

Minutes of Information System Task Force, 23rd July 2015

Local: Maria Alandes (chair, minutes), Julia Andreeva, Alexey Anisenkov, Laurence Field, Edward Karavakis, Maarten Litmaath, Andrea Sciaba, Andrea Valassi.

Remote: Brian Bockelman, Stephen Burke, Alessandra Forti, Andrew McNab, David Meredith, Rob Quick, Samuel Skipsey, Oxana Smirnova, Peter Solagna, Vincenzo Spinoso.

Agenda available in Indico

<https://indico.cern.ch/event/434435/>

1. Introduction

This is a kick-off meeting to start discussing about the evolution of the Information System TF. Two topics have been identified for this first meeting: a review of REBUS known issues and a first discussion on whether installed capacities could be useful to experiments. Next steps are also proposed and identified for future meetings.

2. Action item review

REBUS Known Issues are summarised and tracked in the following twiki:

[https://twiki.cern.ch/twiki/bin/view/EGEE/AllAboutREBUS#REBUS known issues](https://twiki.cern.ch/twiki/bin/view/EGEE/AllAboutREBUS#REBUS_known_issues)

Action items are defined during the meeting as described in the next sections. They will be summarised in a table in future meetings and its progress and completion will be recorded.

3. Report on REBUS Issues

Sites disappear when BDII can't be contacted: it is agreed that this is an easy use case that could be fixed already know no matter what future decisions are taken about the Information System. Eddie agrees that it is a fix easy to implement.

Installed Capacities in REBUS are not validated: It is agreed that this issue depends on a more general discussion on what information is actually needed in REBUS. In any case, it is agreed that the way to address this issue would be to define a set of criteria on what type of checks will be implemented. It is also agreed that these checks should be to avoid obvious wrong values and that it is not possible to fully validate installed capacities. The technical solution could be based on the exiting glue-validator. An evaluation of the glue-validator could already be done to understand how flexible it is to include more tests and not rely on LDAP or GLUE if it's finally decided to accept information described in other ways. Maarten mentions that if sys admins enter installed capacities manually, the validator could already interact at the edition time with the user, so that the user can

react on the feedback given by the validator. It is also agreed that central validation in REBUS would allow to use the same validation mechanism for all sources of information if in the end some information comes from i.e. OSG or GOCDB.

Monthly Average Installed Capacity: REBUS currently displays monthly averaged values of the installed capacities. The question is whether these averages are useful or whether it would be better to just show a meaningful value of the capacities every time. It is agreed to postpone this discussion until it is decided what information and how it will be used is decided.

Installed Capacity per VO: Some sites are not publishing information per VO in REBUS. Brian explains that in particular for HS06 information, MyOSG also publishes the average capacity per core for every OSG site. If REBUS could somehow integrate this value, it will be easy to work around the problem. Maria explains that indeed REBUS currently displays the total site capacity and that this is done within the REBUS code dividing HS06/Logical CPUs. In the case this is taken from the BDII, HS06 is also the average capacity per core. This means Brian's suggestion is easy to fix but it can have implications on how this information is currently used. It is agreed to follow this offline and evaluate this possibility.

There is also a discussion in the Vidyo chat on how total HS06 published for site fluctuates in REBUS: CREAM CEs only define once the number of CPUs. This number only changes when HW is added/removed. For ARC CEs, the number of CPUs is collected dynamically. Oxana says that the refresh rate should be the same for all resources to get sensible results. Stephen reminds that installed capacity was defined to include CPU and disk which are temporarily down, so they should change very little. Stephen adds that all these definitions were made before Clouds and that probably clouds would need a different treatment. Opportunistic resources are also something else, and Alessandra points out that they should not be accounted in the site capacity, but maybe in another column. Stephen adds that installed capacities in REBUS are to be compared with pledges, and that opportunistic resources can't be pledged.

Brian requests whether it would be possible to drop the Physical CPU information from REBUS. It is not clear whether it is actually used at all. It is agreed to follow this offline and evaluate whether it is possible.

There is a general agreement that REBUS has multiple use cases today but that it was originally designed for WLCG management, this doesn't mean it can't be extended to cover other new use cases, like including the T3 sites even if they haven't signed a MoU. It is decided that this should be raised at the MB at some point since experiments would be interested in having the T3 in REBUS.

Action items:

- *Implement a fix for sites disappearing from REBUS when BDII can't be contacted (Eddie)*
- *Once installed capacity information is identified, define a set of validation criteria (Maria Alandes collecting feedback from task force members)*
- *Evaluate feasibility of adding new tests to glue-validator, whether it can be easily extended to understand other formats like json or xml, and how it could be integrated within REBUS (Edward Karavakis)*

- *Evaluate feasibility to publish average HS06 per core instead of total site capacity in HS06. Understand who will be affected by this change and whether the change is desired by everybody (Edward Karavakis)*
- *Evaluate the possibility of dropping Physical CPU Information from REBUS (Edward Karavakis)*
- *Ask MB whether it would be possible to include T3 sites in REBUS. This may require to separate a management view of REBUS from an operational view used by experiments (Maria Alandes)*

4. Installed Capacities

Maria summarises how installed capacities are currently published in REBUS. Two types of installed capacities are currently published with different meanings and calculated in different ways. The T1 and T2 installed capacities are used by WLCG management to generate monthly Accounting reports, and basically rely on the Accounting portal information with the possibility of correcting manually the values for T1s. Site capacities used by experiments (only ATLAS and CMS identified) are coming from BDII and MyOSG and refreshed every hour. Maria also presents the current official definitions of installed capacities based on the Installed Capacity Document: https://twiki.cern.ch/twiki/pub/LCG/WLCGCommonComputingReadinessChallenges/WLCG_GlueS_chemaUsage-1.8.pdf.

Andrea presents CMS input for installed capacities. They would like to see pledges per site, although there is a comment from Brian that maybe CMS sites and WLCG federations are not such a different concept, still today a lot of work is done today to be able to map CMS sites to WLCG federations. ATLAS would also have pledges per sites and not per federation. It is agreed that it will be asked to the MB whether it is possible to provide this, since in any case, federation pledges could still exist if needed by the management by adding up all site pledges belonging to the federation.

Andrea also explains that CMS doesn't require installed capacity information but that it would be nice to have it for operational activities like debugging. Maarten explains that ALICE would make use of installed capacity information if it is reliable since this is very helpful for operations. Maria asks whether operations is the only use case for installed capacities. Alessandra explains that ATLAS relies more on installed capacities than on pledges for future planning.

There is a discussion on whether installed capacities should be better referred to as allocated or available capacities, or even planning capacities. They should be a number that experiments could rely on in a monthly granularity. It should not fluctuate daily and it should be per VO.

Action items:

- *Ask MB whether it is possible to publish pledges per sites (Maria Alandes)*
- *Review current definition of installed capacities by collecting feedback from TF members making a proposal to be approved in a next meeting and presented to the MB (Maria Alandes)*

5. Next steps

Use Case document: Maria will distribute a draft document in the ML for people to contribute. This will be edited and presented at a future TF meeting. The final version will be sent and presented at the MB. Maria reminds it is very important to pass a clear message of what experiments and WLCG project activities need from the IS since he has requested to understand what it is actually needed. As soon as Maria gets feedback from Ian on when he would like to have this presented at the MB, she will pass the information.

Sources of information: Maria suggests OSG, EGI and NDGF present at a future meeting their plans for supporting their current portals and tools that are used as sources of information for today's Information System. Brian agrees to present OSG plans. Peter explains that BDII will be still used within EGI since it's a critical component for non HEP VOs and it's also used by the project to provide real time information on the existing resources. Alexey confirms that ATLAS still relies on the BDII to get queue definitions as the list of CEs is coming from OSG and GOCDB, however it's not a critical use case since in AGIS it's always possible to enter information manually.

Information System Schema: it is decided that a discussion about GLUE should happen at some point. Brian explains that OSG indeed sees GLUE as an interoperability mechanism that would allow different grids to talk among themselves but that information systems should not be tied to it and may choose to have other ways to publish information. Stephen reminds that GLUE is not BDII, and that moreover, other renderings exist other than LDAP, like for example JSON or XML.

Experiment Information Systems: Alexey explains that today it's difficult to aggregate information from such heterogeneous information sources. Adapters indeed work but a lot of effort is needed to get correct information from all the sources. Even if a low level information system is required to define basic services and topology, experiment information systems are also needed to provide the experiment topology and other bits of information needed by the experiments. It would be also good to work towards a common definition of the experiment information systems. A presentation about CRIS (computing resource information system) could be also done at a future meeting.

6. Next meeting

Not decided. To be discussed in the ML depending on absences during vacation period and progress on the defined action items.