Overview of operations at CC-IN2P3

Exploitation team

Reported by Philippe Olivero

11/30/2007

dapnia
cea
saclay

CENTRE NATIONAL
DE LA RECHERCHE
SCIENTIFIQUE

Summary

- BQS
- A few metrics
- Overview of local monitoring tools
- Concrete actions taken
- Issues
- Advices
- Questions

BQS

- **BQS in a few words…**
  - Home built batch system (15 years old) - >2 FT developers
  - Quick answer to problems and demands
  - Under continuous evolution, new functionalities are added for enabling
    - scalability, robustness, reliability (SGBD-R)
    - functionalities required by users
    - GRID compliance
  - Very rich scheduling policy including :
    - quotas, resources status, number of simultaneous spawned and running jobs…

BQS (2)

- **BQS philosophy :**
  - Dispatch of heterogeneous jobs on a worker node (workers not dedicated to one experiment)
  - Usage of BQS resources (kind of semaphores)

- **Current developments :**
  - Addition of GRID functionalities :
    - Managing VOMS groups and roles
    - Storing more GRID information into BQS
  - Support HEP-Unit [SI2K]

A few metrics – October 2007

- **Jobs :**
  - Total submitted jobs : 900 000 -> ~30 000/day
    (October 2006 : ~19 000 jobs/day)
  - 5000 running jobs (October 2006 : ~3500)
  - Total submitted jobs in the LONG class : 64%

- **2 separate farms :**
  - Pistoo : 64 cores (for parallel jobs) – 60 KSI2K
  - Anastasie : ~5000 cores – 7500 KSI2K
    - SL3 – 700 cores
    - SL4 32b – 3700 cores
    - SL4 64b – 600 cores
- **Farm usage :**
  - 62 unix groups, 3000 users, 500 regular users

A few metrics – CMS@October 2007

- **Jobs**

  Total submitted jobs : 126 000 (<15% of all)

  Total submitted jobs LONG class : 121 000 = 96 %

  CPU : 12.3% of all

- **Memory use for jobs on class Long**

  memory consistently requested : 2 GB (local Job-manager)

  72 % used less than 1 GB

  93 % used less than 1.5 GB

- **CPU time use for jobs on class Long**

  93 % of jobs used less than 40% of the class limit

  (but MEDIUM class has only 1GB of memory)

Local jobs monitoring

- **Tools to detect problematic jobs :**

1. **Stalled Jobs** : running jobs which do not consume cpu time.
2. **Jobs « early ended »** : bunch of jobs using much less cpu time than requested
3. **mails :** BQS sends a mail in case of job failure
4. **Manual checks** : by running different scripts

Concrete actions taken

- **Diagnose job failures e.g :**
  - Lack of resource, expired *proxy*, transfers pending, core LCG services unavailable
  - Job environment setting for a given VO
- **Find the Grid Job Identity**

  LCG job IDs, BQS job IDs, globus job IDs
- **Inform the users or the VO admin**
- **Notify the administrator of services involved in the problem**

  mail, GGUS ticket
- **Various tasks for managing the production :**

  Jobs can be deleted, locked in queued, rerun etc.

Concrete actions taken

- **Increasing VO's quota**

  In order to face selective intensive computing : DC, MC production ..

- **Create & modify BQS resources semaphore**

  To cope with internal services unavailability (HPSS, dCache)

- **VO-specific job prioritization**

  To regulate automatically job priorities and resources according to the VO requirements

Issues

- Real needs of resources (CPU, memory) are unknown for GRID jobs
- CMS soft stresses AFS -> number of running jobs is limited
- Users are not well informed about the LCG service status (downtimes)
- Recurrent problems with files access or copy : remote SRM SE unavailable, LFC not responding, transfers failing…
- Sometimes it's hard to find the user email
- Sometimes very low reactivity from users
- Hard to trace jobs which are not submitted through Resource Brokers

Issues

- Zombies processes left by ended jobs on the workers nodes
  - solved in the current BQS version
- Lack of tools which may allow us to manage priorities inside the VO (T1/T2, productions/analysis)
  - solved in a future BQS version (end of 2007)
- Memory wasting with jobs submitted on the class Long
  - Glite ? Cream ?

Advices

To have more running jobs:

- **Have always queued jobs to reach a good score of running jobs**
- **Limit memory request for jobs submitted on the LONG class when it's possible** (direct submissions)
- **Keep us informed as soon as possible about critical production periods**
- **Migration to SL4 64b**

Comments / Questions