WLCG Future Information System Use Cases

Editor: Maria Alandes
Date: 03.12.2015
Version: 1.6
Contributors/Authors: Alejandro Alvarez Ayllón (GFAL2/CERN), Julia Andreeva (Grid Monitoring/CERN), Alexey Anisenkov (ATLAS/CERN), Marian Babik (SAM/CERN), Stefano Belforte (CMS/INFN), Brian Bockelman (CMS/OSG), Concezio Bozzi (LHCb/INFN), Alessandro di Girolamo (ATLAS/CERN), John Gordon (Accounting/RAL), Maarten Litmaath (ALICE/CERN), Andrew McNab (LHCb/University of Manchester), Stefan Roiser (LHCb/CERN), Pablo Saiz (REBUS/CERN), Andrea Sciaba (CMS/CERN), Andrea Valassi (LHCb/CERN)

Table of Contents

1. Introduction ... 3
2. ALICE ... 3
3. ATLAS ... 4
4. CMS ... 5
5. LHCb ... 8
6. REBUS ... 10
7. SAM ... 10
8. GFAL2 ... 11
9. IT-SDC C5 Report ... 11
10. WLCG Google Earth Dashboard ... 11
11. APEL Accounting System ... 12
12. WLCG Operations and Monitoring ... 12
13. Summary ... 13
14. Definitions ... 14
15. References ... 14

1. Introduction

This document describes the use cases for a future WLCG Information System. It is based on the existing use cases described in the document *WLCG Information System Use Cases* [1].

This document reviews the existing use cases analysing whether they are actually needed or not and whether the same information could be taken from a different source. It also classifies the information in static, mutable or dynamic.

The document extends the review of the existing use cases to also include pledges and installed capacity. It describes how pledges and installed capacities will be used by the experiments and how they will be collected and validated.

2. ALICE

2.1. Summary of existing use cases

Use Case	Needed in the future?	Type of Information	Current Source of Information	Other ways to obtain the same information
Status of CEs	See 2.2	Dynamic (once per minute)	Resource BDII of CREAM CEs and site BDII of ARC CEs	See 2.2
Number of waiting jobs in the VOView				
Number of running jobs in the VOView				

2.2. Future use cases

The dynamic attributes currently obtained through BDII queries are not strictly necessary, since it would be possible to let the VOBOXes make use of HTCondor to submit their jobs to the CEs: as HTCondor keeps track of every job (not really needed for pilots), it can report the numbers of running and waiting jobs "for free". A new AliEn module is being prepared that already makes use of a local HTCondor installation on the VOBOX for job submissions to the new HTCondor CE type. Since HTCondor can also submit to CREAM and ARC CEs, the same module could be used for those types as well. However, as long as the BDII-based method keeps working satisfactorily for a particular CE type, a campaign to consolidate VOBOXes toward HTCondor will not have high priority. With respect to the BDII, it would not be a problem for ALICE to switch from GLUE 1 to GLUE 2 instead.

2.3. Pledges use cases

Pledges represent the expected lower bounds on computing and storage resources available to the experiment for a given period and are used for planning the amounts of data and processing that can be accommodated for the experiment. Since operations are concerned with individual sites, it would be desirable to have pledges per site instead of federation, but that distinction is more important for the installed/available capacities at sites (see 2.4). A rough estimate of the pledge per site would be useful to discover significant discrepancies with said capacities and then follow up. Pledges are typically made once per year, but should be updatable in relation to changing circumstances for sites or federations. The pledge info would be regularly collected from REBUS for inclusion in ALICE internal accounting reports. The automatic validation of pledge numbers should prevent the acceptance of values that are obviously wrong, e.g. due to expression in the wrong units, while the validation of reasonable values can only be done at a higher level involving management and/or experts.

2.4. Installed Capacity use cases

For operations it is important to know per site the amounts of computing and storage resources that are expected to be available to the experiment. Hence the term should be *available* rather than *installed* capacities. The expected values are compared with what is observed in the monitoring of jobs and storage per site. Significant discrepancies, in particular shortfalls, are followed upon with the admins of affected sites. ALICE have their own database for recording such information per site, but would like to have it served by REBUS or a new WLCG service, and not only for MoU sites but also T3 etc. For MoU sites the information can be further compared with the corresponding pledges. If all such information is available for all resources in a single place, it would allow ALICE internal accounting workflows to be simplified. Ideally site admins should be able to update the information for their own site, with ALICE being automatically notified of any changes and able to override and lock the numbers for any site. The information could be consulted as often as daily for internal accounting reports. The automatic validation of capacity numbers should prevent the acceptance of values that are obviously wrong, e.g. due to expression in the wrong units, while the validation of reasonable values can only be done at a higher level involving management and/or experts.

3. ATLAS

3.1. Summary of existing use cases

Use Case	Needed in the future?	Type of Information	Current Source of Information	Other ways to obtain the same information
List of sites	YES	Static	GOCDB/OIM	There are no plans to obtain the information in other ways
List of services		Static	GOCDB/OIM	
CE service details		Static	BDII	
Local queues		Static/Mutable	BDII	
HEP-SPEC06		Mutable	REBUS	
Logical CPUs		Mutable	REBUS	
Site downtimes		Dynamic	GOCDB/OIM	

3.2. Future use cases

ATLAS computing model and data structures used by distributed computing applications and services are continuously evolving. For this reason ATLAS requires a flexible information system which can be easily extended and used to describe new types of services and resources. In particular Cloud, HPC and ObjectStore services should be considered as future use cases for the Information System. ATLAS already declares these type of resources in AGIS.

In general ATLAS would benefit from the existence of a Service Discovery tool containing the list of services which can be used by ATLAS. Additional attributes like the type of the service and its production status would be useful to filter out the existing services. This type of service can be an extension of the present GOCDB and OIM tools. The Service Discovery tool should offer a URL which can be queried to get the detailed static information of the various WLCG services like Storage resources and Computing elements, including Frontier/Squid, Perfsonar, SE redirectors and Object Store service definitions with their available protocols.

For what concerns computing resources pledges, installed capacity, and in general how the ATLAS experiment can understand how many resources are available at each site, it would be useful if the pledges could be defined at site instead than at Federation level. It would be also useful if sites would be able to publish how many resources above the pledge “could be used” by the ATLAS experiment (even if the resources are shared with other experiments). The details on how presently the pledge and installed capacity is used are in the next paragraph, where some suggestions on how to improve them are also described.

The information published by the future information system should be validated *before* being published. The information system should also provide a way to export data programmatically (possibly in a REST-full way, preferably using JSON or XML interchange format).

Apart from the static information used by ATLAS to define the topology of resources, it would be useful if ATLAS could also consume dynamic information. The Service Discovery tool could be extended to publish extra URLs where dynamic information could be queried. A clear separation between static and dynamic information is highly beneficial.

3.3. Pledges use cases

AGIS collects and caches federation pledges from REBUS. If the pledges would be defined within WLCG at the site level it could be a useful simplification for ATLAS. It would be also good if pledge information modified in REBUS could contain the modification time, user DN of the person updating the information and a contact email. In this way, the experiment may know who to contact if some clarifications regarding the pledges are needed.

3.4. Installed Capacity use cases

AGIS also collects and caches the HEP-SPEC06 and Logical CPUs information published in the REBUS Site Capacities view. T3 sites that are not declared in REBUS are manually added to AGIS. This is used by the ATLAS dashboard to calculate resource utilisation. It is also used by the PanDA system to make effective task/job brokerage based on the average power of the computing resources available at each site. ATLAS does not need a precise number of the HEP-SPEC06 and Logical CPUs at the sites, it just needs a reasonable average core power value for each job slot provided by the site. The average core power is calculated dividing HEP-SPEC06/Logical CPUs. ATLAS measures the CPU and wall clock time of each job. For each site, ATLAS then calculates, multiplying by the average core power, the HEP-SPEC06 hours provided by the site to ATLAS. Even if the described process is known not to be perfect, it already allows ATLAS to discriminate between sites with different average HS06/slot values.

4. CMS

4.1. Summary of existing use cases

Use Case	Needed in the future?	Type of Information	Current Source of Information	Other ways to obtain the same information
Compute endpoints and associated information	YES	Static/Dynamic	BDII	See 4.3

4.2. General requirements

To ensure maximum flexibility in the coming years, a WLCG information system must:

- Allow to easily add new domain-specific (such as VO-specific) attributes when need arises
- Allow sites to easily publish new information generated by third-party information providers
- Ensure that all the information adheres to a very clear definition
- Validate the information up to a specified level of accuracy
- Provide the ability for central overrides of incorrect data
- Integrate site data from multiple sources, including both the site and central databases

The rationale is that it is not possible to always know years in advance what information will be needed, and WLCG should not be constrained by a rigid information schema or few information providers which are rarely modified and which might not be usable at all sites. At the same time there should be no confusion on the meaning of the attributes, whose definition should be as stable as possible, and whose reliability should be assumed known a priori.

4.3. Future use cases

4.3.1. Service Discovery

There is no use case for storage service discovery from the WLCG information system. The reason is that storage endpoints are already known to PhEDEx and change only very rarely.

Compute endpoints on the contrary must be discoverable via WLCG services, as they are more numerous and subject to changes.

4.3.2. Dynamic Information

Concerning storage resources, the most important dynamic information is the amount of used space, which in CMS is calculated by a specific monitoring plugin run at the site which provides a detailed breakdown of the space utilisation. The amount of total space is known by private communication with the site. If the information can be easily and reliably produced by WLCG, it would likely be used, but it is not formally requested by CMS as the use case might be covered in future by ad-hoc plugins providing more detailed information (e.g. group/user quotas).

Concerning compute endpoints, all information needed to configure a site entry in the glideinWMS pilot factory must be discoverable from WLCG tools. This includes:

- Resource acquisition contact details (information necessary to send pilots to a given resource, such as queue name)
- Resource acquisition requirements (what type of pilots may access this resource?).
- Resource size (for workflow planning purposes; see 1.6 below)
- Resource runtime environment (cores per batch slot, CPU / wall clock time limits, memory limits, HS06 rating)

The use case consists in eliminating the need for manual maintenance of the factory configuration based on alternative communication channels (but leaving open the possibility to supersede the information coming from WLCG in special cases).

Concerning network information, it would be very useful to have from WLCG information about connectivity between WLCG sites, to be used for data routing optimisation.

4.3.3. Downtimes

CMS would like to have a unified WLCG source for downtime information in the form of a calendar, including at least present and future downtimes with all the relevant details for all WLCG sites.

4.4. Pledges use cases

For long term planning it is important to know the amount of resources pledged at CMS sites. Therefore, pledges should be known at the CMS site level. There are cases when sites provide a capacity that deviates considerably from the pledges for a long time but CMS does not require WLCG to publish this information through a common tool.

4.5. Installed Capacity use cases

Information about installed capacity is considered “nice to have” but not essential. We are primarily interested in knowing the capacity for the purpose of planning workflows and storage utilization; typically, this is done at month-level granularity.

For storage capacity, the considerations made before on total available space apply: the information may be used if available but it is not required from WLCG.

For compute capacity, the motivation is that the information can be inferred relatively easily from the past site utilisation figures. This way of estimating capacity proved to be more reliable than other sources. If the information is made available by WLCG, it can be used to help troubleshooting sites when the observed utilisation level is not the expected one. However, CMS does not request its sites to provide this information.

4.6. Information aggregators

CMS identified a clear need for a service to aggregate all information about resources that is needed by the various computing services. Such service would collect information from several sources and allow information to be directly entered into it. The SiteDB service used in CMS is a partial implementation of this concept but it could also be treated as an information source. Such service would supersede a number of custom mechanisms to retrieve information from different sources, which are subject to become unreliable if not properly maintained, and positively affect computing operations.

A second level of aggregator service is being proposed in WLCG to aggregate all information coming from WLCG sites and resources and limited to what is strictly needed by WLCG. Though it might be convenient, what really matters for CMS is that all the needed information can be retrieved through supported clients and plug-ins and the validity and availability of the information is much more important than the way it is presented.

5. LHCb

5.1. Summary of existing use cases

Use Case	Needed in the future?	Type of Information	Current Source of Information	Other ways to obtain the same information
List of CEs	YES	Static	Top level BDII	In the future: GOCDB (ldap from resource level BDII, and/or json/http from another source)
CEs properties	YES	Dynamic/Mutable		
Max CPU Time	YES	Mutable		
SI2K CPU Scaling Reference	YES	Mutable		
Site properties	YES	Static		
Site downtimes	YES	Dynamic	GOCDB	-
Storage space total/used/free	YES	Dynamic/Mutable	SRM	In the future: BDII and/or GOCDB?

5.2. Future use cases

As explained in the existing use cases document [1], LHCb relies on the BDII primarily to retrieve the list of CEs and the detailed properties of each CE (such as Max CPU time and Scaling Reference factors). LHCb also retrieves site properties from the BDII, but already uses (and may eventually rely only on) GOCDB as the primary source of this information. Finally, LHCb is currently using SRM to monitor the total, used and free storage space at sites, but also this information may in the future be taken from the BDII. For all these use cases, the BDII may be (at least partly) replaced in LHCb by a GOCDB-based infrastructure in the future.

All GLUE information used by LHCb is now retrieved by DIRAC using an ldapsearch for GLUE1 attributes and is then inserted into the DIRAC Configuration System. All attributes are retrieved from the BDII by the DIRAC Bdii2CSAgent, with the only exception of two attributes (GlueCEStateRunningJobs and GlueCEStateWaitingJobs) which are currently retrieved by the DIRAC ARC Computing Element through a module that queries the CE directly without using the BDII.

Work is underway within LHCb to prototype several changes to the current infrastructure, including: the possible move from GLUE1 to GLUE2; the possible use of GOCDB instead of the top level BDII as the central access point for information; and the possible use of http/json queries in parallel to ldap searches. The move to GOCDB would start by publishing ldap URLs from resource level BDIIs directly in GOCDB, bypassing the top level BDII; in a second step, these ldap URLs may be replaced by URLs of other locally managed information sources, for instance using json and http. The timescales for the completion of this prototype and its possible production deployment are not yet clear, also because this work might eventually be done for different experiments in the generic context of the DIRAC project and not only for the specific case of LHCb.

The list of GLUE attributes currently queried by LHCb[1] may also change in the future, but this will only become clearer as the migration to GLUE2 and the prototyping work on GOCDB/http/json progress.

[1] The list of GLUE1 attributes currently used by LHCb that was published in the WLCG Information Use Cases document [1] was incomplete. The complete list is the following: *GlueCECapability, GlueCEImplementationName, GlueCEInfoTotalCPUs, GlueCEPolicyMaxCPUTime, GlueCEStateRunningJobs, GlueCEStateStatus, GlueCEStateWaitingJobs,*

Some of these attributes may be dropped, if it becomes obvious that they are not really needed. Some may also be added, in particular new storage-related attributes providing the information that is presently retrieved from SRM, but also new attributes describing multi-core properties of CEs as the use of these resources in LHCb gains momentum.

Finally, it should be noted that LHCb is not currently relying on the OSG Information System, although this may change in the future. LHCb currently uses only one US site (Ohio Supercomputer Center), but DIRAC pilot jobs are submitted directly into its batch queues, even if this is an OSG site.

5.3. Pledges use cases

LHCb uses pledge information from REBUS. Pledge information is only required per federation: LHCb is not interested in having pledge information per site, neither for operational nor for accounting purposes. REBUS information is typically used as-is and without further processing, except for Tier-2 pledges, where the total pledge for a country without a single federation can only be obtained by computing the sum of all individual federations within that country, e.g. the pledged contribution from France is obtained as the sum of the pledges for CPPM-Marseille, GRIF-Paris, LAPP-Annecy and LPC-Clermont-Ferrand.

5.4. Installed Capacity use cases

LHCb relies only on the information published in the BDII (and re-published on LHCb DIRAC) for what concerns the CPU capacities installed and available at sites. As mentioned above, SRM is presently used instead as the primary source of information about installed storage capacities, although this may change and BDII may be used for this use case too in the future. REBUS is used in LHCb only for obtaining the pledged resources at the various sites.

LHCb provided suggestions in the past to improve the validation of published GLUE attributes for installed resources and is now profiting from their implementation and production deployment. One such example [2] is the validation of Max CPU time to ensure that a published value of 999999 indicates an error and not a queue with no limits. For the moment, LHCb has no further suggestions for the validation of published GLUE attributes for installed resources.

LHCb may find a central WLCG Information System useful as a place to get information from, but LHCb doesn't have a requirement for such a system to the created. The DIRAC Configuration Service collects the needed information and in case a reliable central WLCG Information System is available, this could be queried as another information source alongside GOCDB, especially if the information there is better quality.

GlueCEUniqueID, GlueChunkKey, GlueClusterName, GlueClusterUniqueID, GlueForeignKey, GlueHostApplicationSoftwareRunTimeEnvironment, GlueHostArchitecturePlatformType, GlueHostMainMemoryRAMSize, GlueHostOperatingSystemName, GlueHostOperatingSystemRelease, GlueHostOperatingSystemVersion, GlueSAAccessControlBaseRule, GlueSAName, GlueSEAccessProtocolType, GlueSEImplementationName, GlueSEName, GlueServiceAccessControlBaseRule, GlueServiceEndpoint, GlueServiceVersion, GlueSESizeFree, GlueSESizeTotal, GlueSEUniqueID, GlueSiteDescription, GlueSiteLatitude, GlueSiteLocation, GlueSiteLongitude, GlueSiteName, GlueSiteSysAdminContact, GlueSiteUniqueID, GlueSubClusterName, GlueSubClusterWNTmpDir, GlueVOInfoPath.

6. REBUS

6.1. Summary of existing use cases

Use Case	Needed in the future?	Type of Information	Current Source of Information	Other ways to obtain the same information
Federation Pledges	YES	Static	WLCG sites	NO
Installed Capacities	YES	Dynamic	BDII, MyOSG	NO

6.2. Future use cases

WLCG topology, federation pledges and accounting for WLCG reports will still be collected in REBUS.

As far as installed capacities are concerned, once OSG stops publishing in the BDII, installed capacity view will have to be reconsidered in REBUS. It has to be noted that this is already the case for HS06, as the information for OSG is collected from MyOSG and not from the BDII. Moving installed capacities into the future WLCG Information System seems like a more natural place for this information. In this way REBUS could only focus on WLCG topology, pledges and accounting, leaving service discovery and installed capacities to a different tool.

As far as pledges are concerned and in particular for GridPP, it intends to continue using REBUS to publish pledges on behalf of the RAL Tier-1 and UK Tier-2 federations, and to use the requirements gathered from the experiments in its ongoing planning. For operational and procurement reasons, GridPP would not support any move to require per-site pledges. Currently, unforeseen shortfalls at one site can be made up from others within the federation without forcing purchases earlier than had been planned. The current federation model reduces pressure to publish more conservative pledges than if they had to be guaranteed at the level of individual sites.

7. SAM

7.1. Summary of existing use cases

Use Case	Needed in the future?	Type of Information	Current Source of Information	Other ways to obtain the same information
Queue name	NO	Static	BDII	Experiments VOfeed

7.2. Future use cases

SAM is evolving to stop relying on the BDII to get the queue name. In the future, the experiments VOfeed will contain all the information needed to submit SAM jobs. However, there may be experiments who may still rely on the BDII to get the queue name. If this is still the case, SAM will still have an indirect dependency on the BDII.

8. GFAL2

8.1. Summary of existing use cases

Use Case	Needed in the future?	Type of Information	Current Source of Information	Other ways to obtain the same information
Full SURL	YES	Static	BDII	NO

8.2. Future use cases

GFAL will keep on providing the functionality of contacting the BDII when a full SURL is not provided. It must be noted that in practice this does not happen since full SURLs are always provided.

The default BDII configured in the GFAL software is at CERN: lcg-bdii.cern.ch:2170. If this BDII disappeared tomorrow, a new GFAL release pointing to a different BDII will be needed.

9. IT-SDC C5 Report

9.1. Summary of existing use cases

Use Case	Needed in the future?	Type of Information	Current Source of Information	Other ways to obtain the same information
Installed Capacities	Depends on CERN IT needs	Dynamic	BDII	None

9.2. Future use cases

The information contained in the IT-SDC C5 reports summarises a set of installed capacity information related to Batch systems, Operating systems, computing element flavours and storage types. It needs to be understood within the CERN IT department if these numbers are still needed. If this is not the case, it may be still useful to maintain a dashboard with key WLCG metrics, like number of running jobs or data transfers. In this case, this information can be obtained from the monitoring dashboards, which means that in principle there is no longer a dependency on the BDII.

10. WLCG Google Earth Dashboard

10.1. Summary of existing use cases

Use Case	Needed in the future?	Type of Information	Current Source of Information	Other ways to obtain the same information
Site coordinates	YES	Static	BDII	Possible alternatives would be GOCDB/OIM

10.2. Future use cases

The WLCG Google Earth Dashboard will still need the site coordinates to be able to display the geographical location of the WLCG sites in Google Earth. However, since this is static information, there is no need to get it from the BDII. Either a static file with the list of coordinates or eventually a new field in GOCDB/OIM would be enough. This means that the dependency on the BDII could be dropped.

11. APEL Accounting System

11.1. Summary of existing use cases

Use Case	Needed in the future?	Type of Information	Current Source of Information	Other ways to obtain the same information
CPUScalingReferenceSI00	YES	Mutable	BDII	NO
GlueHostBenchmarkSI00	YES	Mutable		NO
List of Message brokers	YES	Static		NO

11.2. Future use cases

The accounting system will still need information available in the BDII for sites who run the accounting software. This affects the majority of WLCG sites reporting accounting data to the central repository. If the BDII(s) disappeared tomorrow, part of the accounting software will have to be rewritten and all the sites running it would also need to upgrade. In any case, since EGI has committed to keep using the BDII, this is not a scenario to be foreseen. As far as OSG sites is concerned, they are reporting accounting through the Gratia system, so they are not concerned by this.

The central accounting repository will keep on using the BDII to get the list of message brokers. The accounting client software will keep on using the BDII for downloading benchmark values. If a site is running different client software - as is the case with OSG - then they may be able to get their benchmark values from elsewhere.

The default BDII configured in the accounting software is at CERN: lcg-bdii.cern.ch:2170. If this BDII disappeared tomorrow, a campaign to get sites to change their configuration will be needed.

12. WLCG Operations and Monitoring

WLCG Operations and Monitoring will benefit from a central information system describing WLCG topology, where both MoU sites and non MoU sites that have a collaboration with any of the LHC experiments could be described.

A central information system will offer WLCG Operations a full picture of the sites and resources belonging to WLCG, whereas WLCG monitoring will have an easy way to get service endpoints to be contacted for monitoring purposes.

A central information system has the following benefits:

- It collects information from heterogeneous sources. That is, it is able to get information from GOCDB, OIM and BDII, which are the main sources of information that need to be queried to obtain information from WLCG sites and resources.

  - It caches information from heterogeneous sources by regularly collecting information from GOCDB, OIM and BDII, which are currently the primary data sources for WLCG service discovery. The list of primary information sources can evolve in the future.
  - This offers an intermediate layer between the sources of information maintained by EGI and OSG on one side and WLCG on the other side and therefore will provide a consistent interface for all interested WLCG clients.
  - It should also include resources that aren't currently provided or published by EGI or OSG, like HPC or clouds, making the service flexible and open to new types of resources.
- It allows to have full control of the cached information:
  - Validating information before it gets published centrally as well as regular validation of objects which were modified.
  - Applying corrective actions without the need of getting the information fixed at the source.
  - Integrating different ways of obtaining the information: manual input, automatic sources…
  - Logging information, namely when, how, by whom information was provided.

The future WLCG Information System should therefore:

- Include information from OSG and EGI, but also other types of resources needed by WLCG, like HPC or clouds.
- Cache information from the different sources to provide persistent information.
- Validate information before it gets published, applying corrective measures if defined and contacting sites in an automatic way to get the information fixed at the source or directly in the information system.
- Logging associated metadata of the published information.

13. Summary

LHC VOs and also other activities like WLCG Operations and Monitoring, have expressed their interest in a more reliable central information system where information like service topology and installed capacities could be gathered and presented to WLCG in a uniform way. Taking into account that WLCG will have to rely on heterogeneous sources of information coming from OSG, EGI and other resources, like clouds or HPC, it seems logical to collect and present this information in a central service owned by WLCG, where caching and validation of information as well as logging of associated metadata should be performed. Flexibility to integrate different types of resources but also to include MoU and non MoU sites is also an important requirement for the future information system that will have to deal not only with grid resources in official WLCG sites, but that should also extend to all types of resources publishing information in many different ways.

14. Definitions

14.1. Static Information

Static information is information that is constant throughout the lifetime of a service. A collection of this type of information is what we call a service registry. Service registries are used for service discovery.

14.2. Mutable Information

Mutable information may change during the lifetime of the service, mainly due to configuration changes. In order to get mutable information, information could be periodically polled or could use messaging to propagate updates in an automatic fashion.

14.3. Dynamic Information

Dynamic information is highly-mutable information, mainly state changes. This is basically monitoring information. Messaging is the technology most suitable to get monitoring information since BDII has shown not to be ideal as it is fairly long to propagate changes.

15. References

1 WLCG Information System Use Cases: https://espace.cern.ch/WLCG-document-repository/Technical_Documents/WLCGISUseCases_1.4.pdf

2 M. Alandes, "BDII Use Cases for LHCb", 3rd LHCb Computing Workshop (May 2014), https://indico.cern.ch/event/278289/#2014-05-22