Using CREAM and CEMon for job submission and management in the gLite middleware

C. Aiftimiei, P. Andreetto, S. Bertocco, S. Dalla Fina, A. Dorigo, E. Frizziero, A. Gianelle, M, Marzolla, M. Mazzucato, M. Sgaravatto, S. Traldi, L. Zangrando, P. Mendez Lorenzo, V. Miccio

www.eu-egee.org

e-infrastructure

What are CREAM and CEMon

- **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
- **CEMon**
  - General purpose event notification framework
  - CEMon core + sensors that can be plugged into the core
    - Sensor for CREAM job information, sensor for CE information, etc.
  - Can be used in synchronous or asynchronous mode
- **Web service interface**
- **Implemented and maintained within the EGEE project by the INFN Padova group**
- **Part of the gLite middleware distribution**
  - CEMon is also in VDT

CREAM: functionality

- **Job submission**
  - Supported job types: normal (sequential batch jobs), MPI, sub-jobs of collection/parametric jobs submitted through the WMS
  - Job characteristics described via a JDL (Job Description Language) expression, based on Condor classads
    - Basically the same used to submit to the gLite WMS
- **Proxy delegation**
  - To delegate a proxy, which can be used by the job to do operations requiring security support (e.g. GridFTP file transfers)
- **Job status**
- **Job cancellation**
- **Job list**
  - To get the identifiers of all your jobs submitted on a specific CREAM CE

CREAM functionality

- **Proxy renewal**
- **Job suspension and job resume**
- **Job purge**
  - To clear jobs from CREAM based CE
- **Disable/enable new job submissions**
  - Useful for example for a scheduled shutdown of the CREAM CE
  - Submissions can be automatically disabled also when a certain condition (on the number of active jobs) specified in the CREAM conf file is met
  - When submissions are disabled the other commands are still allowed

Job submission scenario

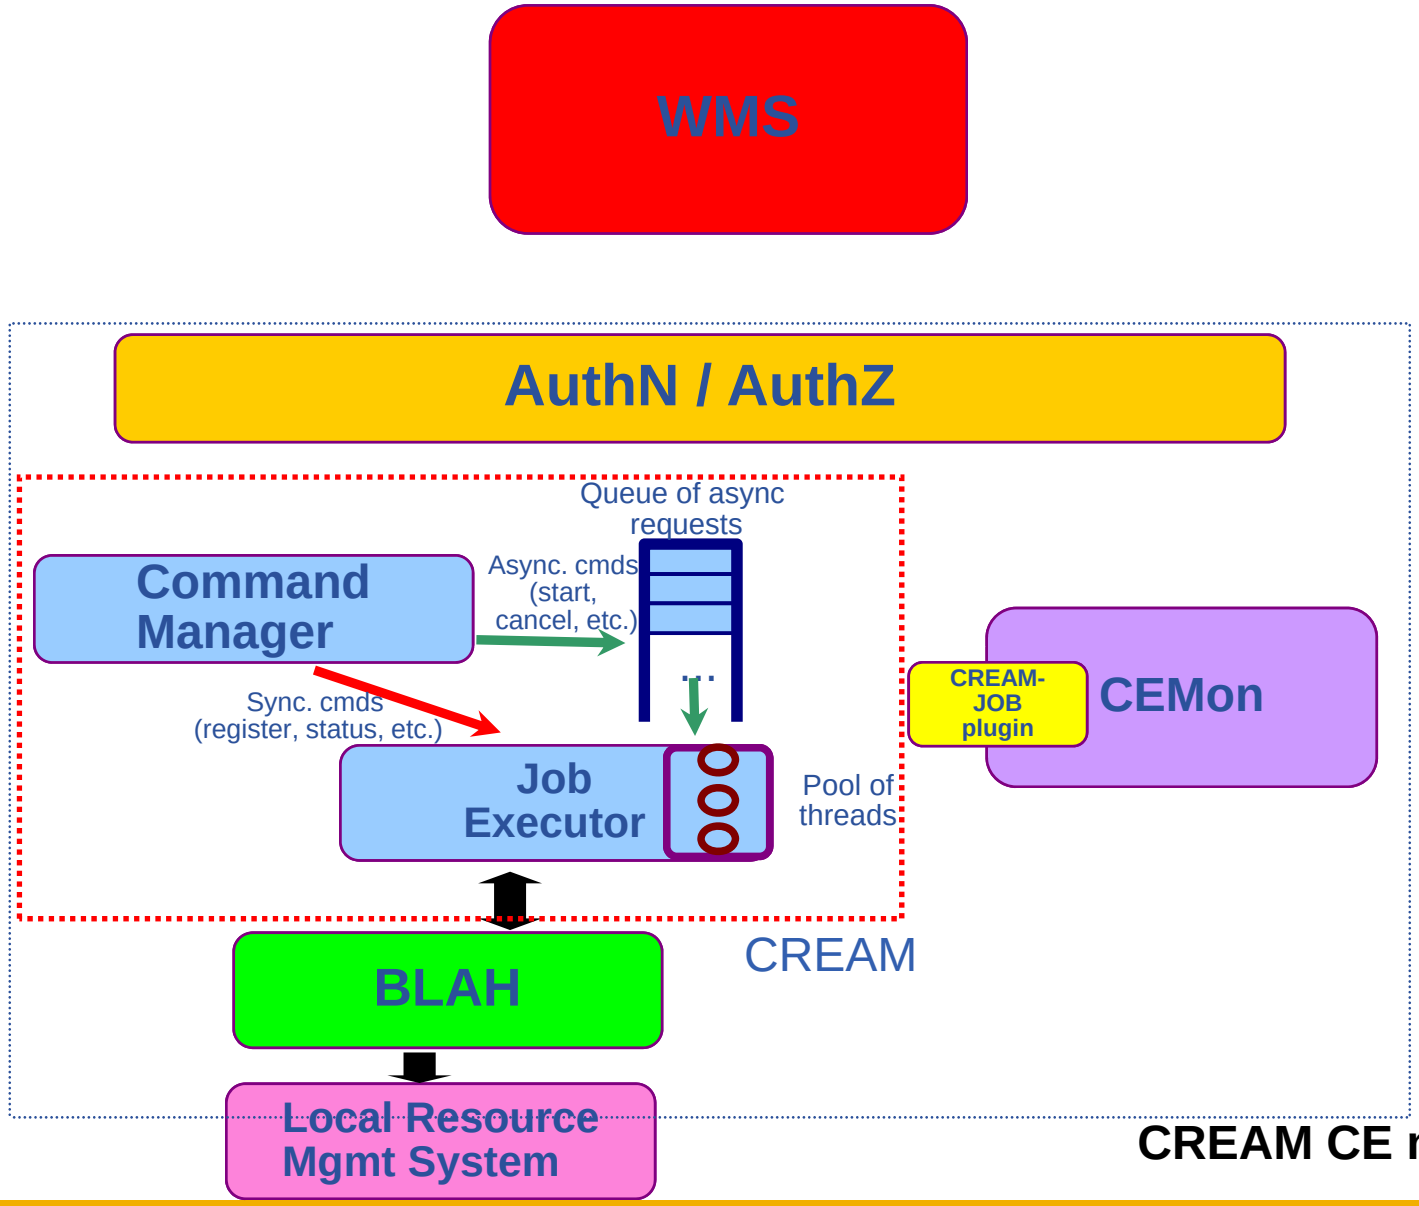

CREAM CE node

Job submission scenario

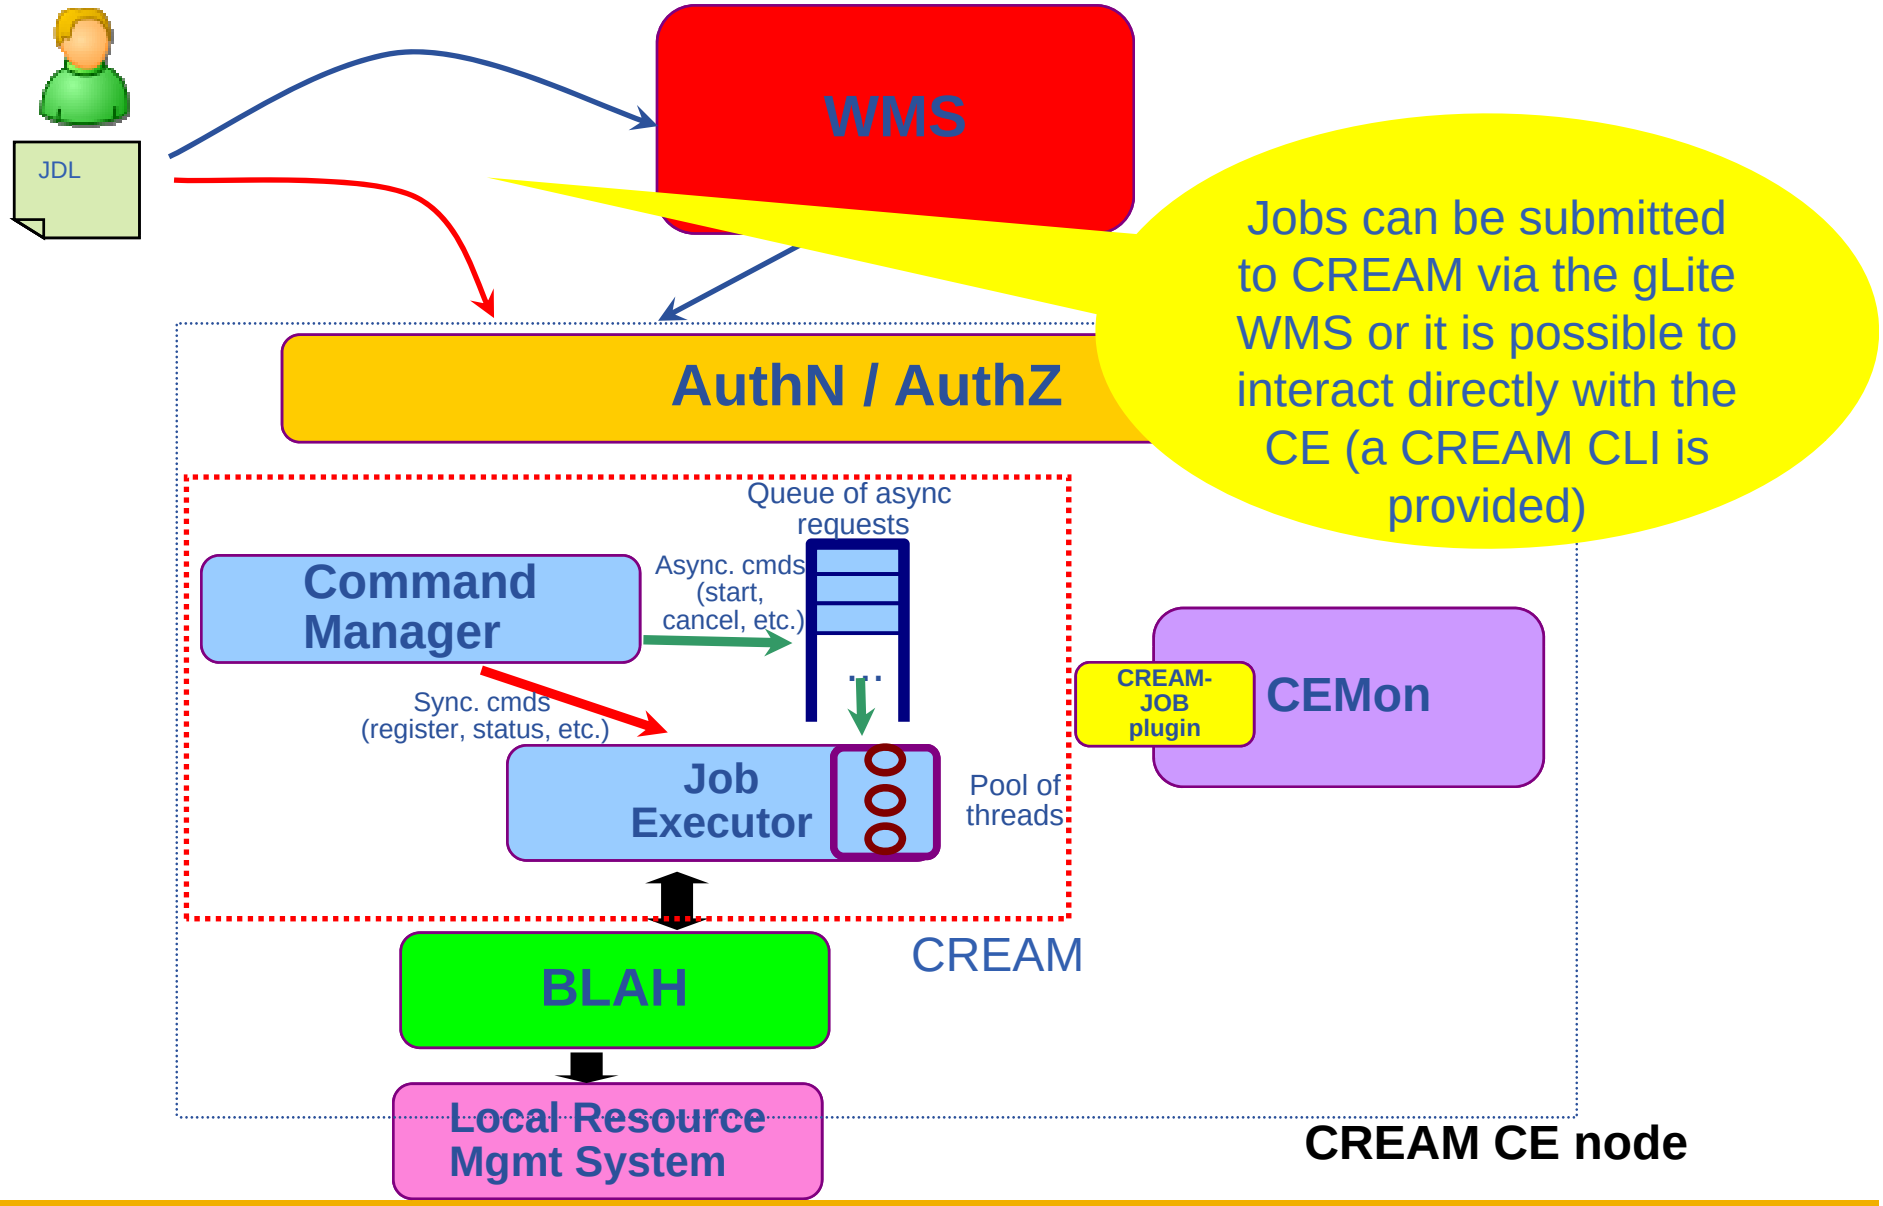

CREAM CE node

Job submission scenario

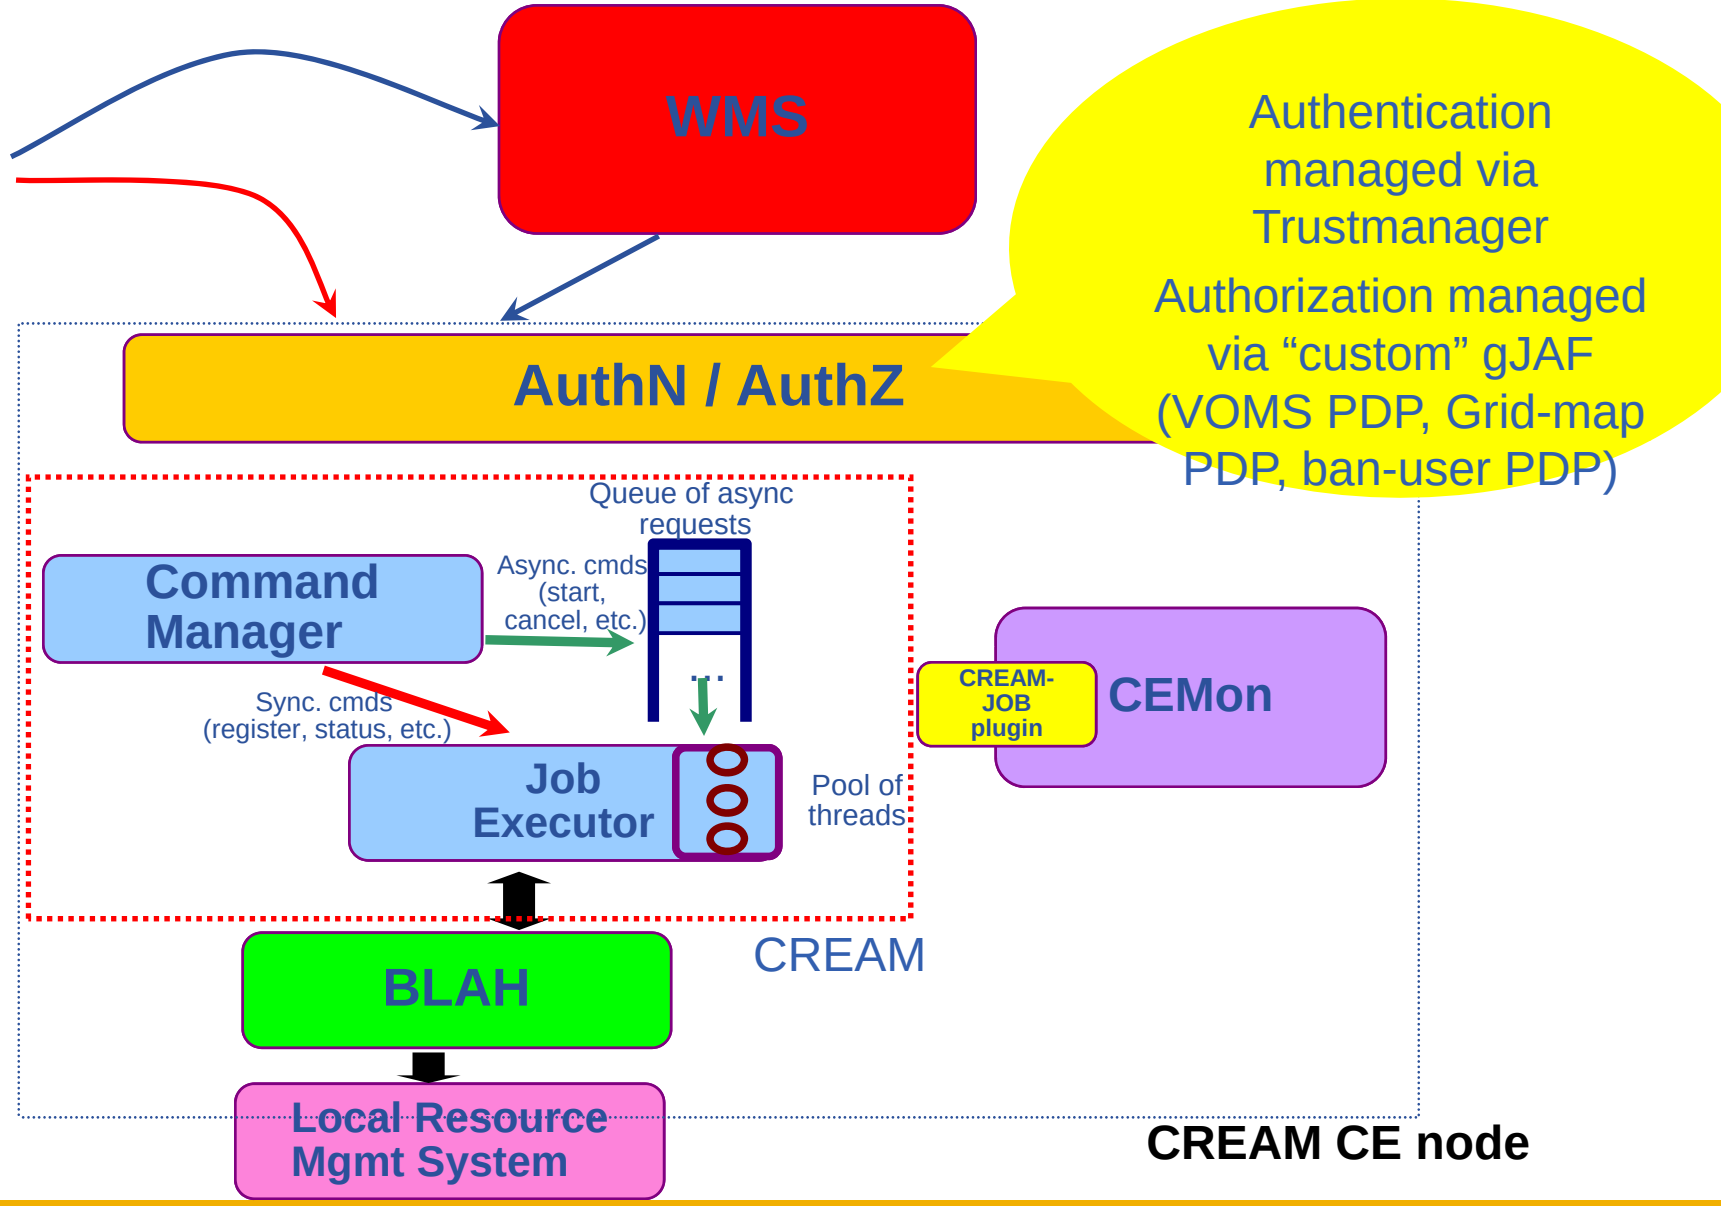

Job submission scenario

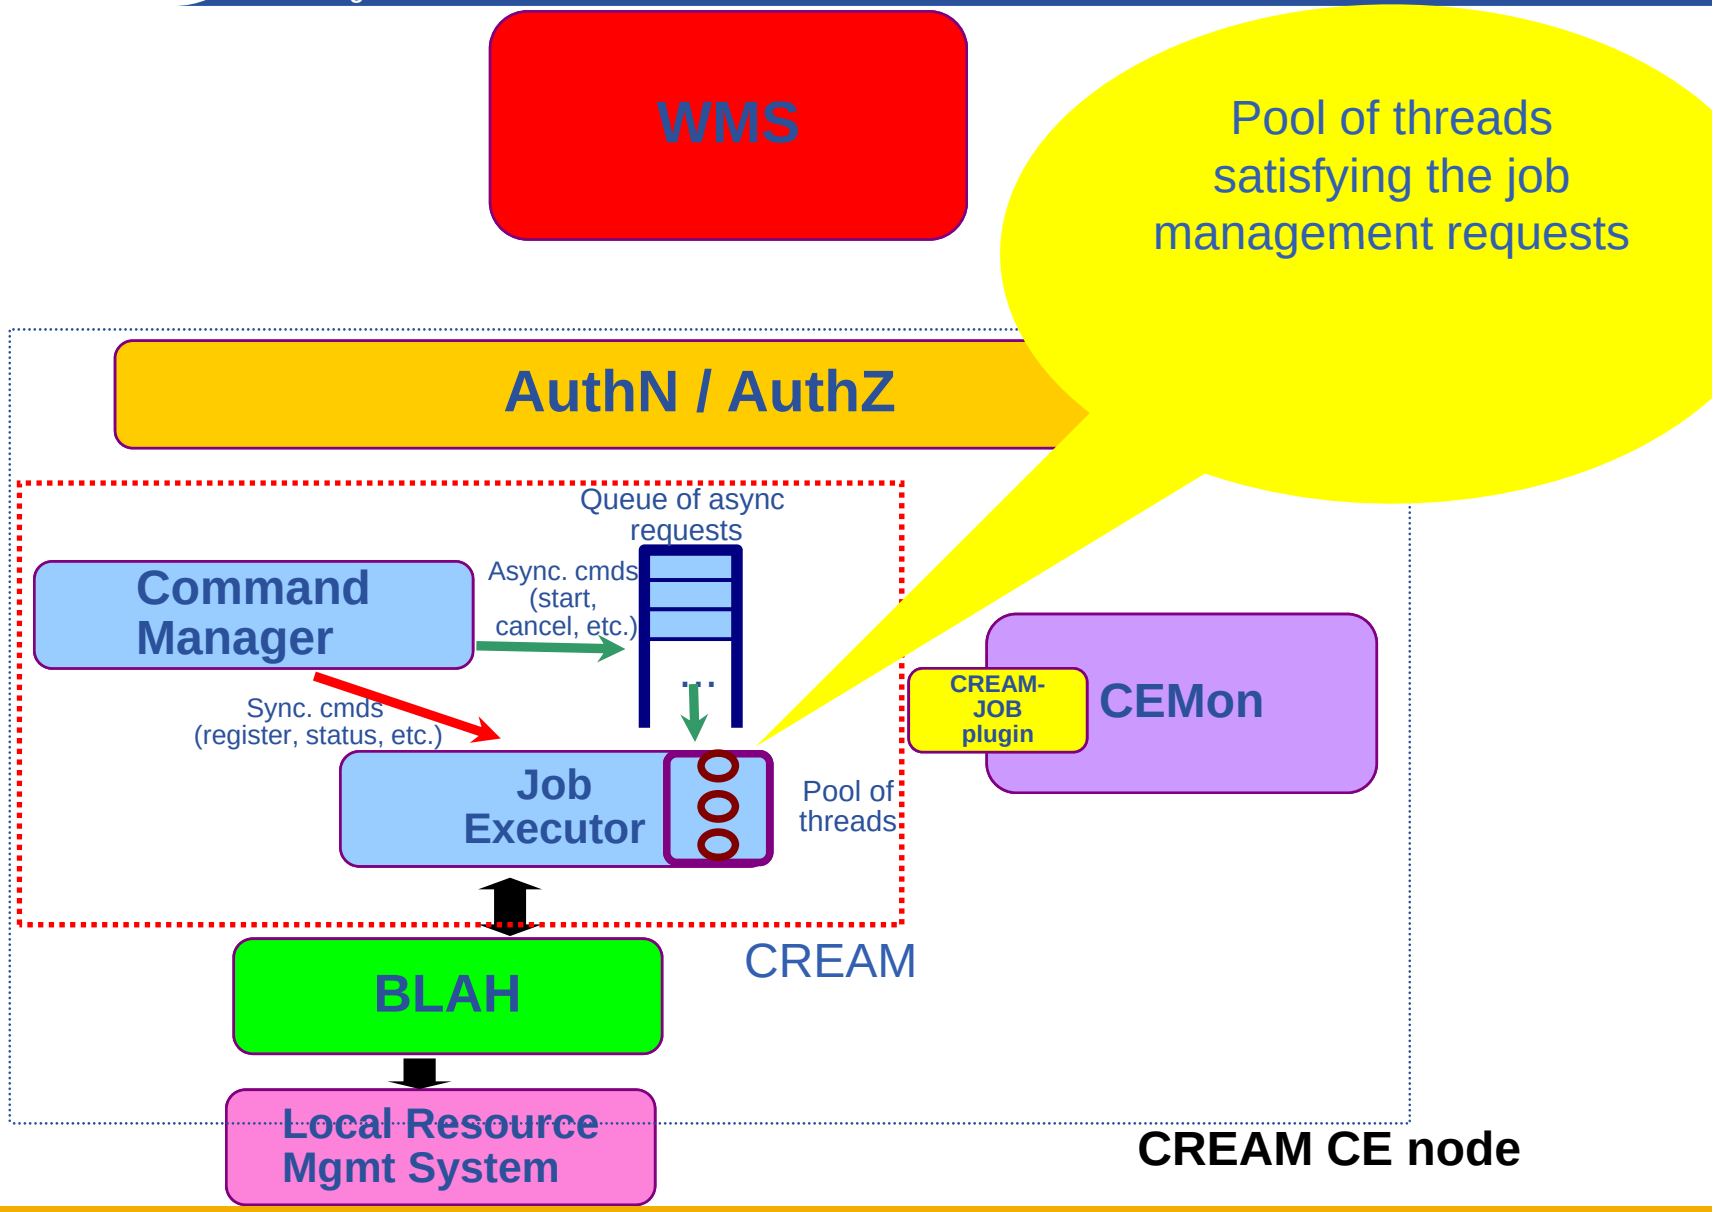

CREAM CE node

Job submission scenario

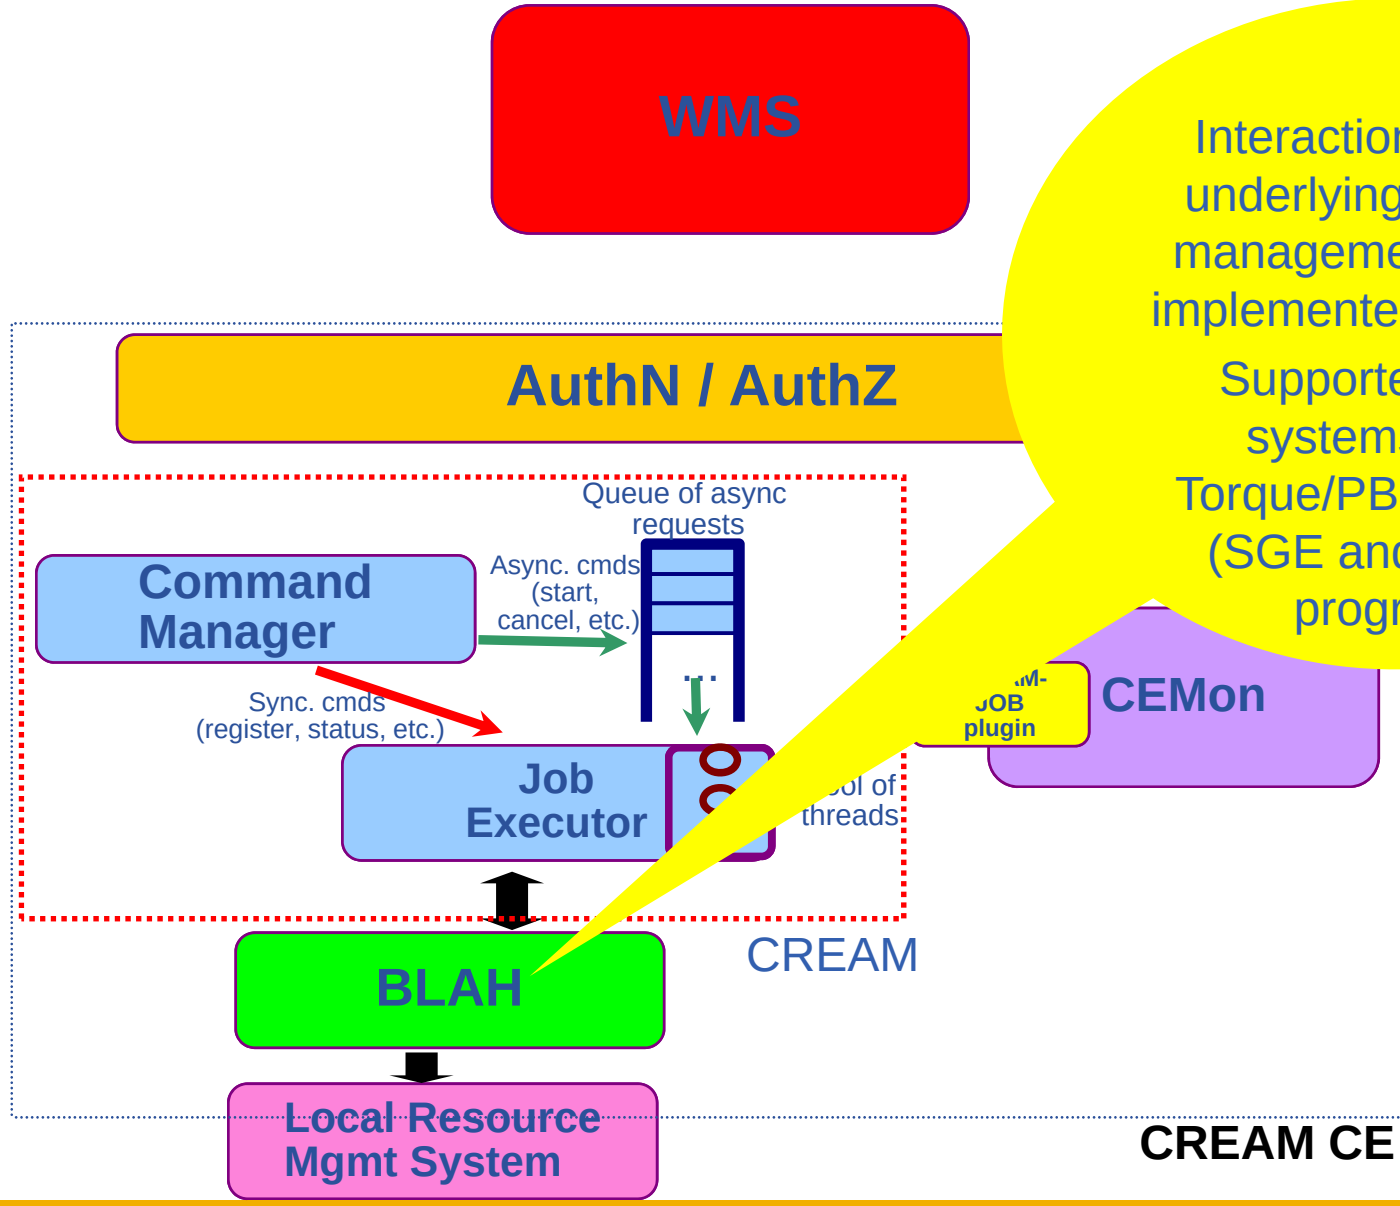

CREAM CE node

Job submission scenario

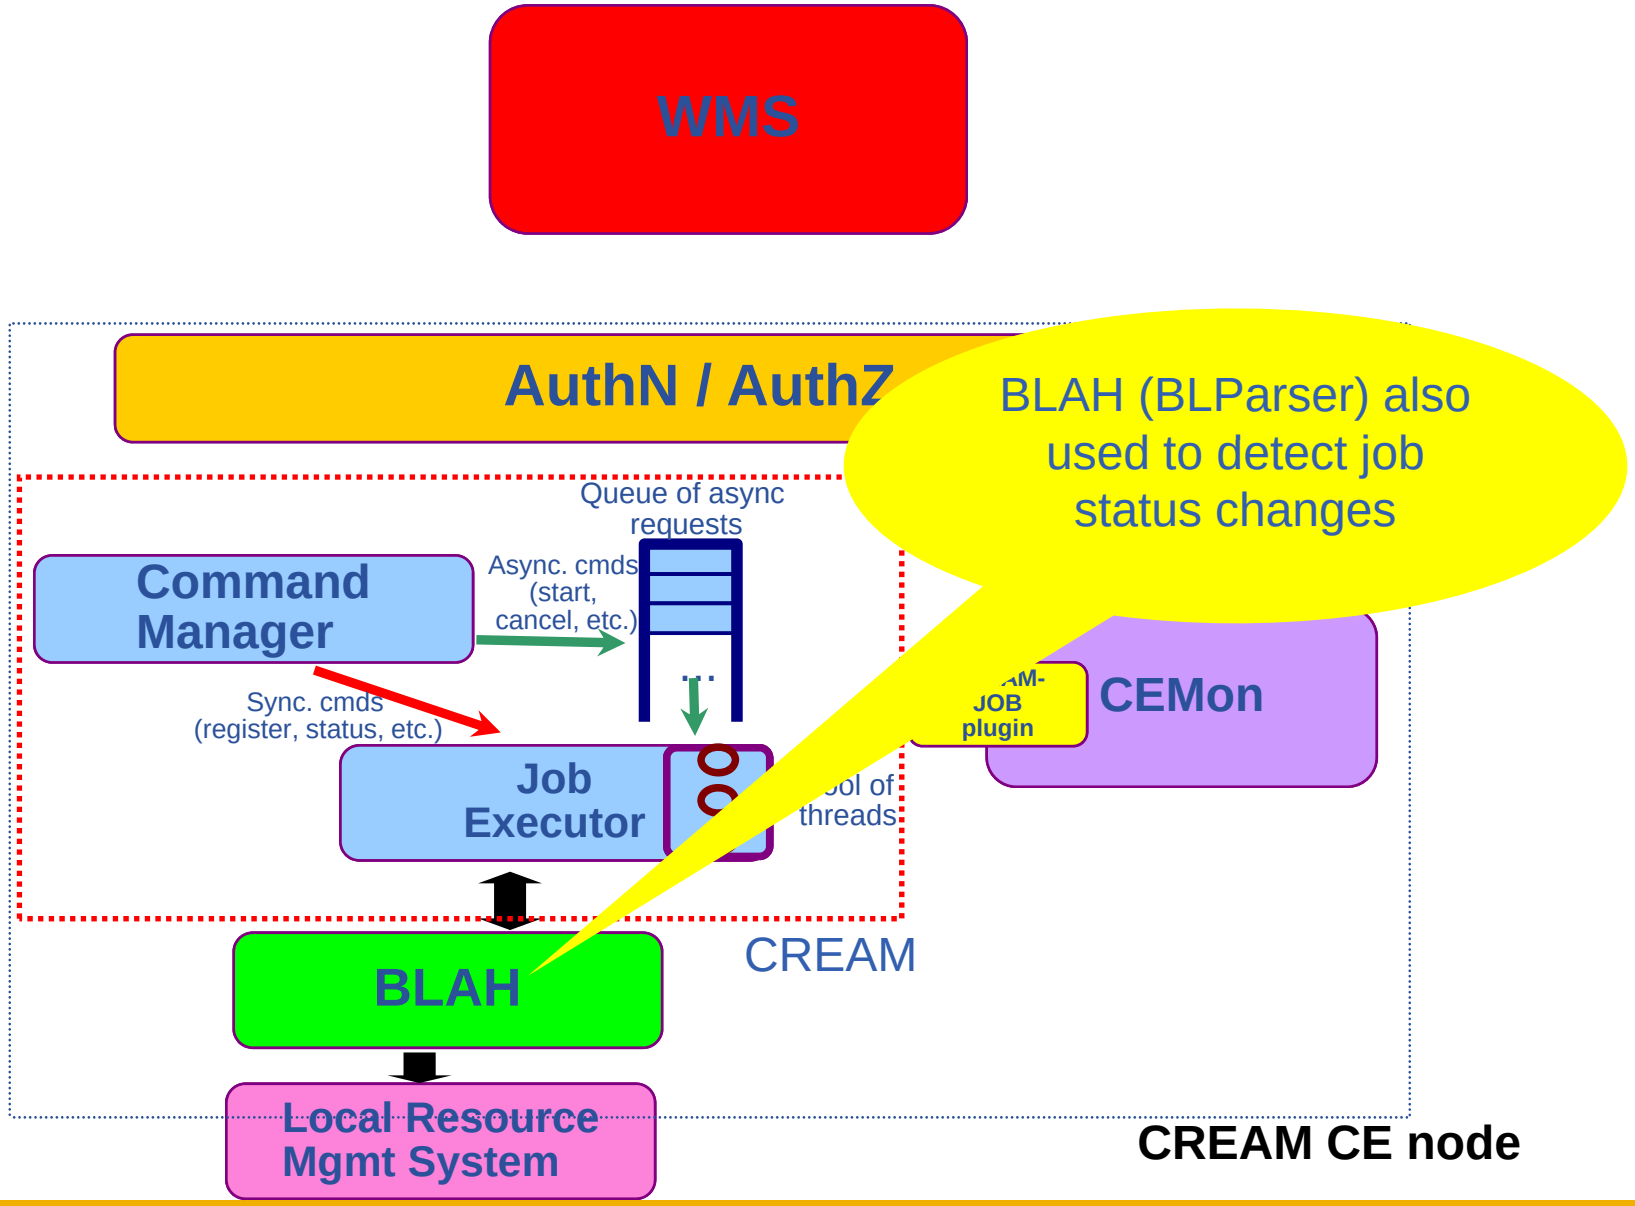

Job submission scenario

WMS

CEMon used to notify interested clients (e.g. WMS) about CREAM job status changes

AuthN / AuthZ

Queue of async requests

Command Manager

Async. cmds (start, cancel, etc.)

Sync. cmds (register, status, etc.)

Job Executor

Pool of threads

CREAM

Notifications about CREAM job status changes

CREAM-JOB plugin

CEMon

BLAH

Local Resource Mgmt System

CREAM CE node

Other info

- **Credential mapping**
  - Implemented via glexec
    - Glexec uses LCAS and LCMAPS
  - To map Grid credentials on local accounts and execute commands on behalf of these local accounts
- **Accounting**
  - Both APEL and DGAS can be used
- **CEMon is used also in OSG but for different purposes**
  - CE sensor (publishes information about CE characteristics and status according to the Glue Schema)
  - Used for resource discovery (ReSS project)
  - See poster #196 for more information
  - CEMon is in VDT

WMS-CREAM integration

- WMS-CREAM integration implemented via ICE (Interface to CREAM Environment)
- Daemon running on the WMS node
- Basically has the role played by JobController + LogMonitor + CondorG in the submission to LCG CEs
- ICE takes the job management requests from its filelist and satisfies them
- ICE also monitors jobs submitted to CREAM CEs and take appropriate actions

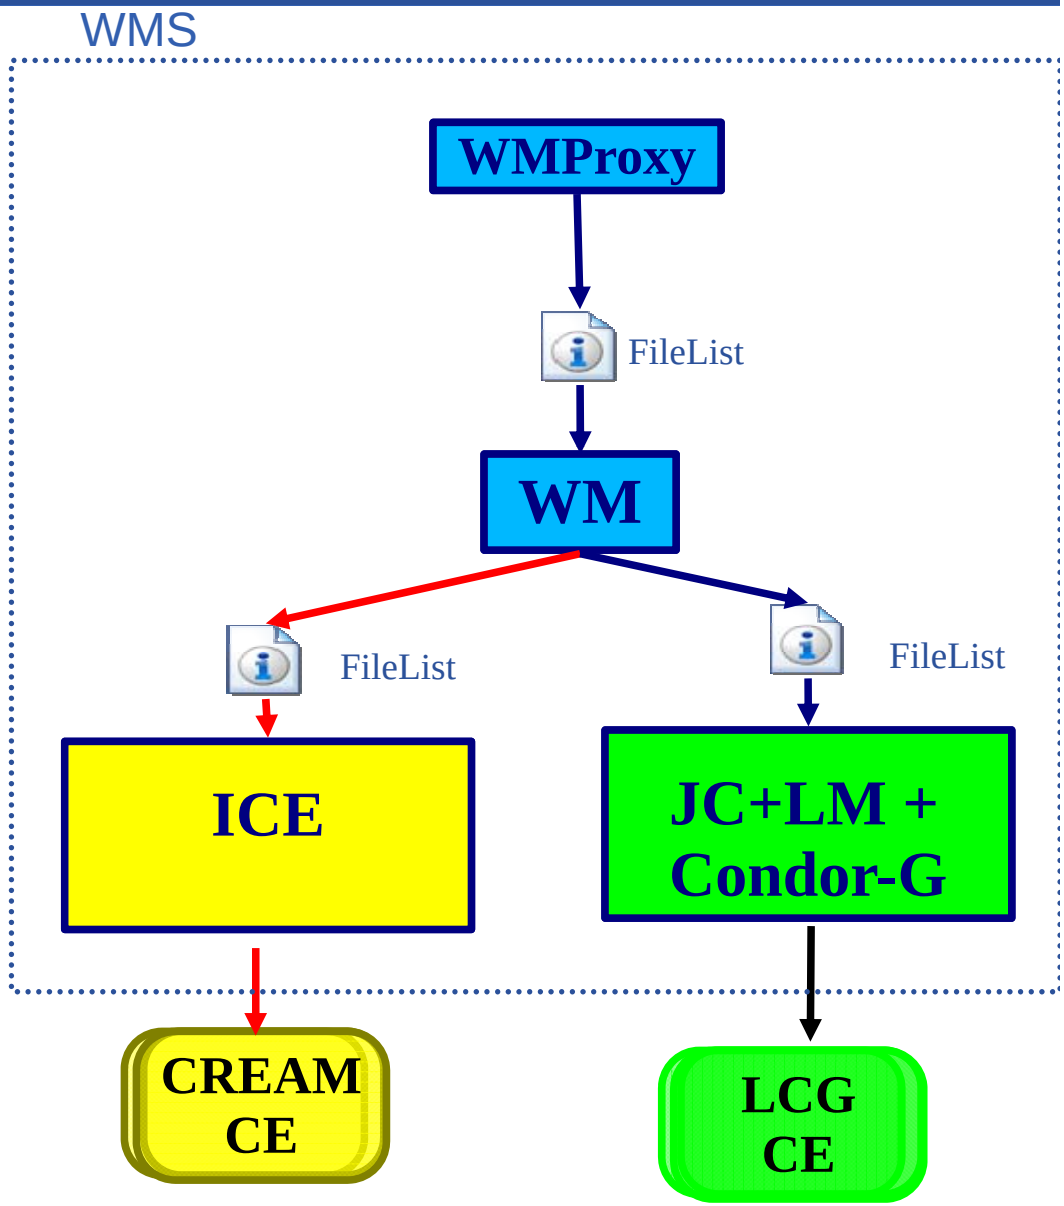

ICE: detecting CREAM job status changes

- CEMon with CREAM-JOB plugin coupled with CREAM
- ICE subscribes to CEMon to be informed about CREAM job status changes
- ICE receives these notifications from CEMon, detects CREAM job status changes and takes the appropriate actions
- As a fail-safe mechanism, ICE is also able to poll CREAM if the relevant notifications are not received via CEMon

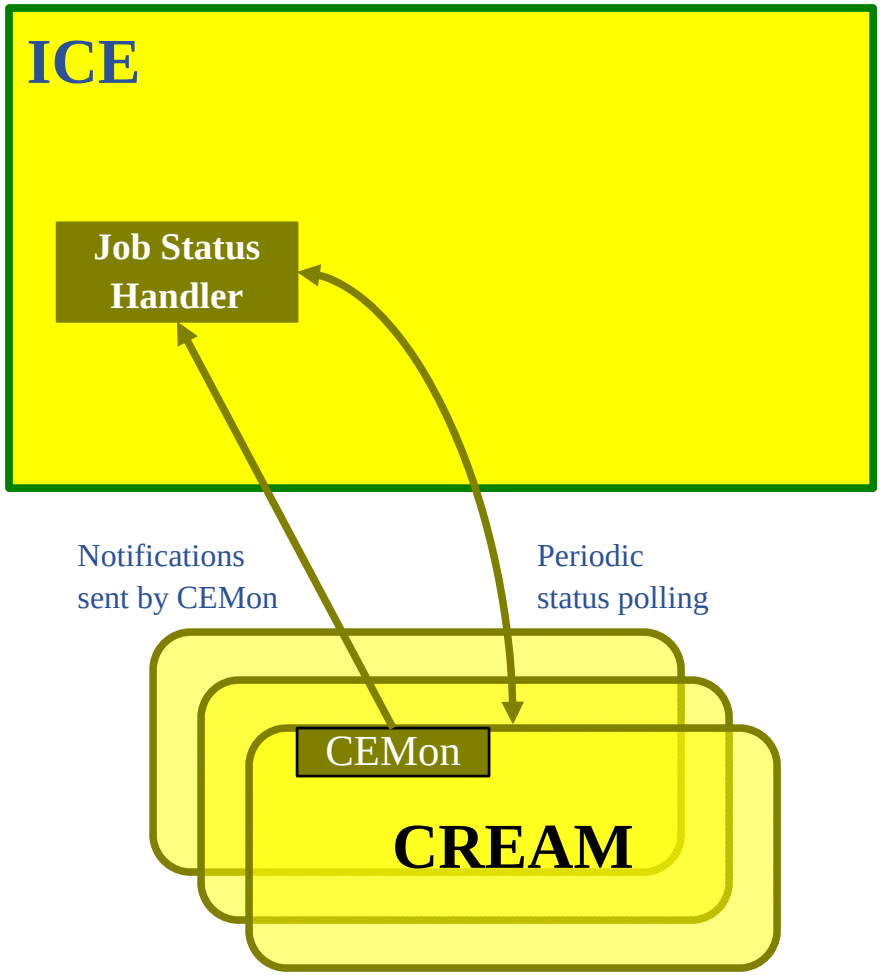

Deployment

- **CREAM released for production in EGEE in Oct 2008**
- **Since that, regular updates with bug fixes and improvements**
- **“Sites are encouraged to deploy a CREAM CE in parallel to their LCG CE”**
- **As of March 18, 14 CREAM CEs (115 CEIds) published in the EGEE production BDII**
  - Used in particular by Alice (see talk #106)
- **Also ICE (enabling submissions to CREAM through the WMS) released (released more recently than CREAM), even if there are still some scalability issues being addressed**
- **Defined criteria that must be met to start the transition from LCG-CE to CREAM**
  - http://twiki.cern.ch/twiki/bin/view/LCG/LCGCEtoCREAMCETransition
  - Functionality and performance criteria
  - Details of how/when/where doing these formal tests being defined

Some test results

- **Test condition**
  - Submissions to 21 CREAM CEs via the WMS
    - 14 CEs @ INFN Padova (LSF, Torque), 7 CEs @ INFN CNAF (LSF)
  - Submission of a collection of 40 jobs every minute
  - Short jobs (~ 5 minutes)
  - Used proxy renewal (initial proxy was 5 hours long)
  - Test duration: 5 days
  - Resubmission was enabled
- **Test results**
  - DONE OK: 99.2%
  - ABORTED: 0.0%
  - Not finished: 0.8%
    - Jobs stuck in Torque (problem in Torque then fixed)
  - Resubmissions: 1.60%
    - Most because of a problem with LSF at CNAF

Submission rate test

- **How the tests have been performed**
  - Submission of 1000 jobs by 4 users to the WMS
  - Different test conditions (proxy renewal enabled/disabled, explit/automatic/delegation)
- **How the measurements have been performed**
  - **Tstart** = LB timestamp of first ICE/JC dequeued event (i.e. request removed from the filelist, i.e. ICE/JC started its work)
  - **Tstop** = timestamp of last submission to batch system
- **Submission rate = # jobs / (Tstop - Tstart)**

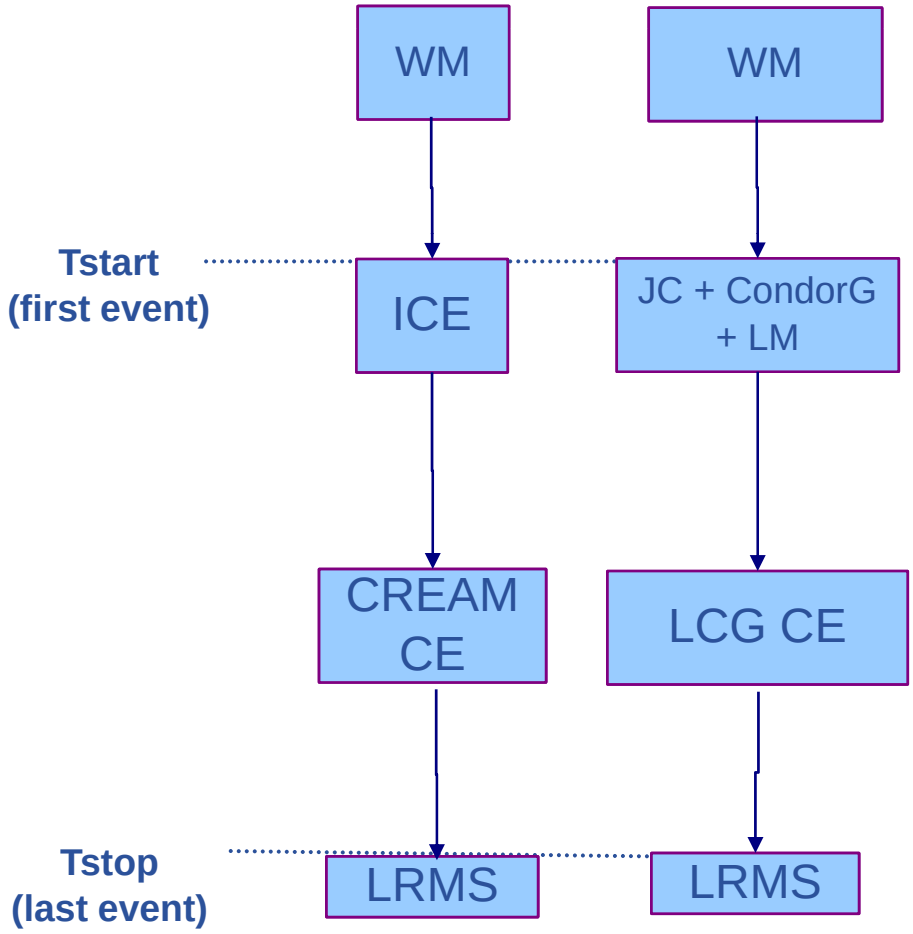

Submission rate test results (jobs/sec)

	ICE → CREAM	JC+CondorG+ LM → LCG-CE
Proxy renewal disabled (MyproxyServer="") Explicit delegation (glite-wms-job-submit –d …)	0.9624	0.3952
Proxy renewal disabled (MyproxyServer="") Automatic delegation (glite-wms-job-submit –a …)	0.1660	0.3633
Proxy renewal enabled (MyproxyServer="xyz") Explicit delegation (glite-wms-job-submit –d …)	0.8976	0.3728
Proxy renewal enabled (MyproxyServer="xyz") Automatic delegation (glite-wms-job-submit –a …)	0.9191	0.3863

Standardization

- **Besides the legacy interface, CREAM exposes also a BES-compliant interface**
  - BES (Basic Execution Service): recent OGF specification for a standard interface for Grid execution services
  - JSDL (Job Submission Description Language) used in BES to describe computational jobs
- **Actually this was done in an old CREAM implementation**
  - To be ported in the current CREAM implementation
- **BES and JSDL specifications are in the final state, but they are not suited for production use**
  - Significant capabilities are missing (e.g. security issues)
  - → A specific OGF WG (Production Grid Infrastructure, PGI) has been set up to define appropriate BES and JSDL profiles including all the important features requested for **production** Grid infrastructures
  - We are now more focused on this activity than in the support of the existing specifications
    - One person of the CREAM team is also co-chairing this PGI WG
  - http://forge.ogf.org/sf/projects/pgi-wg

Interoperability

- **Submission to CREAM from CondorG**
  - Basic functionality implemented and tested by US-CMS (CMS Glidein): see also talks #106 and #220
  - http://hepuser.ucsd.edu/twiki2/bin/view/HEPProjects/CMS-Cream
  - This is also one of the requirement that must be fulfilled for the transition to CREAM
- **Submission to CREAM from ARC**
  - ARC able to submit and manage jobs to CREAM
  - Also implemented Broker able to submit to CREAM CEs besides ARC resources
  - http://www.knowarc.eu/demos/Cream_CE_demo.avi
  - Work done in the context of KnowARC project

Build, Installation and configuration

- **Fully synchronized™ with gLite procedures**
- **Build done using ETICS**
- **Yum repository for installation**
  - For sl4_ia32 at the time being
  - Next will be sl5_x86_64, as decided by the EGEE management
  - Also sl5_ia32, Debian 4 and MacOSX for the client part
- **Yaim based configuration procedure**
- **Manual installation and configuration instructions available as well**

Some next steps

- **Address existing shortcomings and any other issues that will be found**
  - In particular fulfill all the requirements needed for the migration to CREAM
- **Bulk submission**
  - Submission of multiple jobs (e.g. a collection) to CREAM with a single call
- **High availability/scalable CE**
  - CREAM CE front end and pool of CREAM machines doing the work
- **Integration with new AuthZ service**
  - See talk #489

That's all

More info: http://grid.pd.infn.it/cream

http://grid.pd.infn.it/cemon

Contact us: jra1-pd@pd.infn.it