Job submission
Status, issues, plans

Massimo Sgaravatto

INFN Padova

Disclaimer

- Apologies if what I am going to say is already well known
- I am presenting my personal view on the job submission topic
  - I am not representing INFN-GRID and/or EMI position
- I am part of the gLite CREAM and WMS development team

What is the problem

- Submission of jobs (analysis, MC, reconstruction, …) to computational resources available in several sites
- Also management of these jobs (check status, cancellation, output retrieval, etc.)
- Resources in a site consist usually of a set of worker nodes (WNs), accessible via one or more front-end machines
  - Usually these resources are managed by a Local Resource Management System (LRMS)
    - LSF, Torque, PBS, SGE, Condor, …
- Some use to call Computing Element (CE) just the front-end machine, some call CE the front-end machine + WNs

Job submission components in LHC

- Different implementations of CEs
  - LCG-CE
  - OSG-CE
  - CREAM-CE
  - ARC-CE
- Some higher level tools used to submit jobs to CEs
  - gLite WMS
  - Condor-G
- Experiment specific frameworks used by end-user, which rely on some of these job submission components

LCG-CE

- Based on the Globus pre-WS GRAM
- Basically it is the evolution of the first CE used in the DataGrid project (~ 2002)
  - Enhanced by LCG to address some problems (in particular some scalability issues)
- Used in particular in Europe
- Issues with support of this CE
  - Not ported to gLite 3.2/sl5 and no plans to do it
  - Supposed to be phased-out and replaced by CREAM-CE
  - But still widely used
- Can be used by gLite WMS and by Condor-G
  - Submission to the LCG-CE through the gLite WMS is actually also implemented using Condor-G
- Possible but not straightforward to submit jobs directly to the CE without using an higher level tool

OSG-CE

- Also based on the Globus pre-WS GRAM
- Some differences wrt the LCG-CE
  - Different version of Globus
  - Some other different adaptations wrt standard Globus gatekeeper
  - Different accounting systems
  - Different information service on the CE
    - gLite CEMon in OSG-CE
    - resource BDII in LCG-CE
  - …
- .. but same interface
- Used in particular in US Open Science Grid
- Can be used by gLite WMS and by Condor-G
  - Submission to the OSG-CE through the gLite WMS is actually also implemented using Condor-G

CREAM-CE

- Part of the gLite software stack
  - Implemented and maintained in the EGEE projects and now in the EMI project
- Much more newer than the LCG-CE
  - First version released in Oct 2008
- Recent interest in deploying it also in OSG
- Web service interface
- Can be used directly by the end user
  - CREAM CLI available
- Can be used through the gLite WMS
  - In this case Condor is not used
- Submission to CREAM-CE via Condor being finalized
  - There were some issues supposed to be fixed in latest Condor release
  - This is the main issue which prevents the phasing out of the LCG-CE

ARC-CE

- Implemented and used in the NorduGrid
- Old implementation based on gridftp
  - Can be used by gLite WMS and by Condor-G
    - Submission to the ARC-CE through the gLite WMS is actually also implemented using Condor-G
- New web service based implementation (A-REX)
  - Not yet usable by Condor and therefore by gLite WMS
- Both implementations can be used directly by the end user
  - Using proper ARC clients

gLite WMS

- Part of the gLite software stack
  - Implemented and maintained first in the DataGrid project, then in the EGEE projects and now in the EMI project
- Higher level job management service
- Matchmaking functionality
  - Finds the most appropriate CE for a job taking into account
    - Job requirements and preferences
      - Specified by the user in the job description
      - Also possibility to express data requirements
    - Status of the Grid
      - Represented in the information service and in the file catalogue
- Integrated with Logging&Bookkeeping (LB) job tracking system
- Can be used to submit jobs to LCG-CE, OSG-CE, ARC-CE, CREAM-CE

Condor-G

- Leverages on the Condor job and workload management system, which exists since 1986
- Condor-G is just the job management part of Condor, used to submit jobs to the Grid
  - Simplify job submission and management
    - Consistent client interface for job submission and management
  - Make job submission more robust and scalable
- Able to submit to different CEs: LCG-CE, OSG-CE, CREAM-CE, ARC-CE, …
- Also used within the gLite WMS

Job Submission approaches

- Eager planning
  - User job is mapped to the physical resource upon job submission
  - Approach used by gLite WMS
- Very lazy planning
  - Mapping between postponed right up to the moment when the resource is immediately available
  - Approach used in pilot based systems
    - Pilots submitted everywhere
    - When a pilot starts running, it takes a “real” job from a central queue, if any, suitable for that resource

Eager vs. very lazy planning

- Eager planning
  - Maximum throughput if the whole system is always loaded
  - Otherwise there is the risk that some sites get overloaded while other resources are under-used
    - Also because information published in the information service (used to choose the CEs where to run jobs) can not be too accurate
      - E.g. because of misconfigurations
- Very lazy planning
  - Address problems of the eager planning approach
  - But not too suitable if you really have to pay for the resources you use
    - But this is usually not the case (yet)
  - “Ownership” of the jobs
    - Job pilot runs as user X, job to run belongs to user Y
    - Identity switch implemented by some software (e.g. glexec)

Job submission in LHC

- ALICE
  - Pilot based system
- ATLAS
  - Most of the work is based on pilots
- CMS
  - Both approaches used
    - Very lazy planning using Condor glidein
    - Eager planning using gLite WMS
- LHCB
  - Pilot based system

ALICE

- Jobs agents submitted to the resources of that site from the ALICE VOBOX deployed in that site
  - CREAM-CE: agents submitted using CREAM CLI
  - LCG-CE: agents submitted using gLite-WMS
    - Only at Cern, just because there are ~ 20 LCG-CEs vs. 3 CREAM-CEs at Cern
    - Without using any WMS matchmaking functionality
  - Nordugrid: agents submitted interacting directly with the batch system
    - Agreed with the relevant resource owners

ALICE (cont.ed)

- Agents submitted by an “automatic” machinery
  - Depends on number of tasks in the central queue, number of available resources, number of already queued agents
- When an agent starts running on the WN
  - It check if the environment is ok
  - It retrieves and runs a “real” job suitable for that site from a central queue, if any
  - It dies after a few minutes if there are no suitable jobs to run
  - An agent can run multiple “real” jobs
- Same procedure used for all kind of ALICE jobs
  - Analysis, reconstruction, simulation

ATLAS

- Each “organization” is responsible to submit pilots to its resources
  - For Italian T2s this is done via gLite WMS
    - For both LCG-CE and CREAM-CE
    - The target site is already used by the submitting user
      - The matchmaking functionality of WMS just used to select a CE within that site
  - For basically all the other sites, this is done using Condor-G
    - For LCG-CE and OSG-CE
    - Support for CREAM-CE is starting now
      - There were some problems in old Condor implementations
  - In some sites it look like pilots are submitted interacting directly to the batch system

ATLAS (cont.ed)

- User jobs (both production and analysis jobs) submitted through PanDA
- When a pilot starts running on the WN, it retrieves and run a “real” job from PanDA
  - Usually one “real” job per pilot
- NORDUGRID is managed in a different way
  - All NORDUGRID sites are seen as a single “virtual” CE
  - Jobs sent to this “virtual” CE are then managed by ARC middleware
- Some (very few) users use GANGA and submit directly real jobs instead of referring to the pilot model
  - Eager planning

CMS

- ProdAgent used for productions
- CRAB used for user analysis
- Both can use the gLite WMS or Condor glidein
  - gLite WMS used in particular for analysis (~ 75 %)
  - Condor glidein used in particular for productions (~ 75 %)
- gLite WMS
  - Used to submit “real” jobs
  - Eager planning: jobs are sent to the relevant CEs immediately
  - Matchmaking functionality of the WMS used basically to match sites
    - “close” to the relevant storage elements (found by the software experiment framework) where data are
    - where the needed CMS software environment is installed
  - Supports submission to CREAM-CE, LCG-CE, OSG-CE, ARC-CE

CMS (cont.ed)

- Condor Glidein
  - Used to submit Condor pilots
    - The pilot job that is started on the WN is the Condor startd daemon
    - When the Condor startd runs on the WN, that WN becomes part of a Condor pool
      - Two condor pools, one for analysis (master at UCSD) and one for productions (master at FNAL)
    - The real job is then submitted and managed in this WN (which is a Condor executing node) using Condor mechanisms
  - Supports submission to LCG-CE, OSG-CE, ARC-CE
    - Support for CREAM CE being implemented

LHCB

- Pilot jobs submitted to sites using WMS
  - To LCG-CEs, ARC-CEs, and CREAM-CEs
  - Not using OSG-CEs because no LHCB sites in the US
  - For analysis and reconstruction jobs, the CEs where to submit pilots is chosen by Dirac taking into account where data are
    - I.e. the WMS is told which CE must be used
  - For the other jobs, complex requirement and rank expressions given to the WMS
    - To take into account number of running and waiting jobs, free slots, Max cpu time, etc.
- Going to implement submissions of pilots to the CREAM CE using direct submission
  - I.e. without using the gLite WMS
- When a pilot starts running, it retrieves a “real” job from a central queue
  - According to priorities
  - A pilot can take more than one “real” jobs
    - Basically till the maximum allowed wallclocktime / cputime

Interoperability & Standardization

- Many different implementations of Computing Elements
- Interoperability is therefore an issue
- Several attempts of standardization, but without useful results yet
  - OGF BES and JSDL
    - Standards de-jure, but not really usable for real production activities
  - OGF PGI
    - Supposed to define standard “profiles”, really usable to do real work
    - Not too much progress yet
  - Common interface between CREAM-CE, ARC-CE and UNICORE-CE
    - Being done in the context of the EMI project
    - Specification to be finalized by Oct 2010, implementation by Oct 2011
- Interoperability is usually achieved now using clients able to submit to different flavors of CE
  - E.g. Condor able to submit to LCG-CE, ARC-CE, CREAM-CE, …

Some new developments

- For gLite (CREAM and WMS)
  - WMS with “feedback” mechanisms
  - Integration of cloud & virtualization solutions
  - Implementation of common EMI Execution Service interface in CREAM
  - Simplification and harmonization
    - E.g. use of a unique Authorization system (ARGUS) everywhere

WMS “with feedback”

- gLite WMS enhanced with “feedback” mechanisms
  - To address the problem of sites being over-used and sites being under-used
  - Also a problem for computations which can be split in independent jobs
    - They can run in parallel but they must be all successfully completed to consider the overall computation as completed
    - Problem if e.g. 99 % are completed immediately, and 1 % is completed e.g. after some days
  - When a job is stuck in the queue of a site for too long (because there aren’t idle resources) it is migrated to another site
  - The new site is chosen considering some statistics of the past behavior of the CEs
    - Provided by L&B service

Virtualization

- Problem
  - The VOs/experiments can require very different run time environments
    - Different OS, different installed software, etc.
  - Partition the whole set of resources is not very efficient wrt resource utilization
- Solution
  - Create the WN on demand, using an image suitable for that VO / application
  - Something already exists and already used in production in some sites
    - E.g. WNODES used with CREAM-CEs at INFN-CNAF
  - Need to enhance and make this solution more general
    - E.g. besides installing the OS and VO specific software, also installs and runs the VO specific pilot