An implementation plan for top-level BDII hosted by the GOC

Rob Quick,[1] Scott Teige,[2]
Indiana University

1 Introduction

We present a plan to implement a top level WLCG BDII service at the Grid Operations Center (GOC). The GOC currently operates several services for the OSG including a top level OSG BDII service which also acts as a site level BDII for OSG sites in the WLCG. It has been previously proposed and presented to OSG management that a North American site host such a service and that proposal is included here as appendix A. Since that proposal, CERN has formally requested[3] that the GOC host this service on behalf of the North American LHC Tier 1 sites.

To a large extent, this document presents the policies and procedures used to operate the OSG top level BDII. It is anticipated the WLCG BDII hosted by the GOC will operate under identical procedures.

The interconnection of various components of the proposed service is shown in Figure 1.

1.1 Overview

A technical specification for the service[4] has been produced. Appendix B gives the relevant requirements from this specification and points to the relevant sections of this document addressing those requirements.

It is proposed that the OSG purchase four servers to be housed in the Indiana University (IU) data center in Bloomington. A top level BDII service operated by the GOC will be implemented on these servers. In the following, we describe the proposed system from the hardware level to the expectations of end users.

2 Hardware Implementation

The servers hosting the service will be located in the IU data center in Bloomington. This differs from the current implementation used by GOC for its other services in that most GOC services are hosted at both Bloomington and Indianapolis. When initially implemented, the Bloomington data center did not exist and this distributed approach addressed reliability concerns. We believe the IU data center is sufficiently reliable that these concerns no longer exist.

2.1 Physical Facilities

While no facility can ever be 100% failure-proof the IU data center houses vital enterprise infrastructure associated with the daily operations of the University. In the unlikely event of a problem causing failure of the machines housed in the center, it will certainly be the case that the substantial resources of the University will be used to correct it as soon as possible.

The IU data center is designed to withstand category 5 tornadoes. The facility is secured with card-key access and 7 x 24 x 365 video surveillance. Only staff with systems or network administration privileges have access to the machine room. Fire suppression is provided by a double interlock system accompanied by a

[1]rquick@iupui.edu
[2]steige@indiana.edu
[3]https://ticket.grid.iu.edu/goc/viewer?id=9893, https://ticket.grid.iu.edu/goc/viewer?id=9892
[4]https://twiki.cern.ch/twiki/pub/LCG/WLCGISArea/BDII_Deployment_Plan.pdf.

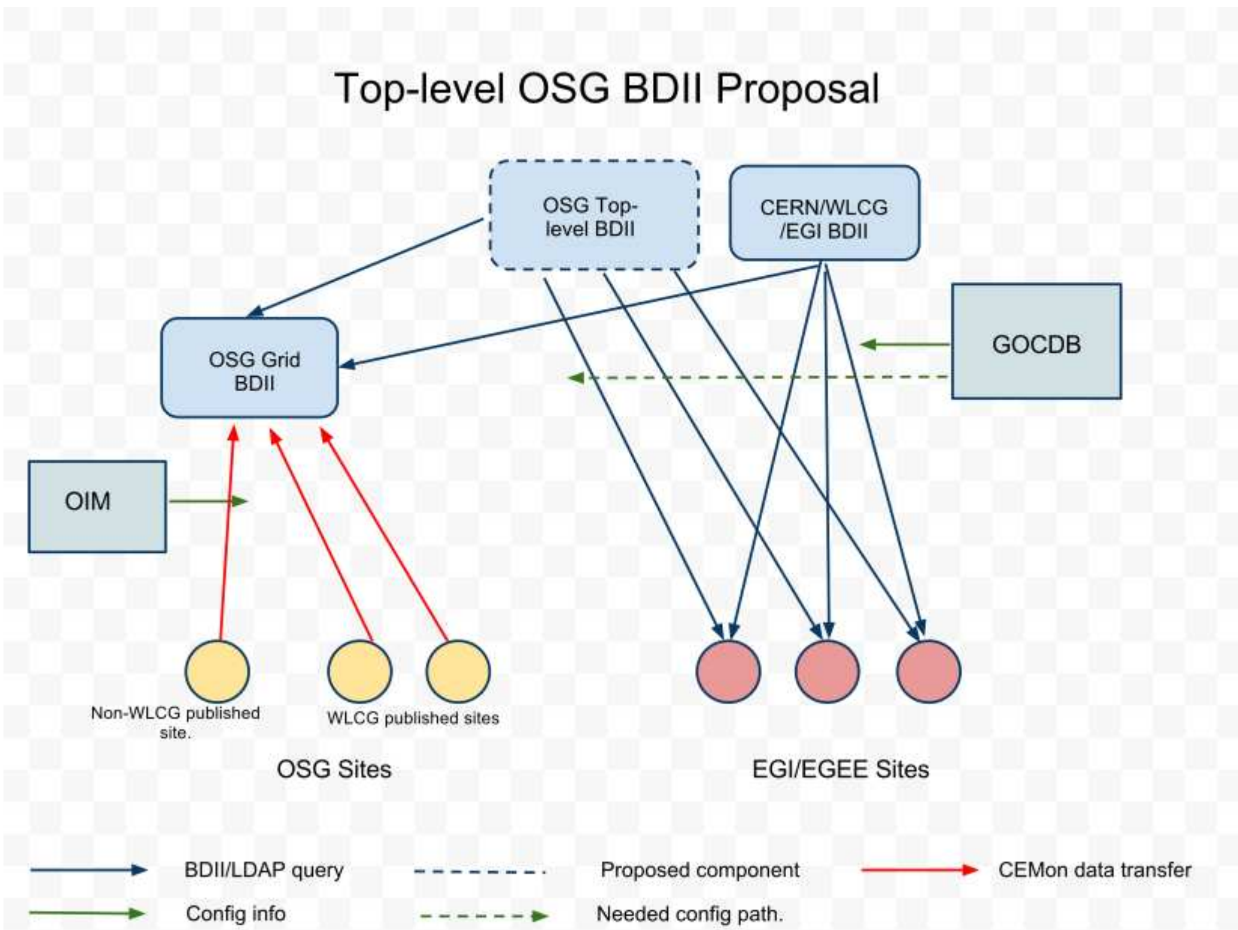

Figure 1: Data flow in the proposed system

Very Early Smoke Detection Apparatus (VESDA). Three circuits feed the Data Center, traveling redundant physical paths from two different substations. Any two circuits can fully power the building. A flywheel motor/generator set conditions the power and provides protection against transient events and uninterruptible power supplies protect against failures of moderate (∼1 hour) duration. Dual diesel generators can provide power for 24 hours in the event of a longer term power failure. In house chillers provide cooling. Externally supplied chilled water plus city water can be used in the event of a failure of this system.

The IU Bloomington campus connects to the IU/Purdue University at Indianapolis (IUPUI) campus via I-Light[5]. Two redundant fiber bundles following different paths implement this connection. From the Indiana GigaPOP in Indianapolis, IU has the following connectivity to external networks:

- Internet2: 2 x 10 GigE
- NLR: 1 x 10 GigE
- TeraGrid: 2 x 10 GigE
- Commodity Internet: 2 x 10 GigE

IU is responsible for the operations of many national and international networks, including Internet2, National Lambda Rail (NLR), TransPac2, the MAN LAN research exchange point in New York City, the Hybrid Optical Network Initiative (HOPI), the Indiana GigaPOP, the TeraGrid's IP-Grid network, and the CIC OmniPoP in Chicago.

2.2 Server Configuration

Because of the rapidly changing nature of the server market we do not specify the actual server configuration at this time. We present here the specifications and price of a server with capabilities described in the technical specification. Sections 5.1 and 2.3.1 for descriptions of other implementations.

- Processor: 2x 4 Core, 2.6 GHz, AMD 4130
- RAM: 8 GB
- Drives: 3x 250GB 7.2K RPM
- Total Storage: 250 GB
- System Cost: ∼$4200

Three such servers would be a minimal requirement giving a total of 24 cores and 24 GB of RAM. We request four to allow spare capacity and allow maintenance to be performed without degradation of service. We note the current implementation at CERN uses five roughly equivalent machines.

The hard drives are the smallest currently available from DELL and the GOC considers RAID plus hot spare to be the minimum acceptable configuration for a production service. This system exceeds the suggested minimum configuration in the technical specification document as justified below.

2.3 Server justification and Constraints defined by CERN

2.3.1 Top level OSG BDII implementation

The GOC hosts a top level BDII service for the OSG. Currently two instances are implemented on physical (as opposed to virtual) machines. A DNS round robin (keep-alive time 60 minutes) makes these available

[5] http://www.iupui.edu/˜ilight/

to users. This system handles ∼2,000 queries per minute averaged over time. A total of 16 2.66 GHz cores and 16 GB of RAM are used. This system can handle current demand with occasional degradation of service quality during periods of unusually high demand. Running with a single server results in degraded performance sufficiently severe there are query timeouts and service alarms.

Figures 2 through 9 show the demand placed on one of the two OSG top level BDII servers used by the GOC. While it is expected that use of the proposed service will be different, we take the demand placed on the existing service as typical. The following discussion demonstrates the proposed hardware configuration could meet the constraints proposed by CERN under the assumption that use of the new service will be similar to use of the existing service.

The request rate (fig. 2) is remarkably stable and shows the service is heavily used. The total network load (fig. 3) reflects this stability of usage and shows the load is not approaching the limit of the network interface card (1000 Mb/sec). The number of processes (fig. 4) and number of open connections (fig. 5) show more variability and demonstrate the random nature of request arrival timing.

2.3.2 Memory

Figure 6 demonstrates application memory use on average is 684 MB, less than the CERN proposed limit of 1 GB. 8GB is the current minimum memory configuration offered by DELL.

2.3.3 CPU

Figure 7 shows the system cannot achieve the required 40% maximum CPU use with a 4 core machine. This motivates the proposed 8-core nodes as opposed to the suggested 4-core minimum in the CERN request.

2.3.4 Network

Figure 3 shows the network traffic generated at the Bloomington instance of the GOC BDII. This machine has 1 GigE connectivity all the way to the site boundary router. Current use is on average 10% of this theoretical maximum easily satisfying the requirement that it be less than 40%.

2.3.5 Response time

Figure 9 tracks the time to complete a request. This request is generated by a machine in the Bloomington machine room to machines in the IUB and IUPUI machine rooms. The network connections between these machines is 1 GigE (minimum) end to end and this time is not typical of requests made by random machines elsewhere. Requests made at locations with inferior network connectivity have been observed to take as long as 1 minute to satisfy. In extreme cases timeouts (at 15 minutes) have been observed. These were all associated with specific sites with severe connectivity problems that were later corrected. Given this variability in conditions under which a request is made it is difficult to state directly that the proposed system will be able to satisfy the requirement that response time be less than 10 seconds.

2.4 Fail-over

Linux Virtual Server (LVS)[6] will be used to load balance and provide fail-over capability.

While the CERN implementation of the BDII uses DNS round robin (RR) for these purposes (and considers it best practice), the proposed system will not. Implementation of DNS RR at Indiana has revealed several problems with using this approach there and for OSG users.

Often, external DNSs do not obey the keep-alive time parameter so when a server is removed from the RR users depending on one of these DNSs often get routed to the server removed from service. Additionally,

[6] http://www.linuxvirtualserver.org/

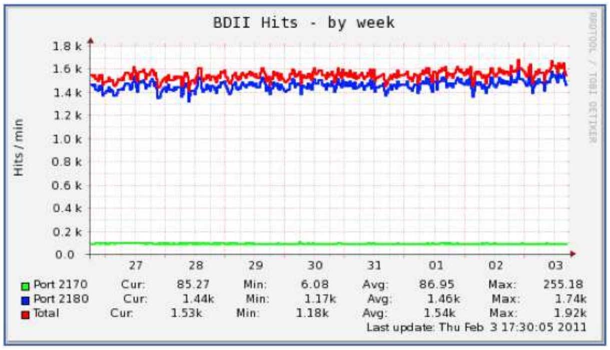

Figure 2: Number of BDII requests per minute

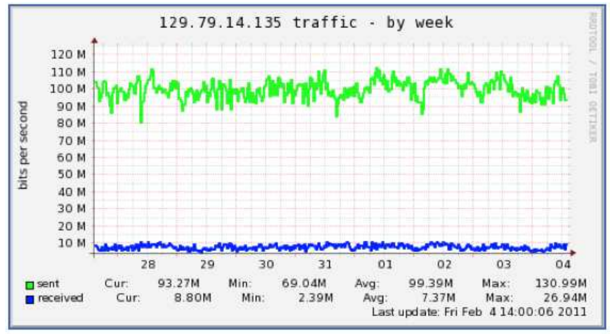

Figure 3: Network traffic volume. On average, 10% of the 1GigE capacity is used, significantly less than the proposed 40%.

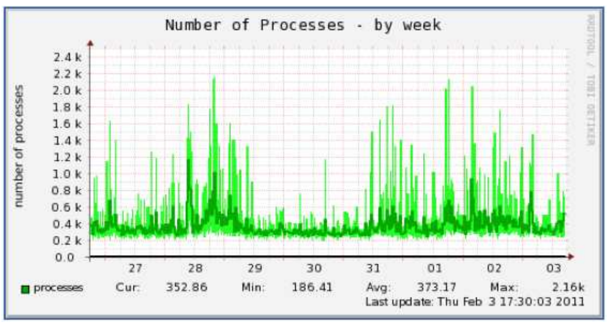

Figure 4: Number of running processes. Shown is the instantaneous value sampled every five minutes

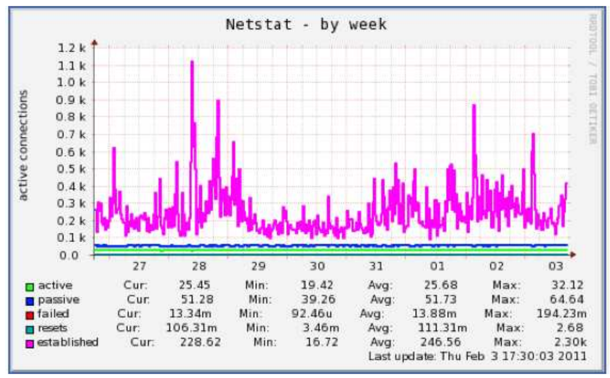

Figure 5: Number of open connections, sampled as in figure 4

the DNS configuration at IU is not controlled by the GOC and changes to it occur only at the top of the hour. This can lead to an outage of up to one hour.

LVS addresses both of these problems. The service will have a single IP address with the work shared transparently across a collection of servers. The configuration of LVS will controlled by the GOC, changes can be made at any time.

Control of the LVS configuration will allow the number of servers handling load to be quickly changed during periods of maintenance or server failure. For maintenance a single server can be removed from the service, modified as required and returned without the delay of an externally controlled DNS.

2.5 Change management

Currently, the GOC adheres to an Information Technology Infrastructure Library (ITIL) inspired change management policy.

For the GOC BDII service, three layers of instances exist, Development, Test and Production. The development instance is held to no standard of service, it may be inoperable at any time, for any length of time. The Test instance is updated on the first and third Tuesdays of the month. The community is notified of the change and requested to test the changed instance. ATLAS and CMS stake-holders are explicitly asked to sign off on the operability of the test instance and no change to the production instance is made

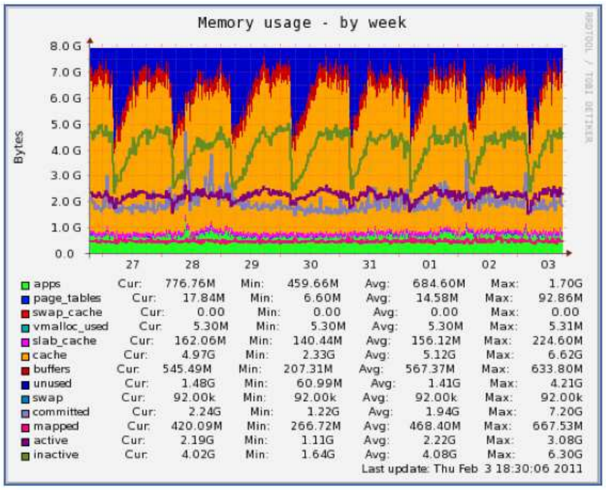

Figure 6: Memory use. The sum of memory used by all applications is, on average, 684 MB, significantly less than proposed 1 GB limit.

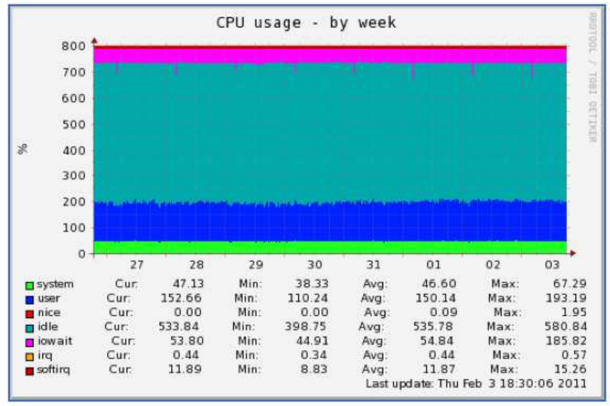

Figure 7: CPU usage (normalized to 800% since the measurement is on an 8-core machine). 18.8% is consumed by user processes and overall the CPUs are idle 67% of the time implying total use of 33%. On a 4-core machine total CPU use would exceed the proposed limit of 40% and user CPU would often exceed that limit.

until they do.

A meeting is held once per week where any proposed changes are discussed. Typically attending are the Project Manager, Production Coordinator, GOC lead, GOC infrastructure lead and GOC service lead. Others often attend and specific people are invited when issues of direct concern to them are to be discussed. No changes are made until consensus is reached at this meeting.

If a change is approved the Production instance is updated on the second or fourth Tuesday of the month. An announcement to the community is made well in advance of the change and the change is made during an announced time window. After the change, various monitors and probes are watched directly by GOC personnel to check for problems. (As opposed to the e-mail or text message alarm methods usually used.) A rollback mechanism is explicitly included in the update procedure should any problems be encountered.

3 User Expectations

3.1 Availability/Reliability

The current Service Level Agreement[7] (SLA) for the GOC BDII states:

> The GOC will strive for 99% service availability. If service availability falls below 99% monthly as monitored by the GOC on two consecutive months a service plan will be submitted to the OSG stake-holders for plans to restore an acceptable level of service availability.
>
> A maximum of two non-scheduled outages will be accepted by OSG during each six month period of service. If the GOC experiences more than the allotted outage, a service plan will be submitted to the OSG stake-holders with plans to restore the service to an acceptable level of operations.

[7] https://twiki.grid.iu.edu/bin/view/Operations/BDIIServiceLevelAgreement

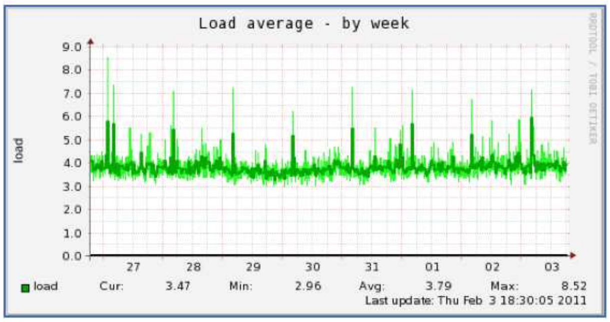

Figure 8: System load

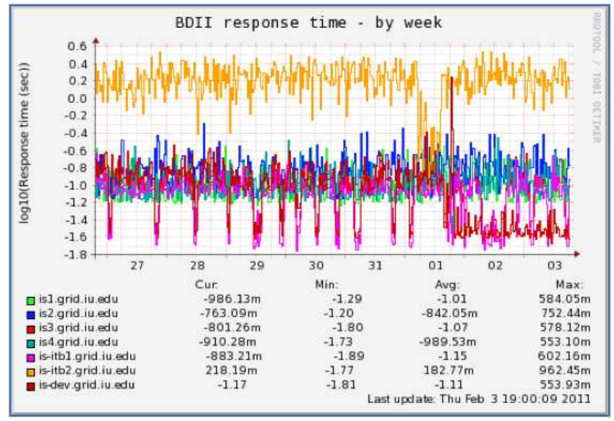

Figure 9: Response time. Note the logarithmic vertical scale. The relevant measurements are green and blue. Average values for theses quantities indicate an average response time less than 0.1 seconds.

It is anticipated the proposed service will be expected to achieve similar performance and be covered by a similar SLA defining this service as critical and requiring it to be covered $24 \times 7 \times 365$.

Table 1 gives the achieved reliability and availability of the GOC BDII for the past six months.

Service	Aug	Sep	Oct	Nov	Dec	Jan
OSG BDII-1	99.66/99.66	99.73/99.89	100/100	100/100	100/100	100/100
OSG BDII-2	99.73/99.73	99.91/99.96	99.95/99.95	99.81/99.81	99.98/99.98	99.98/99.98

Table 1: Availability/Reliability (%) of GOC BDII services for the past six months

3.2 Use cases

A detailed study of use cases for the CERN BDII is currently in draft[8]. An overview of some known use cases is presented as appendix C

Use cases for the OSG top level BDII have been extensively studied. The use of this service, while not expected to be the same as for a CERN BDII, can be used to demonstrate the effect of heavy use on the servers implementing the service.

3.3 Update Frequency

Site level WLCG BDIIs are updated with a period of a few minutes and a top level BDII pulls this data also with a period of a few minutes. Worst case data age is about 10 minutes. Most OSG BDII sites report at 2 minute intervals and for display purposes, data older that 10 minutes is considered out of date.

[8] https://twiki.cern.ch/twiki/pub/LCG/WLCGISArea/WLCG_IS_UseCases.pdf

4 Schedule and effort

The project start date, defined as day 0 below is currently undefined. The schedule, relative to day 0 is given below.

Day 0 Encumberable funds exist, hardware procurement process begins

Day 14 Hardware arrives

Day 16 Hardware racked, OS and Network functional

Day 30 Testable BDII structure in place

Day 31 Testing begins

Day 45 No sooner than after two weeks of testing, release to production.

The timing of the schedule shows the one-time effort required to bring the service into production. Procurement at IU is straightforward and administratively handled by the purchasing department as part of our F&A cost. Installation of the hardware, operating system and network configuration will be handled by the GOC system administrator and is a familiar procedure. We allow two weeks for the installation of the BDII software and will request assistance from CERN personnel for this. Also in this period existing monitor and accounting tools will be configured for the new service. An automated install script, consistent with GOC policies, will be created and tested. LVS will be configured, further tested and implemented for this service.

As a critical production service, extensive testing by stake-holders will be required before it is released for production. The GOC will perform basic stress tests as part of the installation procedures. A minimum of two weeks is allocated for this user testing and approval, five weeks is a probable upper limit on the duration of these tests. Once consensus as to the acceptable operability of the service is reached it will be released to production on the next GOC service release date. (Currently the second and fourth Tuesday of a month)

Recurring effort is required for the operation of the service. We estimate 0.2 FTE will be required. This includes discussion of the service in the weekly GOC operations meeting and OSG production meeting. We anticipate changes to the service will be rare and handled by the largely automated procedures already in place at the GOC. The servers will be incorporated into an existing automated system administration tool, currently *puppet*. Support for the software will be provided by CERN under currently committed effort. Effort associated with monitoring of the service will be largely in the form of a response to a failure event and every effort to minimize those events are part of this implementation plan.

5 Current CERN Implementation

5.1 CERN lcg

The top level BDII service at CERN currently uses 5 machines of the type used as standard batch nodes (1-2 disks). One has 4 cores and 8GB the others have 8 cores and 16 GB of RAM for a total of 36 cores and 72 GB of RAM.

A Previous proposal

A.1 WLCG Deployment Strategy for Top-Level BDIIs

The Top-Level BDII services enable the discovery of Grid services along with further information about their structure and state. They aggregate information from all Site-Level BDIIs to provide a single point which can be queried to find the overall status of the Grid. Each instance of the service must be carefully managed as a critical service for WLCG. This not only requires timely support to ensure that issues are dealt with in a timely manner but also that managed fail over is in place for cases of service unavailability. In addition the service must be scalable to the number of queries which are made against it and also to the volume of data describing all the services. This proposal identifies the critical instances of the service for WLCG in order to ensure that those instance are well managed and that support effort is focused only the instances of the service that are mission critical for WLCG.

A.2 Selection Criteria

The first important consideration is query load. Each service instance must be scalable to exceed the typical query load to which it is exposed in-order to handle peak load. As each computing job can potentially query the Top-Level BDII service there is a strong correlation between the number of logical CPUs in a cluster that is configured to use a specific Top-Level BDII service and the number of queries that it experiences. The distribution of the BDII instances should therefore reflect the distribution of computing resources . Queries also originate from other sources such as the WMS, FTS and Grid Monitoring utilities. Although they do not place a high query load on the service, the queries they use are typical expensive such as returning all information (currently 100Mb).

The other consideration is network latency. Network latency significantly affects the query response time and the reliability of the queries. Due to this, compute resources and other consumers of information should query a 'close' instance from the perspective of network latency and also fail-over to a close instance.

The third consideration is available support effort. The service must be managed as a critical service for WLCG, which requires active support with releases being carefully followed. The services must be intensively monitored, to quickly identify failures, and also to measure the query loading in-order to provide additional resources when required. Currently, five physical hosts are recommend for a Top- Level service.

A.3 Identification of Instances

The latency and fail over requirements suggests that there should be at least two instance per continent: two in the North America, two in Europe and two in Asia. The requirement that the number of instances of the Top-Level service should reflect the distribution of computing resources suggests that the number of instances for Europe should be higher. Although sites in the US contribute a significant amount of computing resources, local policy results in less queries being made from the compute resources themselves and as such this requirement can be relaxed. The initial suggestion is;

- North America: Triumf and BNL or FNAL
- Europe: CERN, RAL, FZK, PIC, CNAF, NIKHEF
- Asia: Taiwan and KEK or Tokyo

This list can be extended as the need is identified.

A.4 Configuration for Sites

Sites should configure their primary Top-Level BDII to be the closest instance from the perspective of the network. For configuration options where fail-over is available, the fail-over instance should be the second closest from the perspective of the network.

B Service requirements

The bullet items are extracted from the CERN generated technical requirements. The italicized comments point to the relevant sections of this document.

- The service must be treated as critical.
  see section 3.1, SLA for this service will be identical
- At least 3 4-core physical (avoid virtual) machines with a minimum of 4GB physical memory in load-balancing (round robin) configuration
  see sections 2.2 and 2.4
- The machines must be strictly monitored so that the average CPU and Network utilizations do not exceed 40%.
  see figures 7 and 8 and discussion in sections 2.3.3 and 2.3.4
- The service must be 99% available, where available means:
  see table 1
  - All BDII processes are running
  - Query response time is less than 10 seconds
    see figure 9 and discussion in section 2.3.5
  - Data is less than 1 hour old
- Sites are encouraged to instrument their own monitoring probes to check data freshness
  Currently, an RSV probe to check freshness of data in the top level OSG BDII exists and is run hourly. It requires data no older than 30 minutes. GOC web pages[9] *require data no older than 10 minutes*
- The service must be strictly monitored
  Figures 2 through 9 are all produced by munin*, a system monitoring tool. Any quantity monitored in this way can generate alarms and configuring new munin monitors is a straightforward task. Figure 9 is an example of a custom munin monitor.*
- Releases must be carefully followed.
  see section 2.5.
- Service administrators should include rolling back to previous release as part of the upgrade plan.
  see section 2.5.

C Some use cases

C.1 ATLAS

- The schedconfig loader, which helps to configure PanDA, uses some basic BDII queries. Explicit details are contained in the source code.[10] [11] The purpose is to locate Compute Elements which support ATLAS and published software tags.
- The File Transfer Service (FTS) uses the BDII to find the Storage Resource Manager (SRM) endpoints. The Glue Classes/Attributes involved are the GlueSite and the GlueService Classes.

[9] http://is1.grid.iu.edu/cgi-bin/status.cgi , http://is2.grid.iu.edu/cgi-bin/status.cgi
[10] https://svnweb.cern.ch/trac/panda/browser/autopilot/trunk/lcgInfositeTool2.py
[11] https://svnweb.cern.ch/trac/panda/browser/autopilot/trunk/lcgLoad.py

- For the distributed data management (DDM) components, the BDII is contacted by Space Monitoring. The relevant attributes are:
  - GlueSATotalOnlineSize
  - GlueSAReservedOnlineSize
  - GlueSAFreeOnlineSize
  - GlueSAUsedOnlineSize
  - GlueSACapability
- Other components that use the BDII are:
  - TiersOfAtlas
  - AGIS - The ATLAS specific information system
  - Service Availability Monitoring (SAM) - Specific ATLAS tests
  - Other monitoring tools, such as dashboards and Site Status Board (SSB)

C.2 CMS

- CRAB uses the gLite, which depends on a top-level BDII. In particular the BDII is used to select production CEs with appropriate CMS software installed and close to SE with needed job data.
- Physics Experiment Data Export (PhEDEx) uses FTS which in turn uses BDII as above.
- Service Status Board (SSB) storage monitoring used/free space information comes from the BDII.
- Software installation tools to find out the sites and publish software tags.
- Other dashboard queries
- One SAM test to list the published software tags
  - GlueHostApplicationSoftwareRunTimeEnvironment

C.3 ALICE

Alice only queries the resource BDIIs of ALICE CREAM-CEs

C.4 LHCb

- LHCb also uses FTS

C.5 DIRAC

- DIRAC also uses gLite