CREAM

Massimo Sgaravatto
INFN Padova
On behalf of the Padova CREAM group

What is CREAM

- **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*

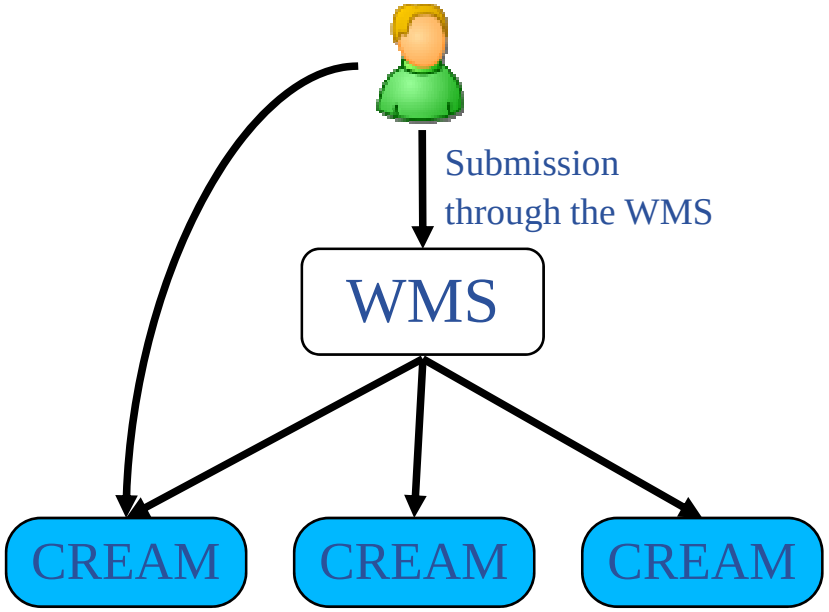

CREAM: functionality

- **Job submission**
  - Submission of jobs to a CREAM based CE
  - Supported job types: normal (sequential batch jobs), MPI, sub-jobs of collection/parametric jobs submitted through the WMS
    - So basically everything but sub-jobs of real DAGs
  - Job characteristics described via a JDL (Job Description Language) expression
    - CREAM JDL is the same (a subset) JDL used by the Glite WMS
  - Possibility to forward requirements to the underlying batch system (via BLAH)
    - See https://twiki.cnaf.infn.it/cgi-bin/twiki/view/EgeeJra1It/ForwardReqTests
    - Actually a couple of bugs to fix
- **Job status**
  - To get status and other info (e.g. significant timestamps, worker node, failure reason, issued commands on the job, etc.) of submitted jobs
  - Different levels of verbosity
  - Also possible to apply filters on submission time and/or job status
- **Job cancellation**

CREAM functionality

- **Proxy delegation**
  - To delegate a proxy, which can be used by the job to do operations requiring security support (e.g. GridFTP file transfers)
  - Possibility to automatically delegate a proxy for each job submission
  - Possibility to delegate a proxy, and then using it for multiple job submissions
    - Recommended approach wrt performance, since proxy delegation can be “expensive”
- **Job list**
  - To get the identifiers of all your jobs submitted on a specific CREAM CE
- **Proxy renewal**
  - To renew proxies for previously submitted jobs
- **Job suspension and job resume**
  - To hold and then restart jobs
- **Job purge**
  - To clear jobs from CREAM based CE

CREAM functionality

- **Disable/enable new job submissions**
  - Can be used only by CREAM CE administrators
  - Useful for example for a scheduled shutdown of the CREAM CE

    Admin > glite-ce-disable-submission grid005.pd.infn.it:8443

    User > glite-ce-job-submit -a -r grid005.pd.infn.it:8443/cream-lsf-grid02 test.jdl

    MethodName=[jobRegister] ErrorCode=[0] Description=[The CREAM2 service cannot accept jobs anymore] Timestamp=[Tue 22 Jan 2008 16:28:47]

    Admin > glite-ce-enable-submission grid005.pd.infn.it:8443
  - When submissions are disabled the other commands are still allowed
  - Submissions can be automatically disabled also when a certain condition (on the number of pending and/or idle and/or running jobs) specified in the CREAM conf file is met
    - E.g. a site administrator can decide to stop accepting new jobs when the site is already managing x jobs
- **Check if submissions are enabled**

  User > glite-ce-allowed-submission grid005.pd.infn.it:8443

  Job Submission to this CREAM CE is disabled

"Simplified"view of CREAM architecture

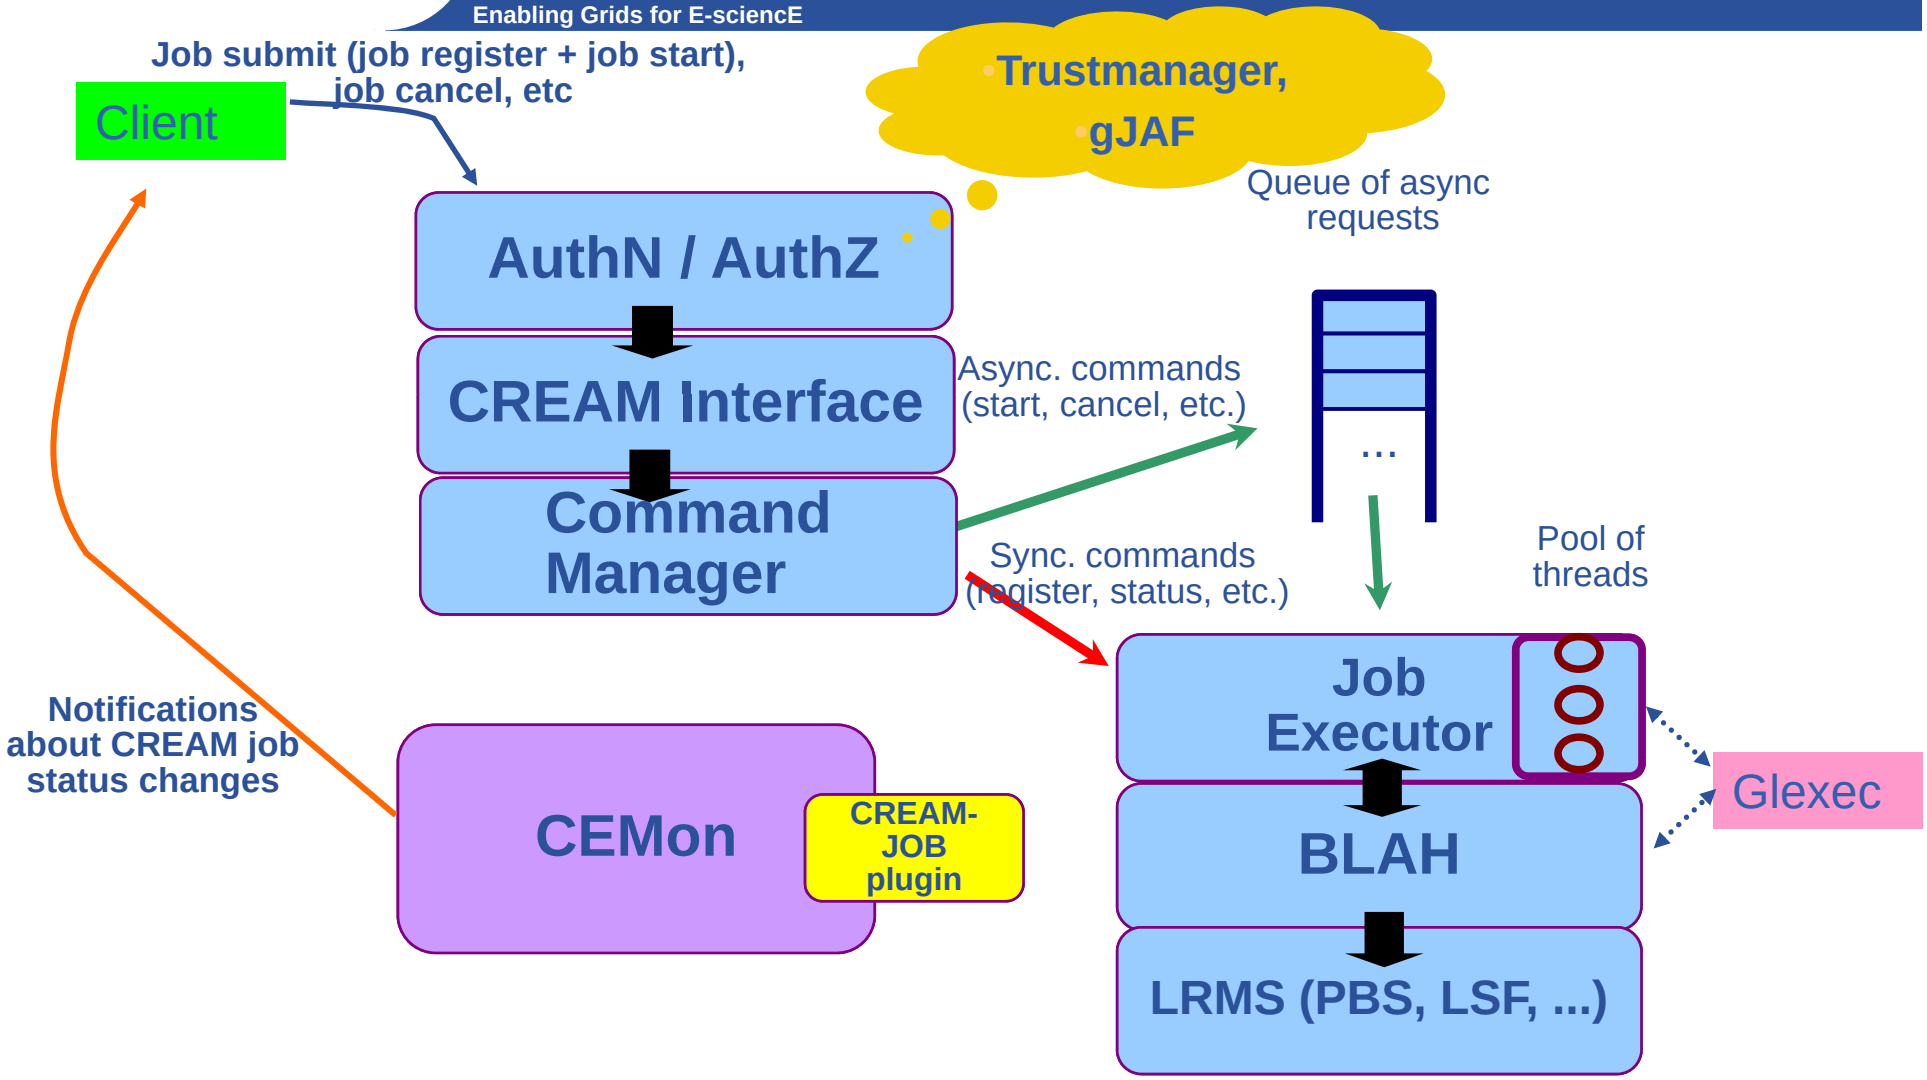

Interaction with LRMS

- **The interaction with the underlying local resource management system (LRMS) is fully managed by BLAH**
- **BLAH used to submit, cancel, etc. jobs on the batch system**
- **BLAH also used, via the BLParser, to notify CREAM about job status changes**
  - Two BLParser implementation models:
    - Old one: works parsing the batch system log files
    - New one: works referring to the batch system status/history commands
      - *New model done also to facilitate the porting to new batch systems*
  - Currently supported batch systems in BLAH:
    1. LSF: using old BLparser model
       - *On-going migration to new parser*
    2. Torque/PBS: using old BLParser model
       - *On-going migration to new parser*
    3. Condor: using new BLParser model
    4. SGE: reported as done, but not yet integrated in the official BLAH software

CREAM & InfoSys

- **No differences wrt LCG CE for what concerns the interaction with the Information Service**
  - CREAM CE characteristics and status published in the BDII, according to Glue Schema
  - Same information providers used in the LCG CE
  - Only some differences in the configuration of the static ldif file (e.g. CEID format) wrt the LCG CE
    - CREAM-CE → *<host>:<port>/cream-<lrms>-<queue>*
    - LCG-CE → *<host>:<port>/jobmanager-lcg<lrms>-<queue>*
- **Querying the BDII, it is possible to check what is "flavour" of a specific CE checking the *GlueCEImplementationName* attribute**
  - *GlueCEImplementationName: CREAM*
- **In the first phase CREAM CEs will not publish "Production" as GlueCEStateStatus (see later)**

Compliance to standards

- **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**
  - Aim: favor interoperability between different Grids
  - BES defines basic operations for job submission and management
    - BES itself does not mandate any specific security implementation
  - JSDL (Job Submission Description Language) used in BES to describe computational jobs
- **Standards addressed**
  - Basic Execution Service (BES) v. 1.0
  - Job Submission Specification Language (JSDL) v. 1.0
  - HPC Basic Profile v. 1.0
- **CREAM-BES related activities done so far in the domain of the OMII-EU project**
  - BES support done on an old CREAM implementation
  - Work in progress to finalize BES support in the current CREAM implementation

CREAM BES

- **Legacy interface**
  - JobRegister
  - JobStart
  - JobCancel
  - JobList
  - JobLease
  - JobInfo
  - JobPurge
  - JobSignal
  - JobSuspend / JobResume
  - JobProxyRenew
  - GetInfo
  - GetCEMonURL
  - EnableAcceptJobSubmissions
  - DisableAcceptJobSubmissions
  - DoesAcceptJobSubmissions

  } Job Submit

- **BES Interface**
  - CreateActivity
  - TerminateActivities
  - GetActivityStatuses
  - GetActivityDocuments
  - GetFactoryAttributesDocument
  - StopAcceptingNewActivities
  - StartAcceptingNewActivities

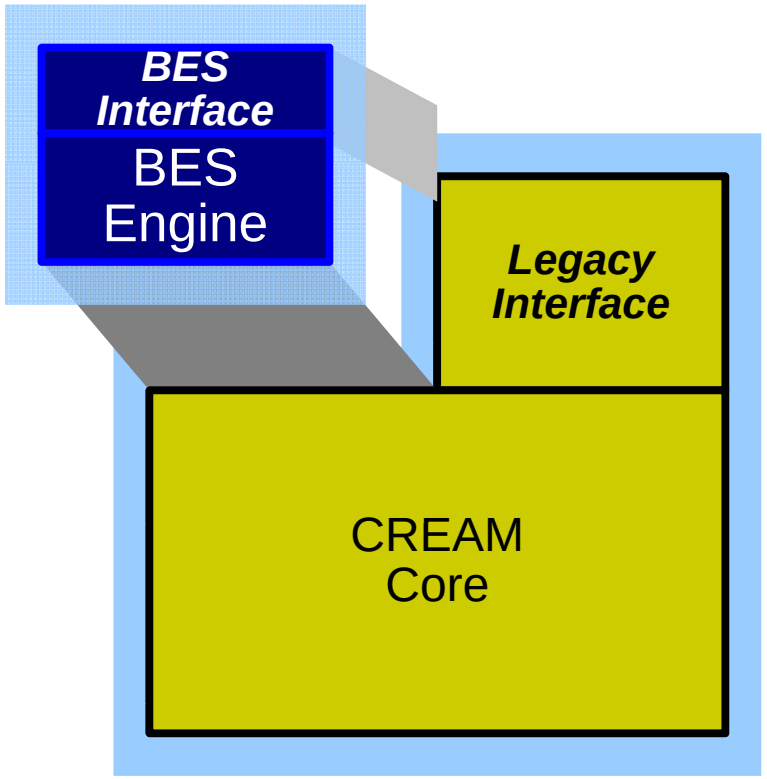

CREAM Releases

- **CREAM (patch #1755) and CREAM-CLI (patch #1790) patches prepared in April**
  - A patch for the WMS (#1854) to avoid matching CREAM based CEs was also released
  - The day before releasing CREAM in production (Aug 27), it was realized that not all production WMSes had patch #1854 deployed
  - So it was decided to not release CREAM and to make CREAM CEs publish something different than “Production” (“Special”), so they are not matched by the existing WMSes
    - This modification is in patch #2201 (yaim-cream-ce)
      - *Now in PPS*
- **Patches #2001 (CREAM) and #2002 (CREAM CLI) now under certification**
  - Bug fixes wrt patches #1755 and #1790
  - Actually we will to resubmit patch #2001, since we found a problem

Alice experiences with CREAM

- **Tests done by Patricia Mendez Lorenzo in the context of CREAM CE pilot service PPS activities**
  - Thanks !!!
- **Tests done directly submitting jobs to CREAM via CREAM CLI**
  - ALICE is interested in the direct submission feature of the CREAM CE (i.e. not via the WMS)
- **The whole setup has been provided at FZK (thanks to Angela Poschlad at the site)**
  - Operated through a VOBOX parallel to the already existing service at the T1
    - CREAM CLI installed on this VOBOX
  - Access to the CREAM CE
    - Initially 30 CPUS (PPS) available for the testing
    - At the end of the first testing period, the standard ALICE production queue at FZK was made available to the CREAM CE

Alice experiences with CREAM

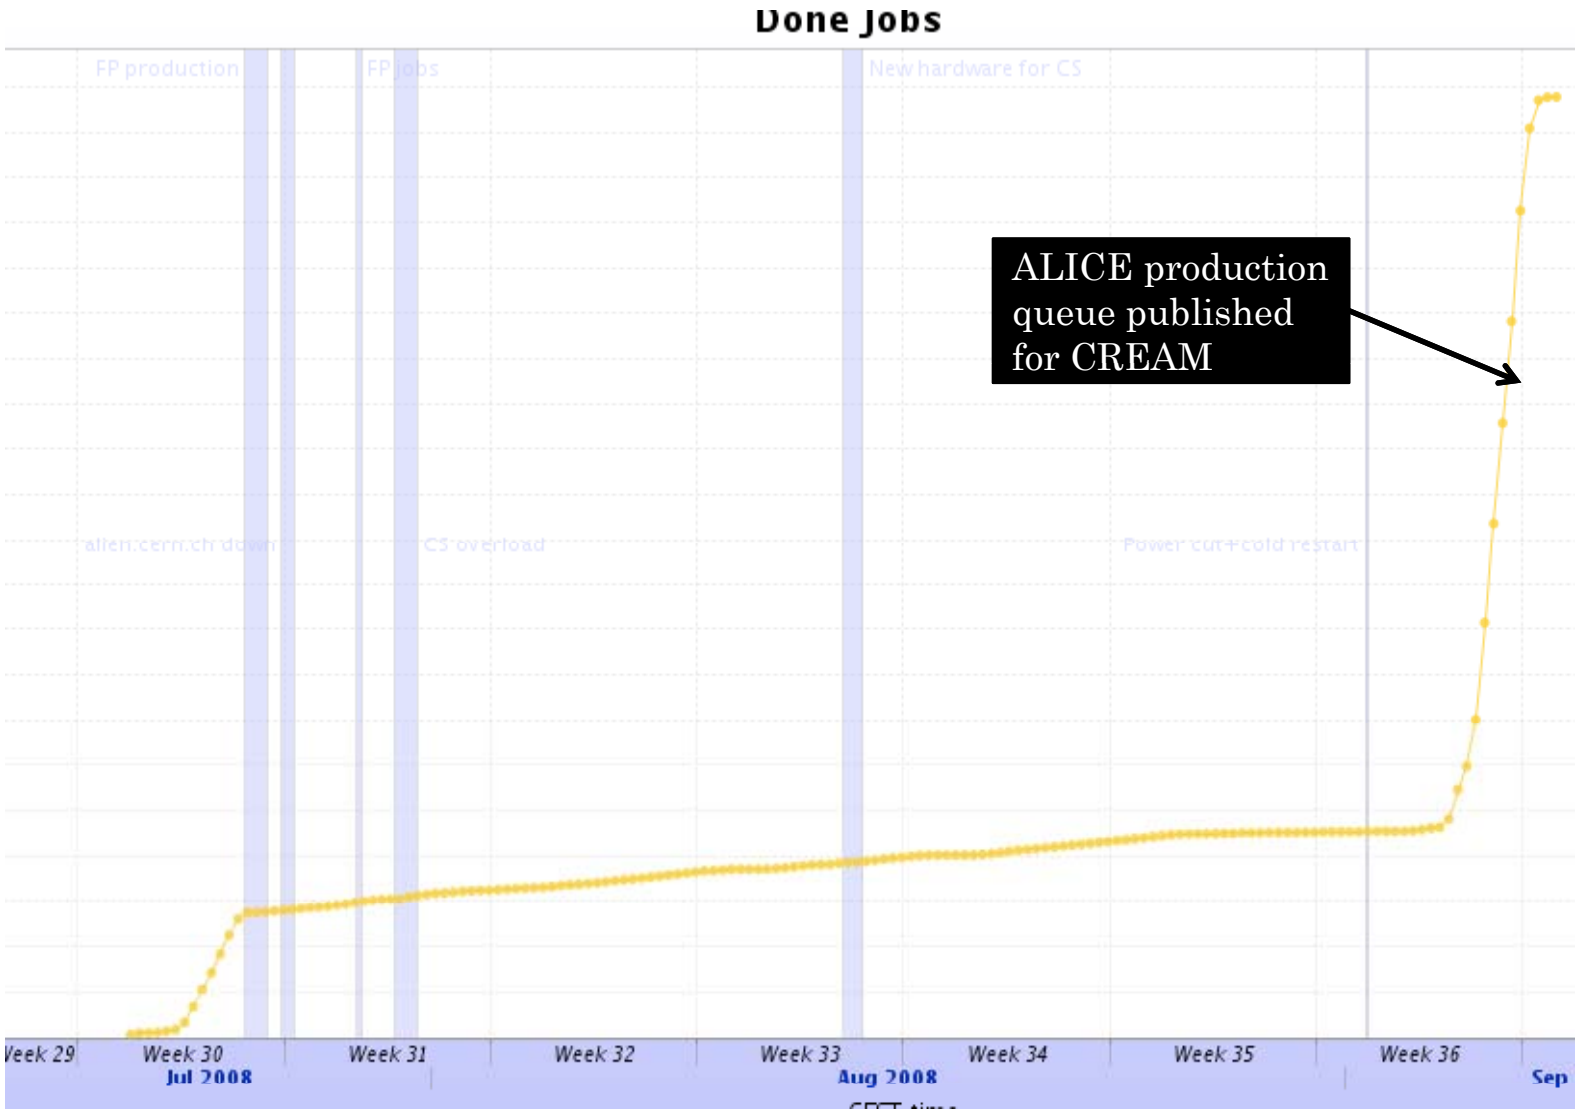

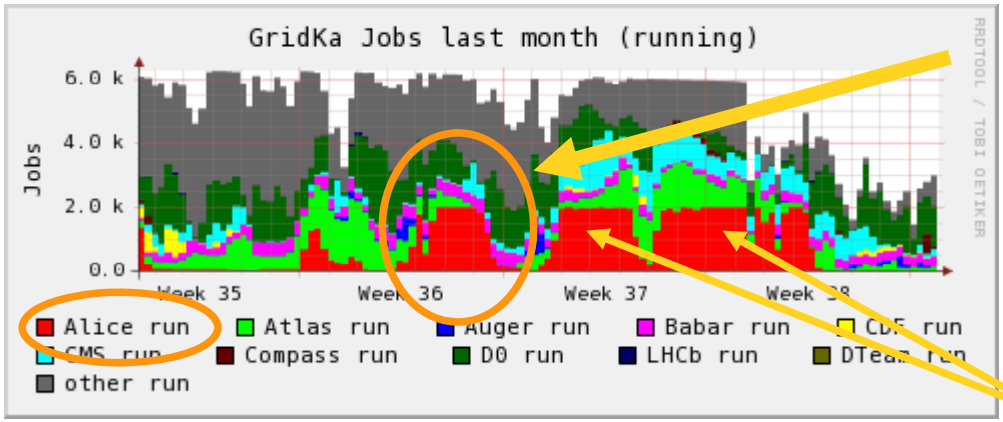

ALICE production jobs via CREAM CE (ca. 2000)

Alice jobs via lcg-CE

The two CEs used have the same hardware

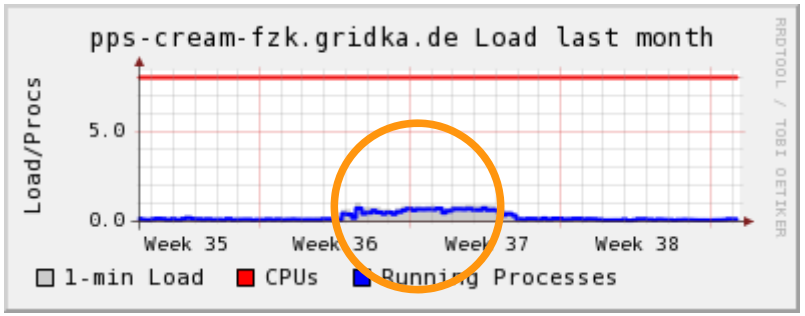

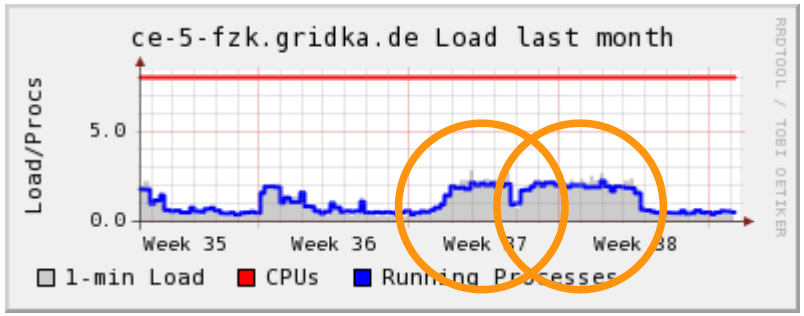

Alice experiences with CREAM

- **More than 55000 jobs submitted**
  - Standard production Alice jobs
  - Rough average: 12 h per job
- **More than 2000 concurrently jobs managed by CREAM**
- **During the entire test period the service showed remarkable stability**
  - No interventions on the vobox
  - No interventions on the CREAM CE
- **Alice wishes that the CREAM CE is put in production**
  - In a secure way, ensure a transition phase by the deployment team
- **Several sites have expressed their interest to provide a CREAM CE if it is in an standard gLite distribution**

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 JC+LM+Condor 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

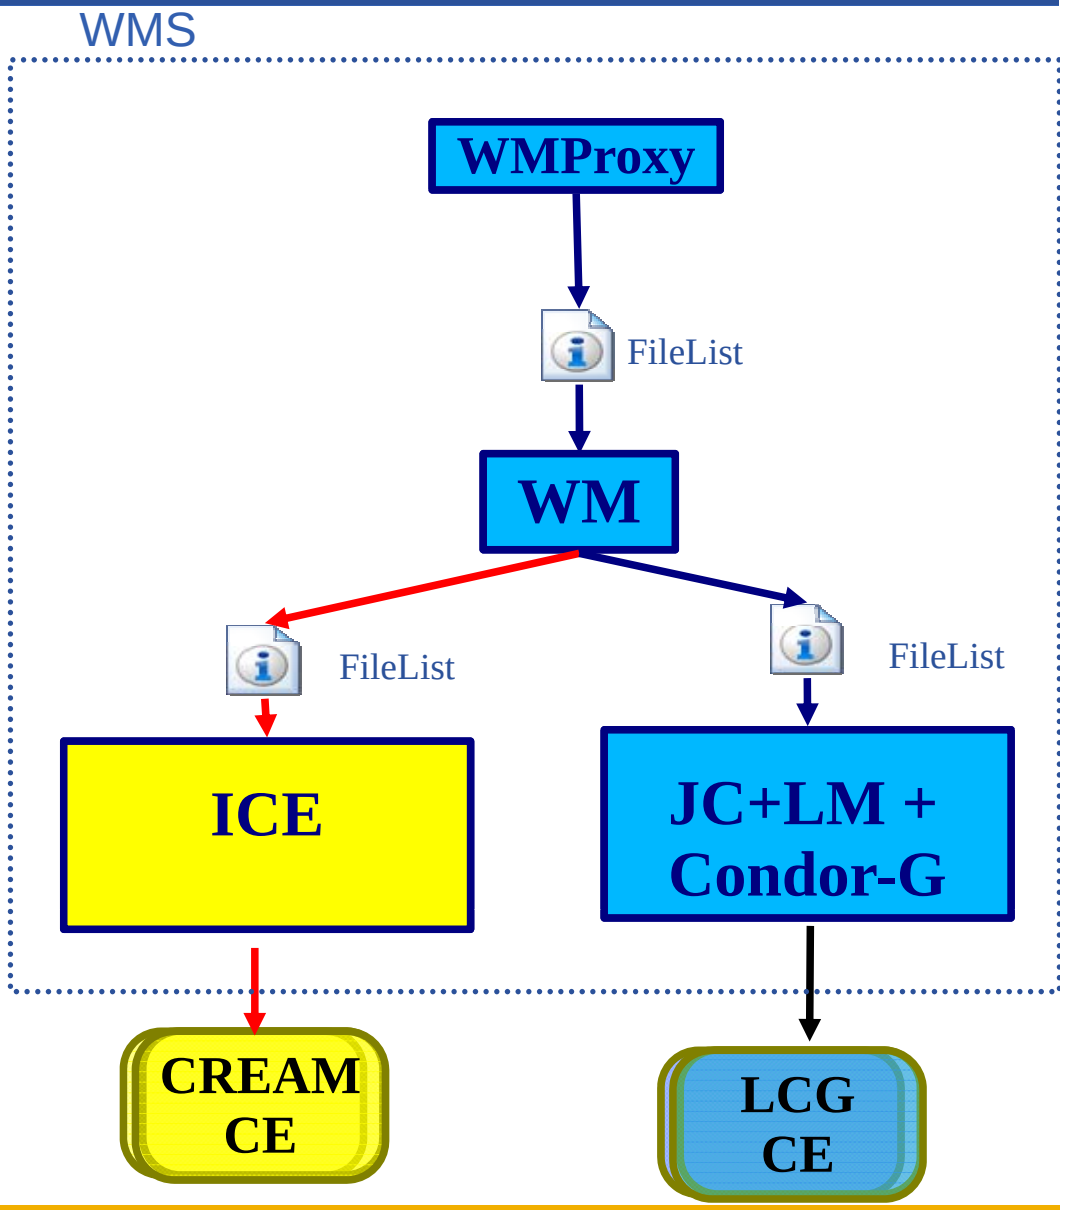

ICE: detecting CREAM job status changes

- 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, detect CREAM job status changes and take the appropriate actions
- As a fail-safe mechanism, ICE also able to poll CREAM if the relevant notifications are not received via CEMon

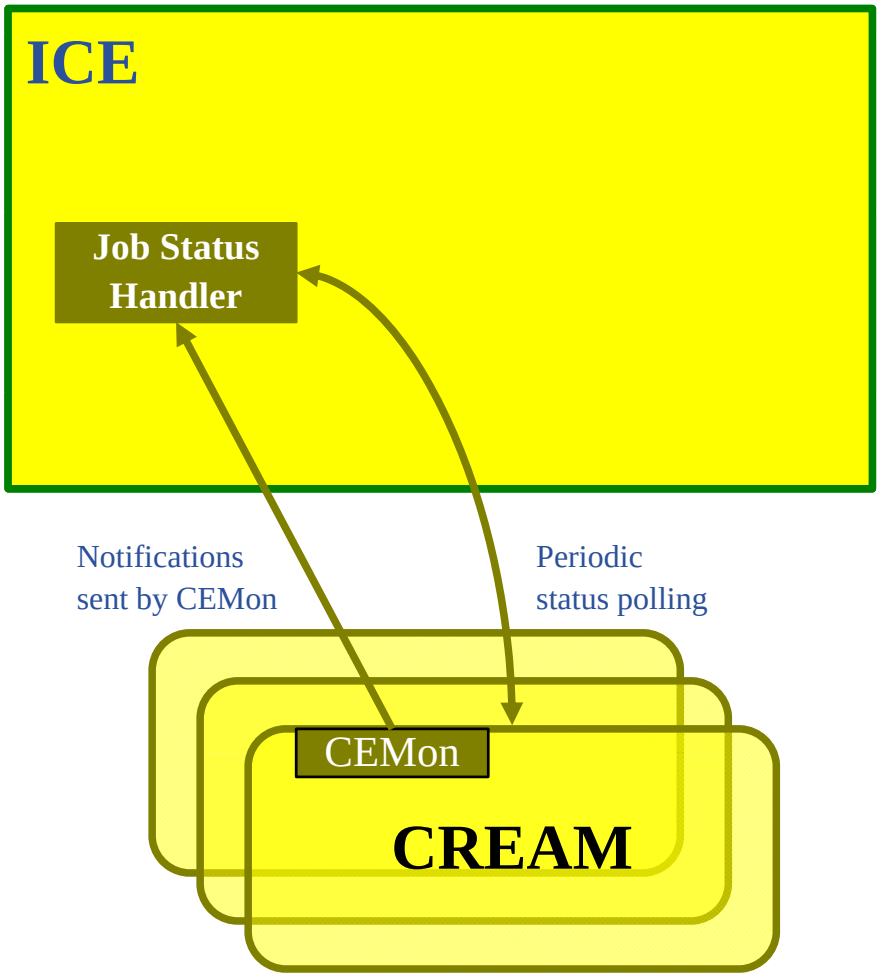

Interaction with L&B

- **No differences wrt LCG-CE**
  - Events logged by the job wrapper running on the Worker Node
    - Running, Really-Running, Done events (*Source* = *LRMS*)
    - Forwarded to LB server via the LB interlogger running on the CE node
  - Events logged by ICE
    - The same logged by JC+LM when dealing with submissions to LCG-CEs
    - Actually they appear with *JobController* / *LogMonitor* as *Source*
    - You'll never see ICE as *Source* in the LB events

Event: DeQueued
- Source = JobController
- Timestamp = Mon May 12 15:01:39 2008 CEST

Event: Transfer
- Destination = LRMS
- Result = OK
- Source = LogMonitor
- Timestamp = Mon May 12 15:01:55 2008 CEST

ICE Release

- **WMS ICE enabled released with patch #1841**
  - Under certification
  - Not only inclusion of ICE, but several bug fixes in other parts of the WMS
  - Acceptance tests passed
    - > 5 days of unattended running
    - About ~5k jobs always in the batch system queue
    - > 10K jobs submitted per day
    - < 0.5 % failures
  - But performance problems seen (in ICE) when several CREAM CEs are used
    - Being investigated

Installation & configuration

- **CREAM Installation & Configuration**
  - YAIM based configuration procedure available
  - Manual installation instructions available as well in the CREAM web site
  - CREAM configuration
    - Several CREAM attributes can be tuned editing the CREAM configuration file *opt/glite/etc/glite-ce-cream/cream-config.xml*
- **CREAM CLI installation & Configuration**
  - CREAM CLI installed on the UI
- **ICE installation & configuration part of the WMS installation process**
  - Yaim based conf procedure available
  - ICE Configuration
    - Specific ICE section in the WMS configuration file (*/opt/glite/etc/glite_wms.conf*)

```
[
WorkloadManagerProxy =  [
…
]
ICE = [
logfile  =  "${GLITE_LOCATION_LOG}/ice.log";
…
]
…
]
```

Some items for the future

- **Promptly address all issues that will be found**
- **Clearly new functionality/improvements will have to be negotiated with the TMB**
- **At any rate we have already identified some items that should be addressed**
  - Review proxy renewal mechanism used by BLAH
  - Migration to new BLParser model for LSF and PBS
  - Submission to CREAM by Condor
    - Some work already done with an “old” CREAM implementation
  - Submission of multiple jobs by the WMS to CREAM with a single call
    - E.g. if in the bulk matchmaking of a job collection n sub-jobs got matched to a certain CREAM CE, these n sub-jobs should be submitted all together to that CREAM
    - CREAM is basically already able to manage such scenario
    - Modifications are needed in the WMS
  - New development model for CREAM and WMS job wrapper
    - CREAM and WMS (the one used for LCG-CE) job wrappers have many common parts
    - Not good and dangerous to have duplicated code

Some items for the future

- **High availability/scalable CE**
  - CREAM CE front end and pool of CREAM machines doing the work
  - Main needed functionality already in place
    - Multiple CREAM CEs sharing the same backend (same DB)
    - E.g. a job can be submitted to a CREAM CE, and can then be cancelled on another CE
  - Still some issues to address
  - Preparing a proposal to discuss at the TMB
    - Which assumptions can be done ?
- **Better integration between CREAM and LB**
  - CREAM able to log information to LB
    - Enhance LB events with further information
  - Use of LB tools to monitor CREAM jobs
    - Also for the non WMS-jobs (i.e. the ones submitted directly to CREAM)
- **Adoption of new authorization service, when ready**

That's all !

- **More info in the CREAM web site**
  - http://grid.pd.infn.it/cream
- **Contact us:**

  jra1-pd@pd.infn.it
  - Paolo Andreetto (EGEE-III SA3, formerly OMII-EUROPE and EGEE JRA1)
    - Internal release manager, Etics confs, CREAM pre-certification
  - Sara Bertocco (EGEE-III SA3, formerly EGEE-II SA1)
    - yaim, CREAM pre-certification
  - Alvise Dorigo (EGEE-* JRA1)
    - CREAM and CEMon C clients, ICE
  - Eric Frizziero (e-NMR, formerly Cyclops)
    - CREAM and CEMon
  - Alessio Gianelle (EGEE-III SA3, formerly EGEE-II SA3 and EGEE JRA1)
    - CREAM pre-certification
  - Moreno Marzolla (EGEE-III JRA1, formerly OMII-EUROPE and EGEE JRA1)
    - CREAM and CEMon C clients, ICE, CREAM-BES
  - Massimo Sgaravatto (EGEE* JRA1)
    - Testing, overall coordination
  - Luigi Zangrando (EGEE* JRA1)
    - CREAM and CEMon