

Computing grids

Bruno Bzeznik

Autrans - Autumn school ANGD/CNRS

05 october 2009

Introduction

Bruno Bzeznik

- CIMENT : Grid engineer, systems expert
- LIG (OAR team) : OAR and CIGRI developer

Grid computing

- 1 Definitions
 - The grid concept
 - From the process to the grid
 - Grid definitions
- 2 Grid classification and examples
 - by objective
 - by infrastructure
 - middlewares
 - around the grids
- 3 CiGri
 - The CIMENT computing center
 - The CIMENT grid
 - The CiGri middleware

History

- Grid... a fashion ? No more, now the fashion is "cloud" computing :-)
- The term "grid computing" was introduced by Ian Foster in the early 1990
- Very popular in late 1990 with Seti@home and Napster
- In France : ACI grid started in 2001

The grid concept



- Comes from the "power grid" concept
- In a power grid, there are several energy sources and the ending user consumes a part of that energy without knowing exactly where it has been produced.
- In a computing grid, there are several computing hosts and the ending user launches tasks that will run on some of them without knowing exactly where.

The grid concept



- Well...
- Computing tasks are a bit more complicated than a simple electrical flow :-)
 - Application code dependency
 - Input data dependency
 - I/O data amount
 - Duration
 - Type of code : parallel/sequential
 - ...

From the process to the grid

Process


Processes are running on CPUs.



From the process to the grid

Jobs


Processes can be grouped into jobs.



From the process to the grid

Nodes


Jobs are running on nodes. Nodes are computers (one or several cpus, a shared memory space, and i/o device).



From the process to the grid

Computing network


Several nodes are connected to a computing network, generally low latency network (Myrinet, Infiniband, Numalink,...)



From the process to the grid

Parallel jobs


Jobs may be "parallel" or "sequential". A parallel job runs on several nodes, using the computing network to communicate.



From the process to the grid

Clusters


A batch scheduler is managing jobs and nodes. We have a cluster.



From the process to the grid

Job submission


Users submit jobs to the batch scheduler.



From the process to the grid

Public network


A cluster frontend maybe connected to a public network, generally not the same network as the private computing network.



From the process to the grid

Public network


Clusters frontend maybe interconnected.



From the process to the grid

Grid middleware


A grid middleware communicates with the different batch schedulers.



From the process to the grid

Grid middleware


The middleware may be responsible of the communication with other external elements to the grid : sensors



From the process to the grid

Grid middleware

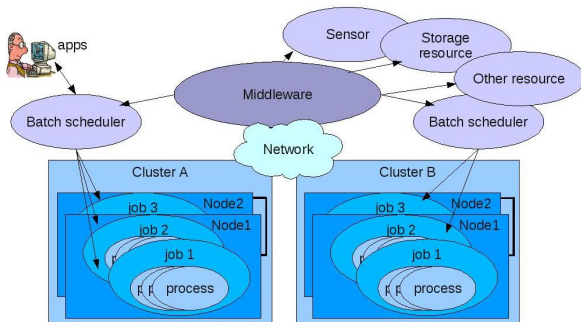
The middleware may be responsible of the communication with other external elements to the grid : sensors, storage



From the process to the grid

Grid middleware

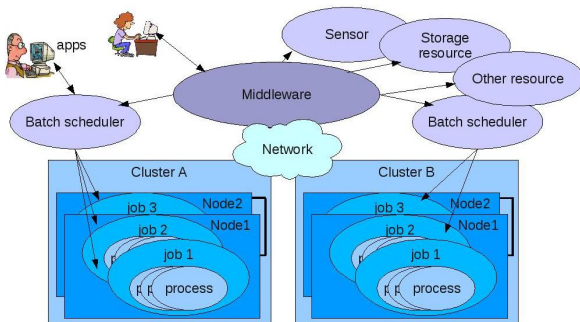
The middleware may be responsible of the communication with other external elements to the grid : sensors, storage, etc



From the process to the grid

Grid job submission


The users interact with the grid through the middleware, for submitting grid jobs for example.



From the process to the grid

Desktop/volunteer computing


But a grid may also look like this...



From the process to the grid

Peer-to-peer grid


...or like this...



From the process to the grid


Naradabrokering

...or like this...



Grid definitions

Common mix-up

- *Grid* \neq *Cluster*
- (french) *Grille* \neq *Grappe*
- *Grid* \neq Cluster of clusters (a grid can't be constructed by simply nesting batch schedulers) 

Grid definitions


In french...

- Grid computing = calcul informatisé sur grille
- GRID = Globalisation des Ressources Informatiques et des Données (ACI GRID : <http://www.recherche.gouv.fr/recherche/aci/grid.htm>). Le terme semble souvent repris dans la presse pour traduire le mot "GRID".

Grid definitions

Trivialisation sites

- The GRID on tryscience : <http://www.tryscience.org/grid>
- Gridcafe (CERN) : <http://www.gridcafe.org>
- Decryphon video : http://www.decryphon.fr/english/ewb_pages/g/grid-computing.php

 The **grid** definition differs depending on the point of view, especially on the way that computing nodes are interconnected (broadband internet, high speed dedicated networks, directly or not,...)

Grid definitions

Wikipedia

"Grid computing (or the use of computational grids) is the combination of computer resources from **multiple administrative domains** applied to a **common task**, usually to a scientific, technical or business problem that requires a **great number of computer processing cycles** or the need to process **large amounts of data**."

Grid definitions

The Grid, I. Foster, C. Kesselman, 1998

"A computational grid is a **hardware and software** infrastructure that provides dependable, consistent, pervasive, and **inexpensive** access to high computational capabilities."

Grid definitions

Ian Foster, again, in 2002

"a Grid is a system that :

- coordinates resources that are **not subject to centralized control** [...]
- using **standard, open**, general-purpose protocols and interfaces[...]
- to deliver **nontrivial qualities of service** [...]"

Grid definitions

The CERN dream : **The grid**

"[...] Now imagine that all of these computers can be connected to form **a single, huge and super-powerful computer** ! This huge, sprawling, global computer is what many people dream "The Grid" will be."

Grid definitions

Michel Louvin, JTR2006

"Partage transparent de l'utilisation de ressources massivement distribuées par des utilisateurs de différentes disciplines"

Grid definitions

01net

"Exploiter la puissance ou les ressources de stockage **non utilisées des PC en réseau**. Tel est le principe sur lequel reposent les grilles de calcul (ou Grid Computing)"

So... When ?

Regarding this panel of point of views,

- when you need more resources than what you can have in one unique place (because of power, conditioned air, area, administrative reasons,...)
- when you want to optimize computers or supercomputers that are not used all the time
- when you have an application that has several paralelism levels and that we easily imagine to naturally use several supercomputers (**code coupling**)
- when it is the cheapest solution for the same service

a grid may be anything you can imagine !

What for? (usage examples)

- Physics : Analyse de donnees d'une experience dans un accelerateur de particules (LHC). Volume de donnees "enorme" de 15 Petabytes par an
- Health : Base de donnee et analyse de millions de mammogrammes repartis sur une centaine de sites.
- Industry : Optimisation de chaines de production avec un algorithme genetique : creation de populations de solutions et croisements sur la grille -j milliards de valeurs de plusieurs parametres
- Environment : Modelisation hydro-meteorologique en Afrique de l'ouest, couplage sol-vegetation-atmosphere
- IT consequence : optimiser l'utilisation des calculateurs d'un projet multicommunautaire en repartissant au niveau grille des calculs multi-parametriques independants (CiGri)
- ...


Challenges

- Standards (OGSA, SAGA, DRMAA, GLUE, GRIDRPC,...)
 - ▮▶ Open Grid Forum (<http://www.ogf.org>)
- Heterogeneous resources
- Data management (amount, synchronization, distribution,...)
- Security and privacy (authentication, encryption,...)
- Availability (monitoring, redundancy,...)
- Sharing (priority, accounting, fairsharing,...)
- Networking performance
- Organization : who is the administrator of the grid ?
- Applications : how to gridify an application ? or how to make the grid transparent to the applications ?

classification by objectives

Information grid

To share knowledge



Examples :

- World Wide Web
- Virtual observatory <http://www.france-ov.org>

classification by objectives

Data storage grid

High scale data storage



Examples :

- LCG (15 Petabytes/an, over EGEE)
<http://lcg.web.cern.ch/LCG/>
- eMule (eDonkey)
- Bittorrent (the grid and it's middleware are merged...)

classification by objectives

Computing grid

Computing power aggregation

Examples :

- EGEE <http://www.eu-egee.org>
- DEISA <https://www.deisa.org>
- CIMENT Grid <https://ciment.ujf-grenoble.fr>
- SETI@home, Folding@home, Decryphon

classification by objectives

Experimentation grids

Distributed computing research

Examples :

- Grid5000 <http://www.grid5000.fr>
- PlanetLab <http://www.planet-lab.org>
- DAS3 <http://www.cs.vu.nl/das3>
- XtremLab <http://xtremlab.lri.fr>
- NAREGI http://www.naregi.org/index_e.html

classification by infrastructure

Institutional grid

Generally grids of clusters, with stable and secured nodes.

Examples :

- EGEE <http://www.eu-egee.org>
- Grid5000 <http://www.grid5000.fr>
- CIMENT Grid <https://ciment.ujf-grenoble.fr>

classification by infrastructure

Desktop and volunteer computing

A lot of nodes (millions), volatiles and not secured

Examples :

- XtremLab <http://xtremlab.lri.fr>
- SETI@home, Folding@home, Decryphon
- eMule, bittorrent
- Computemode

Middleware examples

- The GLOBUS Toolkit <http://www.globus.org> (EGEE, National Virtual Observatory,...)
- gLite : Globus based (EGEE)
- UNICORE (DEISA)
- Oargrid, kadeploy and... ssh (Grid5000)
- CiGri (CIMENT)
- Boinc (*@home)
- CONDOR-G : globus based
- ARC : Globus based (Nordunet)
- eMule
- XtremOS

Cloud computing

- A more recent term to design something less specific than grids
- The idea is that you can use an application or manage data through services without knowing where they are (somewhere in the cloud)
- It's related to an economical model where clients pay for services without worrying about the infrastructure
- Often related to virtualization (you may rent an OS running somewhere in the cloud)
- Also related to scalability : the infrastructure adapts to exactly what you need
- Critics : you lose the control of your own data (Stallmann) ; it's a sexy word to sell something that exists since a long time (Oracle)

Cloud computing

- Famous examples :
 - Amazon EC2 (virtual hosts) and S3 (online storage web service)
 - GoGrid
 - iCloud (free!)
 - Google apps
 - eyeOs

Pervasive computing

- = "Ubiquitous computing", "Informatique ambiante", "Everyware"
- Network devices everywhere at all scale : in your phone, in your clothes, in your skin...
- Involves wireless networks and embedded systems
- Linked to grids due to the distributed aspect and numerous computing resources it may involve

Green grid computing

- Turning off unused nodes
- Best-effort / volunteer computing at low frequency
- ...

Best-effort grid computing

Using the free cpu cycles for parametric computation when the local users don't need them. Zero priority jobs are immediately killed when a local user needs the resources.

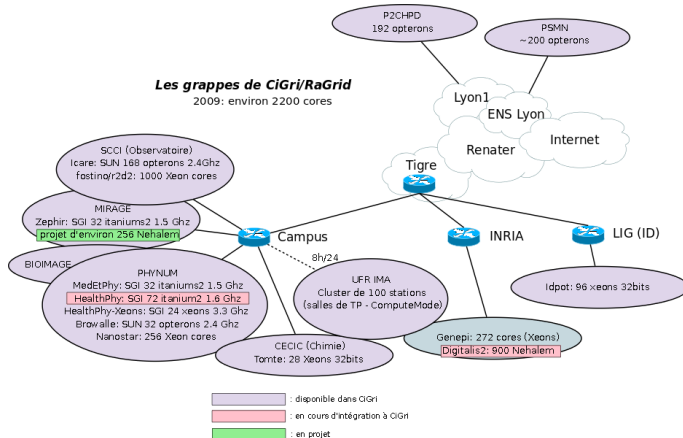
The ciment project

- <http://ciment.imag.fr>
- Calcul Intensif, Modelisation, Experimentation Numerique et Technologique
- Born in 1998 in the Grenoble University (FRANCE)
- Experts and engineers group for HPC usages into the modeling and distributed computing community of the Grenoble's University.
- Distributed computing meso-centre : 5 departments for several subjects. N clusters not especially dedicated to the departments

The ciment "poles"

- SCCI : Sciences de l'Univers, Observatoire
- MIRAGE : Modelisation, Environnement, Climat
- Grappes de PC : Informatique Distribuee
- CECIC : Chimie
- BioIMAGe : Biologie, Imagerie medicale, Sante
- PhyNum : Physique Numerique


The ciment "platforms"



The ciment grid

- Optimizing the use of the CIMENT clusters (which are heterogeneous)
- The lightweight grid concept :
 - Minimizing the grid computing problem : only multiparametric applications
 - Sites are using common administration rules
 - No big authentication problems (ssh and sudo are good enough for us)
 - Lighter than heavy grids like GLOBUS
 - No application deployment, no evolved data management
- Campaigns may be composed of a very high number of small jobs
- Cigri is also used to experiment distributed scheduling in a computer science research context
- Locally idle resources are used by the grid (best-effort concept)

The CiGrid middleware



- A central CIGRI host
- Uses an SQL database as the core model
- Communicates with clusters via ssh
- Non intrusive for local production sites
- Submits jobs into the OAR batch scheduler (maybe coupled with another bs)
- Uses the "best effort" concept of the OAR batch scheduler (next slide)

Best-effort jobs

- Best effort jobs allow us to exploit idle resources of production clusters
- When a node is free, a best-effort job may be placed on it.
- If a "normal" job needs the node, the best-effort job is killed
- as a consequence, a best-effort job must be "short enough" (lt 30 minutes) or there's too much chance for it to be killed
- Checkpointing may help

Cigr efficiency (1/2)



The load of the clusters is not constant and peaks are often not at the same time...

Cigr efficiency (2/2)



CiGrid uses the idle cpus

Thank you

Bruno.Beznik@imag.fr

CIMENT engineer - UJF Grenoble

<http://ciment.ujf-grenoble.fr>

CiGri tutorial for CIMENT users :

https://ciment.ujf-grenoble.fr/wiki/index.php/CiGri_tutorial

Why a grid is not a cluster of clusters?

- A cluster scheduler at grid level has got the vision of all the resources (cpus) \implies not a grid but an heterogeneous cluster
- A special scheduler having an aggregated view of clusters resources and a system to send jobs to underlying batch schedulers \implies a grid middleware and not a simple batch scheduler so, not a cluster; a grid submission results in a job submission, not in an execution on a node.

◀ Back.