The HEP Communities: Grid Infrastructure/Services and Requirements

User Community Transition from EGEE to EGI

Patricia Méndez Lorenzo (CERN, IT/GS-EIS)

Orsay, 17th December 2008

www.eu-egee.org

Information Society and Media

Agenda

- LHC preparing to record pp collisions events in 2009 (5+5 TeV)
- **15PB of data to be collected per year, analyzed by thousands of physicists placed all over the world**
- Done using a worldwide (virtual) grid – federating resources from several large-scale grid infrastructures – to offer a high (performance, quality) production service
- **The perspectives of the 4 LHC experiments, the large amount of data and users to handle and their world-wide distribution defines the level of SERVICES, OPERATIONS and SUPPORT required by this community**

Outlook

- The WLCG/EGEE project
- The Support infrastructure for HEP
- The NA4-HEP cluster in EGEE
- The Services infrastructure for HEP
- The Operations infrastructure for HEP
- Summary

The WLCG project

- **Worldwide LHC Computing Grid (WLCG) is the Grid project created to cover the computing expectations of the 4 LHC experiments**
- **Working in close collaboration with EGEE in terms of services setup, middleware creation, support, operations, etc.**
- **This formally consists of a collaboration between the 4 experiments and a set of institutes that provide computing resources and services to these communities**
  - Defined in a "Memorandum of Understanding" signed by all parties;
  - Includes services and service levels offered, resource pledges for coming years
- **EGEE also includes other communities in the picture – other (VOs, sites, services) which are supported at CERN** igb1
  - e.g. GEANT4 – main simulation tool in HEP and (way) beyond…
  - SIXT – simulation tool for LHC accelerator (sixtrack)
  - Lattice QCD simulation – a number of serious scientific publications
- **Many physicists / institutes also involved in other experiments and / or disciplines**
- **Flexible environment**
  - e.g. CERN Grid Support group works with a wide range of disciplines;
  - Grid Data & Storage Management products (dCache, DPM, …) – even if developed at HEP labs for HEP users are also used by – and extended for – many disciplines…

Slide 4

igb1 Strictly speaking, WLCG is explicitly limited to the 4 LHC experiments. These VO's are supported by EGEE.
Ian Bird, 12/16/2008

The Experiment Support

- **There are two fundamental concepts to distinguish in terms of support:**
  - User support
    - Managed by GGUS as a central «helpdesk» - tickets managed by TPMs
      - *Responsible to categorize, dispatch and follow up*
      - *TMPs are ROC responsibilities*
    - Not only in terms of user support, also for network, services and operation support
    - Fundamental tool for users.
      - *It acts as a unique access point*
      - *Enabling a managed and tracked process in order to have a reliable operatins procedure*
  - Application support
    - Concept I will refer always to during the following slides

The Experiment Support

- **HEP support infrastructure needs to deal with two middleware stacks**
  - Supported middleware
    - Common to all applications – computing model agnostic
  - Application middleware (HEP Computing models)
    - Requires a deep knowledge of the computing models and their Grid use
- **The support team needs to become expert in both middleware stacks**
  - Each support member needs to know the computing model of the supported application
  - This requires a big effort by the Grid support team and labor intensive interaction with the communities

igb2

igb2 effort

Ian Bird, 12/16/2008

The Experiment Support

- **Fundamental tasks of the HEP Grid support team**
  - Liaison between the community and the services-operations and middleware teams
    - The application might not know many Grid related issues
  - Provides Grid expertise for the middleware implementation into the applicatiom computing model
    - Software development towards final implementation
    - Creation of missing components which might be shared by other communities
    - Evaluation of the middleware functionality and scalability tests based on the community needs
  - Management of VO-specific services
  - User-support of needed (a la GGUS)
  - Expertise to solve Grid related issues for running applications
    - To the site and middelware level
    - Deep knowledge of bothGrid infrastructure and VO computing model
  - Documentation and training (application oriented)
  - Liaison with the sites
  - Operations activities if needed
    - Moreover in terms of services and resources negotiation

NA4-HEP cluster

- The Grid team at CERN helps different communities to be merged in the Grid environment in and beyond HEP
- Our "Gridification" structure is based on 3 major blocks igb3

RESOURCES	SUPPORT	TOOLS
→ Basic and full Grid infrastructure is guaranteed to the users → Built based on the experience gained with previous communities	→ Follow up of the gridification and productions with the communities → Dedicated support person for each HEP community → Training and dissemination	→ Production • Ganga and/or DIANE are proposed as tools → Data management: • AMGA, xrootd → Monitoring • Dashboard → Analysis facilities

Slide 8

igb3 major

Ian Bird, 12/16/2008

Ganga Communities

ATLAS

LHCb

HARP

COMPASS

Cambridge Ontology

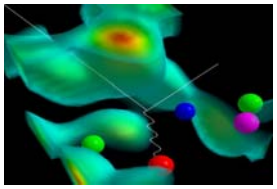

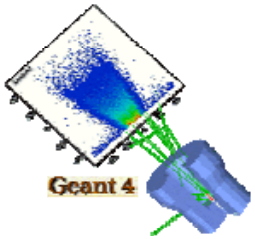

Geant 4

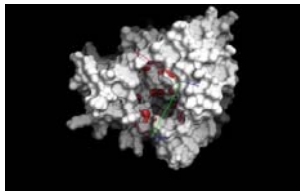

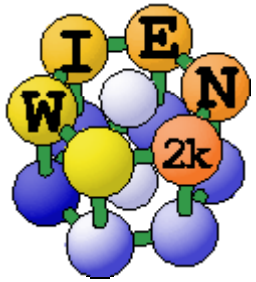

Academia Sinica
Genomics Research Center

LHC Collimation Project
CERN

unitar
United Nations Institute for Training and Research

UNOSAT
satellite solutions for all

CERN PS Multi Turn Extraction

med austron

ITU International Telecommunication Union

Garfield

HEP/Fusion collaboration

- During EGEE08 Fusion and HEP clusters shared a common session to establish a future collaboration based on the tools created and supported by HEP
  - 1st step: Ganga/Diane set up for the Fusion team
  - The results will be shown during the UF4 (Catania)
- Discussions continued in Geneva during the Fusion Conference
- The data management experts (for AMGA) and the monitoring team (for Dashboard) have also offered their help to contribute to the future collaboration

WLCG Services: General View

- **Level of the Grid services provided by WLCG**
  - Signed by all sites in the MoU
    - Contains the services provided by each site per experiment, the time schedule for recovery and the level of service
  - Targets proposed for the T0 services (Critical services)
    - Tier1s: 95% of problems resolved < 1 working day
    - Tier2s: 90% of problems resolved < 1 working day
    - **Post-mortem triggered when targets not met!**

Time Interval	Issue (Tier0 Services)	Target
End 2008	Consistent use of all WLCG Service Standards	100%
30'	Operator response to alarm / call to x5011 / alarm e-mail	99%
1 hour	Operator response to alarm / call to x5011 / alarm e-mail	100%
4 hours	Expert intervention in response to above	95%
8 hours	Problem resolved	90%
24 hours	Problem resolved	99%

- **We must consider the specifications of the experiments: Critical services**

Experiment	Down	Seriously Degraded	Perturbed
ALICE	2 hours	8 hours	12 hours
ATLAS	As text	As text	As text
CMS	30'	8 hours	24 hours (72)
LHCb	30'	8 hours	24 hours (72)

The Current Service Status

- **Has the current service stood up to real production usage?**
  - One year ago we had not yet demonstrated a sustained production workflow for the 4 LHC experiments
  - This lead to the « Common Computing Readiness Challenge(s) (CCRC'08) »
    - Combined effort between the 4 LHC experiments plus the sites serving the experiments
      - *The 4 experiments together testing all Grid services*
    - The goal has been the measurement of the readiness of the Grid services and operations before the real data taking
      - *Identify problems early and allow time for fixing*
    - Complementary to the experiment Full Dress Rehearsals
      - *Experiment exercises testing the full data chain*
    - Two slots foreseen: Feb and May 2008
      - *Goals have been met but real data taking will be the real challenge*
    - Future CCRC foreseen during the LHC shutdowns

Experiment example: ALICE

❖ Results during the Feb phase of CCRC'08

❑ Running simultaneously the 2nd Commissioned Exercise

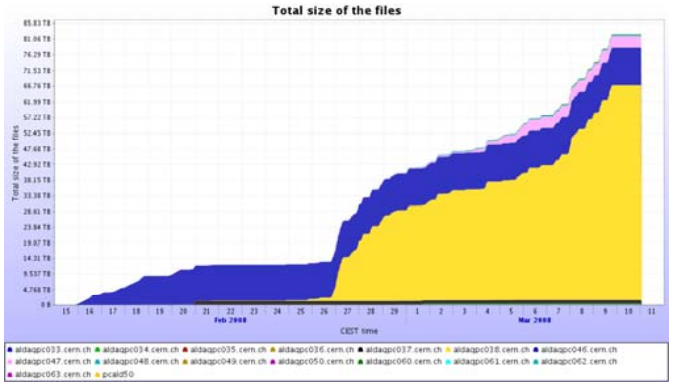

Total size of files copied from pit to Castor2@CERN: 70% of the nominal data rate achieved

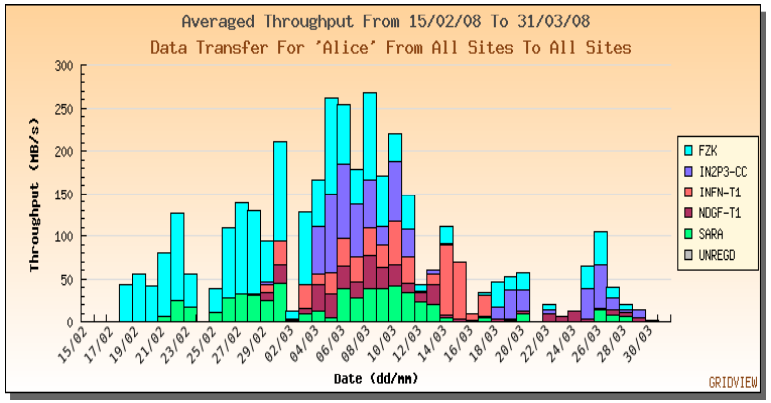

Gridview: daily averaged throughput to all ALICE T1 sites

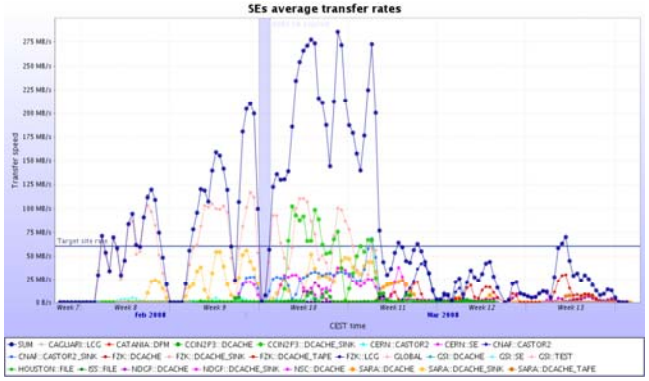

MonaLisa: SE transfer rates

Expected: 60MB/s

Achieved: 125 MB/s

CCRC'08 post-mortem

- **Success measurements**
  - Explicit "**scaling factors**" set by the experiments for each functional block: discussed in detail together with sites to ensure that the necessary resources and configuration were in place;
  - Targets for the lists of "**critical services**" defined by the experiments – those essential for their production, with an analysis of the impact of service degradation or interruption
  - WLCG "**Memorandum of Understanding**" (MoU) targets – services to be provided by sites, target availability, time to intervene / resolve problems …
- **The experience in 2008 confirms that we have a working service model and we are ready to face the data taking challenge**
  - ☺ Most aspects f the service working well and issues resolved in an acceptable short time
  - ☺ Middelware process and DB services have properly worked
  - ☹ Still to ensure robustness in those services associated to SE and data management
  - ☹ Communication still an issue, this will required extra effort
  - ☹ Not all experiment activities fully finished in May

WLCG General Operations

- **The experience gained at CERN for many years shows the value of running daily operations meetings**
  - No longer than 20min (week-days at 15:00)
  - The current alarms and issues are checked and discussed
  - Assignment of open issues to the responsible
  - If any weakness in the system, this will be exposed
  - Experiments are also welcome to join!
- **Weekly operations meeting with the ROC responsible**
  - Escalation of open problems mostly opened by the experiments
  - Each experiment is represented during the meeting
- **Besides this procedure any operation in the system must be performed as quick and transparent as possible**
  - Otherwise the degradation of the system can cause serious problems to the users and other services
- **Also once the intervention is finished a post-mortem and documentation procedure is required for future interventions**
- **Experiments also have their own operations teams and infrastructures**
  - WLCG collaboration workshops: 200-300 people
  - ATLAS “jamborees”: closer to 100….

WLCG VO Operations

- Experiment Dashboards, VO-specific SAM tests together with specific monitoring systems
- Very close collaboration between Grid support team, the experiments and the sites
- Very close collaboration with the Grid and Services experts at the T0 and the rest of the sites
- This schema has demonstrated to work
  - Experts in both sites (computing model and Grid infrastructure knowledge) are needed
  - This procedure ensures an efficient and effective support and experiment production

Summary

- **At the end of EGEE-III, LHC will be in full data taking regime**
  - HEP is a clear example of an international application with a central institution
- **The Grid infrastructure created for this community is the result of many years of work in many fields**
  - Closed specific and dedicated experiment support
  - Services stress tests (Data challenges, service challenges, FDR, CCRC'08, etc)
  - Specific Operation protocols
- **It is extremelly dangerous to change the current infrastructure known by the experiments**

At no time can there be an interruption to the WLCG service