

Information system and the experiments

Alessandro Di Girolamo
Andrea Sciabà

Pre-GDB
January 15, 2013

- Motivations
- Current limitations
- Current and future use cases
- Possible strategies
- Conclusions



- Since some months there is a renewed interest in WLCG for improving the IS
 - Monthly meetings among experts, experiment contacts, OSG and CERN IT
 - Coordinated by Maria Alandes
- Several motivations
 - Follow up on the Operations and Workload Management TEG's recommendations
 - Better adapt to current and future use cases
 - Better focus the limited available effort
 - Try to consolidate parallel developments of different experiments

- **Instability:** information disappearing due to stuck services, network glitches, etc.
 - Largely mitigated now, thanks to caching in top BDII (provided caching times are long enough!) and better availability of NGI BDII's
- **Validity:** some information has invalid or not GLUE-compliant values
 - Will be solved by proper validation following the EGI profile for GLUE 2 using a new glue-validator
- **Accuracy:** some information stale or wrong
 - Need to understand whether this comes from misconfigurations, wrong versions of information providers or sysadmin mistakes
- **Slowness in fixing issues:** sometimes need to wait for a new information provider, or sites are slow in answering
- These problems led over the years to a mistrust in the system
 - Home-made alternatives developed by experiments to
 - Aggregate from various sources (BDII, GOCDB, OIM, etc.)
 - Cache and validate the information needed by the experiment services
 - But things are starting to improve

- **Service discovery**
 - CE, SE discovery (ALICE, ATLAS, CMS, LHCb)
 - FTS and LFC discovery (ATLAS)
- **Software installation**
 - Publish installed software releases at sites
 - Dynamic CE information on OS, maximum wallclock time, RAM used by software installation jobs (CMS)
- **CE status**
 - Dynamic attributes such as CE state, max total jobs, max CPU time, close SE, estimated response time, . used for **gLite WMS CE matchmaking**
 - Waiting/running jobs, real/virtual memory per slot, CE state used to control ALICE pilot submission

- Publish information for new services
 - E.g. squid for Frontier, xrootd redirector
- SEs without SRM
 - GridFTP, etc.
 - But standalone xrootd should be already there...
- SE capacity
 - Used, free, total online/nearline space
- Opportunistic resources
 - E.g. HPC, clouds, non-WLCG sites
 - Work in progress in EGI?
- Publish information useful for multicore jobs
 - Already covered in a pre-GDB talk last year

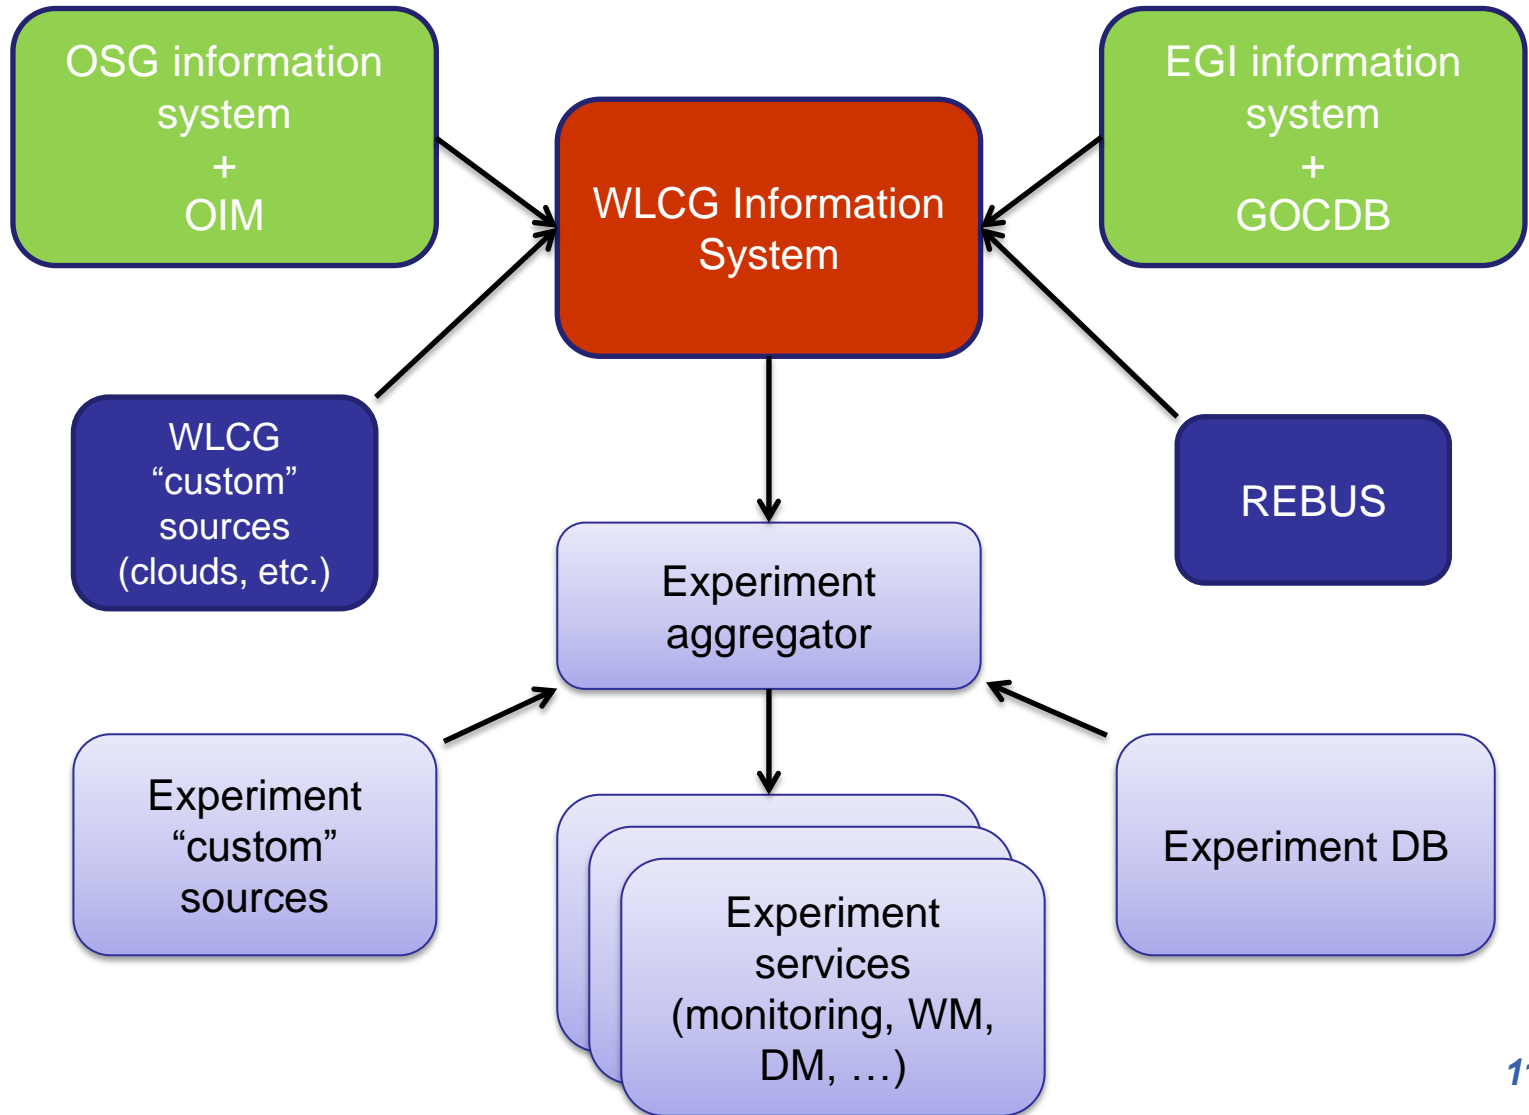
- Software installation
 - With the full adoption of CVMFS, software tags will become unnecessary
- CE status
 - With the decommissioning of the gLite WMS, matchmaking against dynamic CE attributes will not be done anymore



- All experiments developed their own “information systems” using information from several external sources or directly entered by hand
 - AliEN LDAP server, DIRAC Configuration System, AGIS, CMS SiteDB
- The external sources are mostly common
 - BDII
 - GOCDB, OIM
 - REBUS

- Discovery and properties of resources must be the main purpose
 - What sites and resources exist?
 - Everything which is needed to use a service
 - Best if generated by the service itself!
 - Resilient to glitches and downtimes
 - Guaranteed to be correct
- Aggregates information from EGI, OSG and NorduGrid
 - Not all of it, only what is used by experiments
- Uses REBUS, GOCDDB and OIM
- Can publish non-Grid resources (e.g. clouds) but considered part of WLCG
- In other words, performs all VO-independent information aggregation

- Experiment aggregators (e.g. AGIS for ATLAS) just add the experiment-specific information to the WLCG information
 - VO site names, custom services, experiment contacts, etc.
 - Annotations: experiment-specific properties of services
 - External resources uncorrelated to WLCG
- Other requirements
 - REST interface with standard output formats
 - A simple query language
 - A convenient and user-friendly query command



- From the experiments point of view, the schema is an implementation detail
 - but a common schema is essential for information sharing across EGI, OSG and Nordugrid
- Even if not critical, dynamic information should be published only if it is correct
 - Validation efforts are very welcome
 - It might also be retrieved directly from the services themselves, if supported

- Short term
 - Improve the glue-validator and implement the EGI profile checks
 - Extend the functionality of the ginfo client
- Medium term
 - Separate publication of static and dynamic information and possibly move the former to REBUS

- Experiments can successfully use the site resources
 - But a lack of a good information system increases the operational effort and lowers the efficiency
- Resource discovery is by far the most important functionality
- New challenges already now
 - New common services already used (squid, xrootd) for which information (ports, endpoints, etc.) is not easily available
 - New resources (e.g. clouds) need to be discovered and used
- Opportunity to review the information system concept to evolve into a common WLCG service
 - And reduce the need for parallel developments in different experiments