CE v. 1.5
 - On gLite 3.1/SL4_i386
 - On gLite 3.2/SL5_x86_64

- **Submission to CREAM from Condor is a requirement**
 - In particular by ATLAS
 - → No phase out of LCG-CE before
- **(Un)fortunately this is not in our control**
- **Condor team is working on it**
- **Last news: implementation done and given to Rodney Walker (ATLAS, Uni-Muenchen) for test**
 - CREAM GAHP crashes seen from time to time

- **LCG defined a set of criteria to be met before the phase out of the LCG-CE can start**
- <https://twiki.cern.ch/twiki/bin/view/LCG/LCGCEtoCREAMCETransition>
- **Most of them already satisfied (at least according to us)**
- **Some of them are not under our control**
 - E.g. support for submission to CREAM from Condor
 - On-going activity by Condor team
 - E.g. SAM/NAGIOS
 - E.g. SGE support
- **Some of the criteria are not too well defined**
 - Most of them waiting to be confirmed by operations/users
- **The known issues that we are aware of and for which we are responsible, addressed in next CREAM CE (1.6) and ICE (WMS 3.3) releases**
 - E.g. job status changes detection by ICE, limiter in CREAM CE

- **New operation QueryEvent**
 - To be used by ICE
 - To improve ICE's job status changes detection
 - Scalability problems in current (CEMon + polling) approach
- **Glexec → sudo**
 - Glexec used only once per job submission (just to find the uid to be used in sudo calls)
 - Better performance
 - Will facilitate the integration with Argus
- **BLAH release 1.14**
 - Support for sudo
 - New BLparser model for LSF and Torque/PBS
 - Using the status/history commands instead of parsing the log files
 - Also easier yaim configuration
 - Old parser still supported
 - Support of SGE (as contributed by CESGA)

“Simplified”view of CREAM-CE v. 1.6 architecture



- **Limiter to protect CREAM when the machine is overloaded**
 - New job submissions are blocked when this happens
 - Taking into account load, memory usage, # of file descriptors, etc.
 - Very similar to the limiter used in the WMS
- **Support of file transfers from/to gridftp servers started with user credentials**
 - Asked by Condor for Condor glidein
- **Several bug fixes**
- **Status**
 - Implementation done
 - Now testing and fixing what is wrong

- **Integration with Argus**

- Argus used to decide if a certain operation on a CREAM CE is authorized
- Also used to get the local user id
- Single AuthZ system in the CREAM CE
 - Now there is gJAF, LCAS/LCMAPS for glxexec, LCAS/LCMAPS for gridftp
 - *Because of bugs/misconfigurations inconsistent authZ decisions could be taken*
- Also gridftp integrated with Argus
- Not using gJAF and glxexec (and dependencies) anymore
 - Initially the old code will be maintained
 - At configuration possible to decide if Argus or “old system” has to be used

- **Bug fixes**

- **Initial integration with LB**
 - CREAM (and not only JobWrapper) logging to LB
 - Allows to monitor via LB all CREAM jobs
 - Not only the ones submitted through the WMS
 - Allows integration with dashboards (Alice request)
 - Depends on availability of LB 2.1
- **Support for bulk job submissions**
 - Submission of parametric jobs, job collections
 - In CREAM and CREAM CLI
 - To be used also by ICE
- **Bug fixes**

- **BES (Basic Execution Service) was supposed to be the standard interface for job management services**
- **Actually, even if the BES and JSDL specifications are in the final state, they are not suited for production use because they lack significant capabilities**
 - Used only in demos
- **A specific OGF working group, called Production Grid Infrastructure (PGI) WG has been set up to address these issues**
- **First proposal defined in the AGU (ARC Glite, Unicore) Lund workshop**
- **Being discussed and completed in the PGI WG**

- **An old CREAM implementation had some support for BES**
 - Implemented in the context of the OMII-EU project
- **Decided not to implement it in the current CREAM implementation for the above reasons**
- **We are instead following the PGI activities**
- **Activities going to be continued in the EMI project**

- **Support for high availability/scalability**
 - Pool of CREAM CE machines seen as a single CREAM CE
 - Some work already done
- **Common clients and APIs**
- **Better support for MPI jobs**
- **Some support for interactive jobs**
 - In order to monitor and/or steer the job while it is running
 - Hopefully possible integrating existing solutions

- **Service responsible for providing information in sync or async mode**
- **CEMon core + specific sensors, responsible to collect and provide specific information types**
- **Existing sensors**
 - CE sensor
 - CE information wrt Glue schema
 - Implemented as requested in the first phase of the EGEE-I project for the pull mode
 - *Possible to select policies specifying when the CE is willing to get jobs*
 - *The CEMon tells the WMS when it wants/it doesn't want jobs*
 - Actually never used
 - LHC-B from time to time expresses its interest in it
 - *They want the CE notifies their client when the CE is free*
 - GridICE sensor
 - A sensor publishing GridICE information
 - Implemented as requested by INFN-GRID
 - Never used

- **Existing sensors**
 - Job sensor
 - Provide information about CREAM job states
 - Used in the current implementation by ICE
 - As reported, not used anymore with CREAM CE v. 1.6 (QueryEvent used instead)
 - OSG-CE sensor
 - Provide CE information according to Glue schema for OSG needs
 - Used in ReSS
 - Maintained (just the sensor) by OSG people
- **CEMon core + OSG-CE sensor used in OSG**
 - Part of VDT
 - Deployed in several sites
 - They want to be assured about support after EGEE

- **“Cluster of competence” organization in place since a while**
 - Precertifications done since a while for all CREAM patches
 - Testing done by different persons than the ones who implemented
 - Testing done only on tagged (and locked) software
 - Testing done using CREAM CLI and WMS
 - Testing done on a testbed of about 20 CREAM CEs
 - CNAF and Padova
 - LSF and Torque
 - Glite 3.1/SL4 and Glite 3.2/SL5
- **Cluster of competence → product team**
 - Now also responsible for metapackage management
 - Now also responsible for certifications

<http://grid.pd.infn.it/cream>

[cream-support \[at\] lists \[dot\] infn \[dot\] it](mailto:cream-support@lists.infn.it)