

LHCb WLCG Critical Services

1. INTRODUCTION

In this document I've attempted to summarize, categorize and define some metrics for all our LHCb dedicated services in WLCG. This list might be still not fully accurate and does not include all DIRAC services whose single weight has been aggregated into a final, symbolic "10", in the criticality of the generic VO-BOX service (at CERN).

Of course finer grained information about all DIRAC services running on VO-BOXEs will require further iteration.

For each of these services I've tried - driven by the experience but also taking into account what already assigned and discussed early in February for CCRC (see for more information <https://twiki.cern.ch/twiki/bin/view/LHCb/CCRC08>) to give a description of the usage we effectively do.

The description was also to justify the criticality assigned. A not exhaustive list of plausible "service-oriented" tests that should mimic the daily activity of the services is also sketched. Of course the limitations of these tests are known:

"ad hoc WN setup for SAM", "no stress conditions as it happens in reality", "we take just a snapshot of the status of the service".

Final remark: I have applied a "tier distinction" on the grid services used by LHCb because different are the services (i.e. their weights) involved at different Tiers. The importance associated to each service is absolute and for that I've really used the gradation of numbers suggested by Nick with some approximation.

It means that by rating the criticality of a service I start by considering the importance of the general service (that a given instance is part of) and I weight it by dividing the general service criticality by the number of potential instances that would make the service transparently up and then contribute to its effectiveness.

For the LFC RO services (our LFC distributed service), the result of this operation represents the single LFC mirror instance importance.

2. INVENTORY

List of Services at the T0 and their use in LHCb

- Network Campus: PIT-TO data transfers, data transfers from CASTOR to local WN, upload of data to SE from CERN computing resources, file access in reprocessing/recons activity and analysis (what else? who never knows the impact of the CERN LAN?)
- AFS servers: serving the Shared Area on the WN at CERN and m/w software on our UIs.
- CE: access to Computing Resources via GRID

- WN: Computing Resources for data recons, reprocessing, analysis MC productions
- SRM/SE: raw data archiving from PIT, raw data redistribution, rDST upload (fraction competing to CERN), DST redistribution (to/from T1s), MC data upload, data access recons/stripping/analysis
- VOMS: accessed each time a fresh proxy is required.
- LFC: Master instance for replica information registration + RO instance for replica information access.
- FTS: Data movement service used for raw data redistribution and for DST data uploads (from T1 to T0).
- VO-BOX: hosting all critical DIRAC services. The basic functionalities (left to FIO people like /var full /tmp full load and so forth) are monitored but not the services running behind the scenes
- RB/gLite WMS: (Pilot job submissions, pilot jobs status and logging info query and output sandbox retrieval)
- Oracle Streams (serving both Condition DB and LFC distributed service)
- AMGA service (underlying the BK service for accessing the Oracle BE in the job provenance service)
- SAM service (used for monitoring all Grid resources)
- ConditionDB: Master r/w and read-only instance
- Log SE: (used for uploading and surf through the job log)

List of Services at the T1

- CE (for accessing Computing Resources)
- WN (Computing resources for recons, stripping and user analysis)
- Shared Area (served via distributed file system NFS/GPFS/AFS)
- SRM/SE (receiving raw from T0, upload rDST form WN, DST redistribution, data file access on recons/stripping/analysis)
- LFC mirror instance (for accessing replica information)
- FTS (for DST redistribution across T1)
- VOBOX at T1 (for messaging service and failover mechanism redundancy)
- Eventual RB/gLite WMS for redundancy
- ConditionDB replica mirror.

List of Services at all other centers

- CE and WN for production MC jobs
- Some large T2 centers might offer SE service for hosting analysis too.

3. TESTS

Envisaged tests @T0

Network Campus: under SLS monitors

AFS servers: Lemon or SLS

CE: SAM jobs

- jobs submissions test job retrieval test (as part of the WN test chain)
- VOMS FQAN correct mapping check
- Information System publication: queue length advertized and consistency with SAMDB information (see recent checks from Joel)

WN: SAM jobs

- Run all GAUDI application in cascade (as part of the Installation and Validation of Application Software process)
- Check of the shared application area
- Check of minimum h/w requirements
- Upload of a file to the various space tokens at CERN

SE/SRM: SLS from FIO (basic metrics) + SAM dedicated tests

- Querying metadata tests:
 - Lcg-ls and loop over all these files retrieving meta information
- Staging tests: a file (surely not in the disk cache) must be staged from tape
- Access test: opening a SURL using gfal/ROOT application from within a WN (already in place)
- Uploading from a UI a test files on all Space tokens defined in the tier
- Exporting a test file from T0 to T1 using FTS (in place)

VOMS: under SLS + some basic SAM tests

- Using all CAs in LHCb (at least CERN,IT,UK FR,DE and ES) run for each proxy holder voms-proxy-init and ask role production

LFC: under SLS + SAM

- Browsing the lfc namespace
- Asking for a replica through gfal
- Verifying the replication of new information across LFCes at T1
 - ACL replicated correctly
 - New users or groups replicated correctly
 - Replicas replicated correctly
- Testing the insertion of new entries in the Master LFC
 - Verification of the correct VOMS FQAN mapping (e.g normal users cannot write into production area)

FTS: SAM tests

- Transfer job submission
- Transfer job status

VO-BOX: under Lemon for basic metrics (hardware and OS)

- **NEED to expose Raja's sensors**

gLite WMS: SAM Test

- Fake job submission, job status & logging info, and output retrieval

Oracle Streams: Under SLS monitoring + Information from some LFC SAM

AMGA: Lemon of the BOX (SLS in future) + Some LHCb specific tests?

(Using Elisa API it would be nice to setup some basic functionality tests)

Log SE: Implementing SLS sensor checking the basic box matrices (FIO) gridftp status and apache.

Envisaged tests @T1

CE: SAM jobs

- jobs submissions test job retrieval test (as part of the WN test chain)
- VOMS FQAN correct mapping check
- Information System publication: queue length advertized and consistency with SAMDB information (see recent checks from Joel)

WN: SAM jobs

- Run all GAUDI application in cascade (as part of the Installation and Validation of Application Software process)
- Check of the shared application area
- Check of minimum h/w requirements
- Upload of a file to the various space tokens at the T1 and CERN SE

SE/SRM: SAM dedicated tests

- Querying metadata tests:
 - Lcg-ls and loop over all these files retrieving meta information
- Staging tests: a file (surely not in the disk cache) must be staged from tape.
- Access test: opening a SURL using gfal/ROOT application from within a WN (already in place)
- Uploading from a UI a test files on all Space tokens defined at the tier
- T1 to T1 matrix connection using FTS (in place)

LFC: SAM

- Browsing the lfc namespace
- Asking for a replica through gfal
- Verifying the replication of new information across LFCes at T1
 - ACL replicated correctly
 - New users or groups replicated correctly
 - Replicas replicated correctly

FTS: SAM tests

- Transfer job submission
- Transfer job status

VO-BOX: Local Site Monitoring System?

- **NEED to expose Raja's sensors for all DIRAC services at T1**

gLite WMS: SAM Test

- Fake job submission, job status & logging info, and output retrieval

4. CRITICALITY

Criticality of Services and (rough) motivation @ CERN

1. **CERN Network Campus** (other top critical services depend on that): **10**
2. **SE at CERN:** (All activity from the PIT, to the T1 and on the T0 depends strongly on the shape of this service): **10**
3. **VO BOX at CERN:** (RunDB, AdtDB, IntegrityDB, Processing DB, transfers client, DIRAC WMS and so forth, all run on this box): **10**
4. **AFS at CERN:** (Data recons/analysis at CERN won't happen) The CERN VO Boxes –running central services- also rely on AFS for the gLite m/w : **10**
5. **VOMS:** (VOMS server is not accessed constantly: As soon as production shifter got a 7 days valid voms proxy a 8 hours disruption of this service could be accepted): **7**
6. **CE at CERN :** (Data recons/analysis at CERN won't happen): **5**
7. **WN at CERN:** (Data recons/analysis at CERN won't happen): **5**
8. **LFC Master at CERN** (Critical service for all steps of the data flow: Data coming from PIT, data exported to T1s, data registered from recons&stripping.: **10**
9. **LFC Replica at CERN** (~10/7: all RO LFC replicas together are **critical**, a single replica is just a reduced effectiveness): **3**
10. **AMGA:** (Critical service but failover mechanism is in place): **7**
11. **Oracle Streams:** (Condition DB and LFC at T1 depends on that, recons at T1 won't happen): **7**
12. **FTS at CERN:** (Replication of RAW to T1's, recons at T1 won't happen): **7**
13. **RB at CERN:** (Whichever processing/MC production activity on LCG sites won't work **but** other instances geographically distributed at T1): **5**
14. **ConditionDB:** Reprocessing of data. At CERN the master will necessarily have high criticality: **7**

Criticality of Services and (rough) motivation @ T1

1. **CE/WN at T1:** (Recons/stripping/analysis at single T1 won't happen): **5**
2. **SE at T1:** (Recons/stripping/analysis at single T1 won't happen but failover mechanism for redistribution/upload in place, modular share data distribution in place): **3**
3. **FTS at T1:** (DTS data redistribution won't happen): **3**
4. **RB/gLite WMSes** (despite RB at T1 were supposed for a redundant service, the outage of any RB at T1 turns into a global reduced effectiveness being the RB at CERN not able to support all the load): **3**
5. **LFC at T1:** (72 hours of downtime would not be acceptable at all) **3**

6. **ConditionDB at T1:** (The same downtime tolerance as the master at CERN – 8 hours – but this would imply a major reduced effectiveness rather than a serious disruption)
5
7. **VO-BOX at T1:** (DIRAC redundancy services, best effort): 1

5. METRICS and CURRENT STATUS

By their nature several LHCb specific services have a complex (hierarchical and/or distributed) nature and the metric definition is not always as trivial as like for a single and well identified service instance.

In this chapter I will try to propose some numbers and how do they can be elaborated as well as a status of the sensor available (or not) by putting a color in the name of the service. Some ancillary services (Oracle Streams, AFS, CASTOR and so forth) are not listed there. It means that the very top view Gridmap should provide is *only* made of the following services.

1. **LFC.** For LFC we have **both** SAM tests and SLS sensors running every hour. The SLS sensor https://sls.cern.ch/sls/service.php?id=LFC_LHCB_distributed reflects the distributed nature of the LFC for LHCb and also its hierarchical structure about various sub-services and dependencies involved. The service is composed by a Master and 7 read only instances that are Streams Replicated. The service is fully available (M=100) if all instances succeed the test described above. The service is affected (100<M<90) if some read-only is not fully available while the service is degraded for 90>M>50. The service is not available M<50. M is the weighted average of the 8 instances. The Master weights 50 while each other read-only instance weights 7. This roughly reflects the criticality assigned above. The same SLS sensor exposes information on other services like Physics Streams and the `lfc_lhcb` RAC (from DM), and the LFC Front End sensors maintained by FIO. This allows to better drill into problems in case some issue happens. *Grid Map should represent LFC as a very large case and should allow to expand this case into another map that reflects this hierarchy.*
2. **VOBOX:** We do have a SLS sensor that is a hierarchical service containing in cascade at least three layers of DIRAC specific services exploitable because of the DIRAC3 ping() method. <https://sls.cern.ch/sls/service.php?id=DIRAC-VOBOX> We assigned top criticality to the central VOBOX service and this comes because it collects all DIRAC services. It means that the overall metric (number that quantifies the status of this service) should consequently reflect the status of each service proportionally to its criticality. The VOBOX service offers a hierarchical view being composed by different setups (production/analysis) in turn composed by different systems (WMS,DMS and so forth) that are indeed composed by different service

instances (JobManager, TransferDB). I propose here the “flat” hypothesis: each component (at each level) gets a criticality equal to $(100/N) \%$ the total criticality (N =number of subcomponent at the level considered: ex. Number of “systems” in the specific “setup”). This same weight is used to compute the final metric M for the VOBOX. Correspondingly, at visualization level, all cases of the generated sub-Grid Map (referring to the map of all Dirac setups inside the VOBOX service) will get the same size. The same exact principle must be considered applied by clicking the case of a given “Dirac setup”; this will generate a sub-sub-Grid Map with all systems of the given setup (and so forth so on by clicking on each sub-cases of this last Map new GridMap will come). The atomic unit used for compute the overall metric of the VOBOX (and then the color of the corresponding Grid Map case) is the service instance bit that might be 1 or 0 and available via ping () method

3. **CE/WN (CERN and T1):** we made a clear distinction here between CERN CE and single T1 CEs because different activities are supposed to happen through. Here we do not see any complex structure for this service being already in place a distinction between T0CE and T1 CE. Graphically this can be translated into 7 different cases whose size is proportional to the criticality we gave above. For each CE/WN service we already have (historically) a fairly exhaustive SAM suite running smoothly every 12 hours. (<https://lcg-sam.cern.ch:8443/sam/sam.py?sensors=CE®ions=AsiaPacific®ions=CERN®ions=CentralEurope®ions=France®ions=GermanySwitzerland®ions=Italy®ions=NorthernEurope®ions=OpenScienceGrid®ions=Russia®ions=SouthEasternEurope®ions=SouthWesternEurope®ions=UKI®ions=USATLAS®ions=USCMS®ions=Unknown&vo=lhcb&order=SiteName&func=ShowSensorTests>) The Metric (M) here is: overall SAM test result (OK or ERROR). M might be then 100 or 0. We do not see the reason to have finer grained results. Unless we introduce here some weights on the different (tests) we currently run on our CEs and we build the final M for this service by aggregating all these results (opportune weighted). To be defined.
4. **SE/SRM at T1:** Several SRM sensors are in place in SAM but are not fully covering all aspects (because thoughts for SRM v1 while we need here rather greater complexity). LHCb refuses to distinguish SRM by SE or whichever service constituting a site Storage System. We generically refer to SRM by indicating (and hiding the complexity underneath with a convention) the whole system for stocking data. It means that we ignore the SAM SE sensor but we (we plan to) test through SRM (v2) all components exposed to the users. Currently in place a test checking the SRM front end status (whenever available accordingly the information from StagerDB) and a test accessing data using file access protocol.

sam.cern.ch:8443/sam/sam.py?sensors=SRM®ions=AsiaPacific®ions=CERN®ions=CentralEurope®ions=France®ions=GermanySwitzerland®ions=Italy®ions=NorthernEurope®ions=OpenScienceGrid®ions=Russia®ions=SouthEasternEurope®ions=SouthWesternEurope®ions=UKI®ions=USATLAS®ions=USCMS®ions=Unknown&vo=lhcb&order=SiteName&funct=ShowSensorTests . With SRMv2 and a bigger granularity introduced with the concept of Space Token we envisage to run unit tests for each space token and the overall metric (M) for each SRM endpoint (=site storage system) is the weighted average of the results of all these sub-service tests results. GridMap must offer 7 different SRM service (for 7 different sites) each offering a hierarchical view of provided space tokens defined.

5. **SE at CERN:** it must be dealt separately as different service than SEs at other centers. It's worthwhile to note that for CERN there are already SLS sensors that would help to debug eventual problems: (real) SRM sensor: https://sls.cern.ch/sls/service.php?id=CASTOR-SRM_LHCB service classes sensors: <https://sls.cern.ch/sls/service.php?id=CASTORLHCB> which in turns depend on CERN fabric sensors that should also be linked by Grid Map like the Oracle DB used by castor https://sls.cern.ch/sls/service.php?id=DB_castor https://sls.cern.ch/sls/service.php?id=DB_c2lhcbstgdb or the tape service <https://sls.cern.ch/sls/service.php?id=CASTORTapeService> or even AFS <https://sls.cern.ch/sls/service.php?id=AFS>
6. **FTS:** Straightforwardly FTS SAM tests have been put in place this week testing the capability of submitting (querying and cancelling) transfer jobs and the channels setup. <https://lcg-sam.cern.ch:8443/sam/sam.py?sensors=FTS®ions=AsiaPacific®ions=CERN®ions=CentralEurope®ions=France®ions=GermanySwitzerland®ions=Italy®ions=NorthernEurope®ions=OpenScienceGrid®ions=Russia®ions=SouthEasternEurope®ions=SouthWesternEurope®ions=UKI®ions=USATLAS®ions=USCMS®ions=Unknown&vo=lhcb&order=SiteName&funct>ShowSensorTests> This SAM test is fairly complete for evaluating whether a FTS service is OK and the metric in this case is simply 100 (available) or 0 (not available) for each independent site instance. Again Grid Map should expose 7 different cases for 7 different sites whose size is proportional to the criticality assigned above. For CERN there are also service specific metrics from SLS that might be also taken into account or linked to the T0 FTS case. <https://sls.cern.ch/sls/service.php?id=FTS>
7. **RB:** Like for FTS SAM we have simple tests probing basic features of the LCG RB (job submission/list match and job cancel) that exhaust the functionalities LHCb expects from such service. <https://lcg-sam.cern.ch:8443/sam/sam.py?sensors=RB®ions=AsiaPacific®ions=CERN®ions=CentralEurope®ions=France®ions=GermanySwitzerland®ions=Italy®ions=NorthernEurope®ions=OpenScienceGrid®ions=Russia®ions=SouthEasternEurope®ions=SouthWesternEurope®ions=UKI®ions=USATLAS®ions=USCMS®ions=Unknown&vo=lhcb&order=SiteName&funct>ShowSensorTests>

[ons=CERN®ions=CentralEurope®ions=France®ions=GermanySwitzerland®ions=Italy®ions=NorthernEurope®ions=OpenScienceGrid®ions=Russia®ions=SouthEasternEurope®ions=SouthWesternEurope®ions=UKI®ions=USATLAS®ions=USCMS®ions=Unknown&vo=lhcb&order=SiteName&funct=ShowSensorTests](https://sls.cern.ch/sls/service.php?id=RB-LHCb) The metric we suggest here is again 100 or 0 depending on the overall result of SAM tests for each RB. Grid Map should display 7 different instances of RB (each for any site providing one) and for CERN it should also allow for a hierarchical structure (having more instances available) and link also to SLS sensor also available <https://sls.cern.ch/sls/service.php?id=RB-LHCb>

8. **gRB** : for the gLite WMS we have exactly the same situation depicted for the RB service. SAM tests already available (checking the same functionalities as LCG RB) <https://lcg-sam.cern.ch:8443/sam/sam.py?sensors=gRB®ions=AsiaPacific®ions=CERN®ions=CentralEurope®ions=France®ions=GermanySwitzerland®ions=Italy®ions=NorthernEurope®ions=OpenScienceGrid®ions=Russia®ions=SouthEasternEurope®ions=SouthWesternEurope®ions=UKI®ions=USATLAS®ions=USCMS®ions=Unknown&vo=lhcb&order=SiteName&funct=ShowSensorTests> and also SLS sensor (for CERN instances) should be linked in Grid Map <https://sls.cern.ch/sls/service.php?id=WMS-LHCb>
9. **CERN network CAMPUS**: We just need to link the SLS sensor <https://lemonweb.cern.ch/sls/service.php?id=GeneralNetwork> into Grid Map . This service might be ancillary to almost all T0 higher level services.
10. **VOMS**: We don't have any SAM sensor yet though we have SLS lower level sensors for the backend https://sls.cern.ch/sls/service.php?id=phydb_lcg_voms . It has however been agreed that this is a common to all VOs service and then we could link our Grid Map to the dteam tests for VOMS. VOMS administrators are also looking at the possibility to implement some dedicated service sensor in SLS. Steve Traylen will provide something soon
11. **BK**: We have a SLS sensor for low level components for the backend https://sls.cern.ch/sls/service.php?id=phydb_lhcb_bookkeeping and another sensor (SLS) must be registered (but I can anticipate the info are obtainable from <https://sls.cern.ch/sls/service.php?id=LHCb-BK>)
12. **LOG SE**: for this service we have in place SLS sensors (result of a query to lemon for the status of the gridftp http and certificates) that are part of a generic sensor that needs to be imported from SDB to SLS. https://sls.cern.ch/sls/service.php?id=log_se_gridftp , https://sls.cern.ch/sls/service.php?id=log_se_httpd , and https://sls.cern.ch/sls/service.php?id=log_se_cert
13. **Condition DB**: The condition DB is mainly its underlying DB and Physics Streams. <https://sls.cern.ch/sls/service.php?id=PhysicsStreams> and https://sls.cern.ch/sls/service.php?id=phydb_LHCB . However

I've asked developers to think about some higher level tests that verifies the status of the tiny application on top of that. *Grid Map should be able to display a hierarchical structure in a similar way as for the LFC distributed service.*