



# REBUS – input from experiments, sites and project office

pre-GDB Accounting Review  
Apr 12, 2016

Maarten Litmaath  
v2.1

# Current functionality

- **Display of topology**
  - Only MoU sites
  - Input by WLCG project office
- **Input and display of pledges**
  - Input by site/federation representatives
- **Display of installed capacities**
  - Automatically collected from the BDII
  - Often wrong or absent for many sites
- **Input / correction of accounting data**
  - Currently only by T0/T1 representatives
- **Generation of accounting and trend reports**

# Installed capacities

- The concept has been discussed to a great extent in the Information System TF
  - <https://twiki.cern.ch/twiki/bin/view/EGEE/WLCGISEvolution#Definitions>
- Preliminary conclusions
  - **Installed** Capacity: physical HW which is in place
    - Not very interesting in comparison to the next item
  - **Available** Capacity: HW which is actually usable (i.e. not offline for maintenance) for a period of time longer than a few days.
- Whenever a significant amount of resources are added or removed, the site should ensure the relevant available capacity gets updated
  - Automatically where possible and desired, else manually
  - Could be as little as once or a few times per year

# Requirements and wishes

## *executive summary*

- Pledges
  - Input and display as before
- T2 data also per site, not only federation
- Historic views for everything
- Aggregation views also by country
- Wish: allow inclusion of non-WLCG and opportunistic resources
- Available capacities
  - Need to be trustable → directly inserted by sites
  - Automatic where possible and desired, else manual
  - Display them, also in comparison with pledges and accounting
  - The precision for computing may be 10%
- Notification on information changes
- Accounting and trend reports as before

# Implementation considerations

- Significant overlap with CRIC →
  - Use that service to get the following:
    - Topology
    - Capacities
  - Values are inserted through admin API and/or GUI
  - Repository keeps track of history
- Accounting data from the improved EGI accounting portal
- REBUS itself handles pledges and report generation
  - Current and historical views
  - T1, T2, ...
  - Countries, federations
  - Possibly experiment-specific callouts for inclusion of non-pledged resources

- The following slides give details provided by the various stakeholders
  - Experiments
  - Sites
  - Project office
- The experiment and site inputs were essentially taken from the April 7 Ops Coordination meeting
  - <http://indico.cern.ch/event/514077/>
  - <https://twiki.cern.ch/twiki/bin/view/LCG/WLCGOpsMinutes160407>

# REBUS for ALICE

- REBUS is used for pledged and required resources accounting
- A historical view within one year would be very useful for accounting purposes
- A view of non-WLCG resources should be possible (opportunistic or sites without WLCG affiliation)
- Currently installed capacity needs to be recorded and can go in the historical view within one year
- T2s should be entered in REBUS per site, not as a federation

# REBUS for ATLAS

- ATLAS use REBUS as the place where the pledge are collected from the sites. we need it (or similar).
- ATLAS use installed/available capacity to have a guess of the core power.
- Would be useful to know how much a site can provide to each VO, e.g. a working installed capacity by site.
  - This would most probably need to be extended because as of today this installed capacity per federation is used to match with pledges, while ATLAS would be interested to know how much it is really available, including over the pledge and opportunistic resources.
- Additional note: for accounting numbers for computing a +/- 10% can be considered ok.



# REBUS for CMS

- REBUS is use for pledges
- There is a general agreement in CMS that capacity information is useful if it is correct and can be trusted
  - Useful also if sites could rely on the storage systems themselves instead of the sys admin to report about storage capacity in a reliable way
- A missing functionality is to know when information changes in REBUS
  - Although this functionality could be moved to the new IS if pledges are also included there

# Rebus for LHCb

- Storage
  - Using information of pledged resources
    - In addition to pledges information about single sites also aggregation by country
  - Less interested in installed and capacity available to the VO
    - Better be done “live” - see space accounting
- CPU
  - Rough comparison between pledged resources and executed work is done but not to the % level
  - In operations need to take into account different benchmarking and fare share mechanisms of sites.
- Request to provide aggregated pledges across federations in addition to individual sites [ *vice versa!* ]

# REBUS for sites

- Reports pledges correctly as they are inserted manually.
- Reported capacity depends very much on what sites publish in the BDII and it is not always correct nor validated by sites.
- There is a discussion about replacing "installed capacity" with "available capacity" which is considered more accurate and get the sys admins to manually validate it every month.
  - This maybe ok for T1s, for T2s needs to be thought out IMO.
  - Also it is not clear if this requirement is for experiments ops teams which may want this information to plan production campaigns and need the information in advance or if it is an accounting requirement and in that case it is similar to validate afterwards the numbers of the reports circulated by the WLCG office which are generated from REBUS.

# REBUS for project office

- Not a great fan of the look/layout, but the raw data and functionality is there, and it is easy to use.
- Feedback from a T1 that a .csv upload function for accounting be made available (instead of manually editing).
- EGI Accounting portal used for REBUS T2 reports, but historic data is not available for decommissioned sites.
- We should review what stakeholders want from the formal reports, and adapt the Portal accordingly.