

GS

Grid Support

CERN IT
Department

Critical services in the LHC Computing


Andrea Sciabà

CHEP 2009
Prague, Czech Republic
21-27 March, 2009


CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it



- Introduction
- Computing services
- Service readiness metrics
- Critical services by experiment
- Readiness status evaluation
- Conclusions




- All LHC experiments work on the WLCG infrastructure
 - Resources distributed over ~140 computing centres
 - Variety of middleware solutions from different Grid projects and providers
- WLCG
 - coordinates the interaction among experiments, resource providers and middleware developers
 - Enforces a certain level of service availability and reliability



A map of the worldwide LCG infrastructure operated by EGEE and OSG.

- Computing systems are built on a large number of **services**
 - Developed by the **experiments**
 - Developed (or supported) by **WLCG**
 - Common **infrastructure** site services
- WLCG needs to
 - Periodically evaluate the **readiness** of services
 - Have monitoring data to assess the service **availability** and **reliability**
 - ⇒ SAM, GridView



Experiment layer

VO data management

VO central catalogue

VO workload management

Middleware layer

FTS

xrootd

LFC

MyProxy

WMS

SRM

VOMS

CE

Fabric layer

CASTOR

Oracle

Batch

Infrastructure layer

AFS


CVS


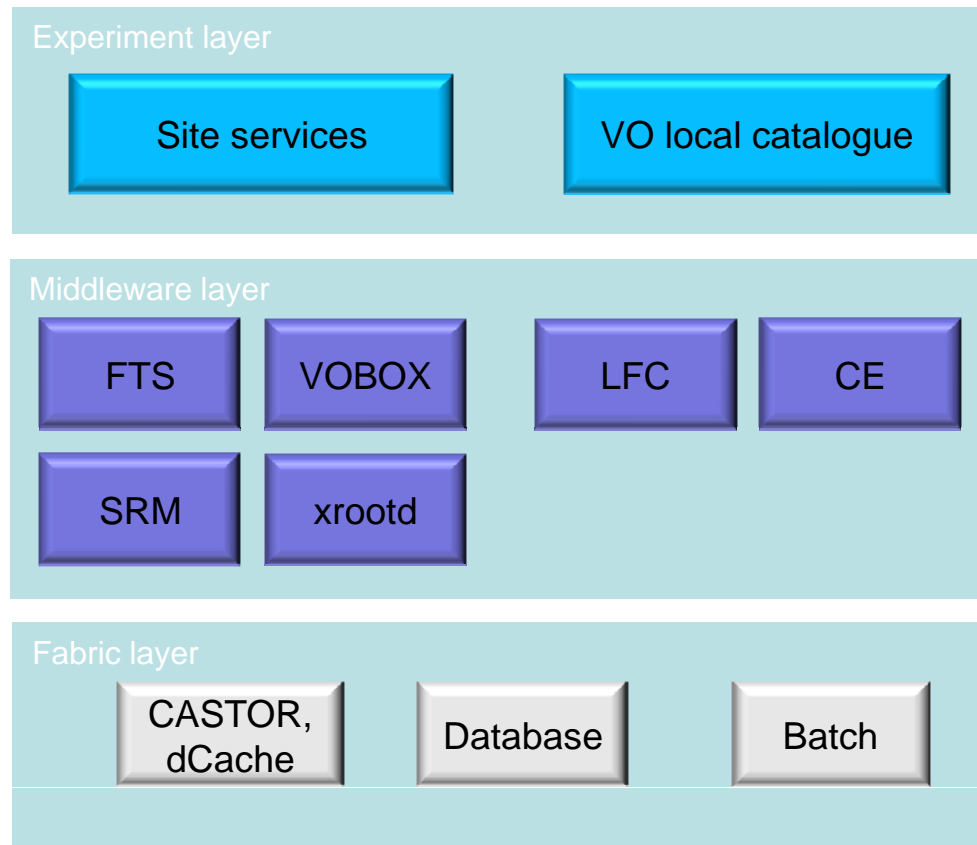
E-mail

Twiki

Web servers

Savannah





- **Service readiness** has been defined as a set of criteria related to several aspects concerning
 - Software
 - Service
 - Site
- Service **reliability** is **not** included and it is rather measured *a posteriori*

Software readiness

- High-level description of service available?
- Middleware dependencies and versions defined?
- Code released and packaged correctly?
- Certification process exists?
- Admin Guides available?


Service readiness

- Disk, CPU, Database, Network requirements defined?
- Monitoring criteria described?
- Problem determination procedure documented?
- Support chain defined (2nd/3rd level)?
- Backup/restore procedure defined?

Site readiness

- Suitable hardware used?
- Monitoring implemented?
- Test environment exists?
- Problem determination procedure implemented?
- Automatic configuration implemented?
- Backup procedures implemented and tested?

- Each experiment has provided a list of “critical” services
 - Rated from 1 to 10
- A survey has been conducted on the critical tests to rate their readiness
 - To see where more effort is required




Critical services	Rank	Comment
AliEN	10	ALICE computing framework
Site VO boxes	10	Site becomes unusable if down
CASTOR and xrootd at Tier-0	10	Stops 1 st pass reco (24 hours buffer)
Mass storage at Tier-1	5	Downtime does not prevent data access
File Transfer Service at Tier-0	7	Stops 2 nd pass reco
gLite workload management	5	Redundant
PROOF at Tier-0 CERN Analysis Facility	5	User analysis stops

Rank 10: **critical**, max downtime 2 hours

Rank 7: **serious disruption**, max downtime 5 hours

Rank 5: **reduced efficiency**, max downtime 12 hours




10: Very high	interruption of these services affects online data-taking operations or stops any offline operations
7: High	interruption of these services perturbs seriously offline computing operations
4: Moderate	interruption of these services perturbs software development and part of computing operations

Rank	Services at Tier-0
Very high	Oracle (online), DDM central catalogues , Tier-0 LFC
High	Cavern→T0 transfers, online-offline DB connectivity, CASTOR internal data movement, Tier-0 CE, Oracle (offline), Tier-0/1 FTS, VOMS, Tier-1 LFC, Dashboard, Panda/Bamboo , DDM site services/VO boxes
Moderate	3D streaming, gLite WMS, Tier-0/1/2 SRM/SE, Tier-1/2 CE, CAF, CVS, Subversion, AFS, build system, Tag Collector

Rank	Services at Tier-1
High	LFC, FTS, Oracle
Moderate	3D streaming, SRM/SE, CE

Rank	Services at Tier-2
Moderate	SRM/SE, CE

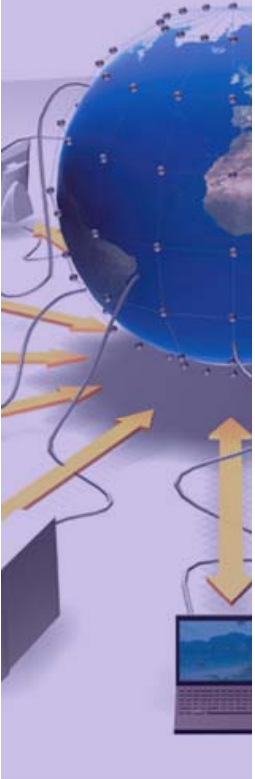
Rank	Services elsewhere
High	AMI database



CMS gives a special meaning to all rank values

Rank	Definition	Max. downtime per incident (Hrs)	Comment
11	CMS Stops operating	0.5	Not covered (yet) here
10	CMS stops transferring data form Cessy		Cessy output buffer time
9	T0 Production stops		min(T0 input buffer/CESSY output buffer) or defined time to catch up
8	T1/T2 Production/analysis stops		defined time to catch up
7	Services critical when needed but not needed all the time (currently includes documentation)	0.5	
6	A service monitoring or documenting a critical service	8	
5	CMS development stops if service unavailable	24	
4	CMS development at CERN stops if service unavailable	24	
3	Services not critical for CMS	24	
2	Services required for CMS	72	
1	Used by a significant fraction of CMS	72	
0	Not used or discouraged by CMS	forever	

Rank 10: 24x7 on call
Rank 8,9: expert call-out




Rank	Services
10	Oracle, CERN SRM-CASTOR, DBS , Batch, Kerberos, Cavern-T0 transfer+processing
9	CERN FTS, PhEDEx , FroNTier launchpad , AFS, CAF
8	gLite WMS, VOMS, Myproxy, BDII, WAN, Non-T0 prod tools
7	APT servers, build machines, Tag collector , testbed machines, CMS web server, Twiki
6	SAM, Dashboard, PhEDEx monitoring , Lemon
5	WebTools , e-mail, Hypernews, Savannah, CVS server
4	Linux repository, phone conferencing, valgrind machines
3	Benchmarking machines, Indico

Rank ▲	Definition ▲	Max downtime (hrs) ▲	Comment ▲
10	Critical	0.5	
7	Serious disruption	8	
5	Major reduction in effectiveness	8	
3	Reduced effectiveness	24	
1	not critical	72	


Rank	Services
10	Tier-0 CASTOR, AFS, CERN VO boxes (DIRAC3 central services) , Tier-0 LFC master, Oracle at CERN
7	VOMS, CERN FTS, Tier-0 ConditionDB, LHCb bookkeeping service , Oracle Streams, SAM
5	Tier-0/1 CE and batch, Tier-0 gLite WMS, Tier-1 ConditionDB
3	Tier-1 SE, Tier-0/1 Replica LFC, Dashboard, Tier-1 FTS, Tier-1 gLite WMS
1	Tier-1 VO boxes

Service	Readiness
Data services	
CASTOR, SRM, FTS, LFC	100%
Oracle	No piquet service
Computing Services	
CE, batch services	No expert piquet service
gLite WMS + LB	Insufficient monitoring and problem detection, no expert piquet service, no backup
Other Grid services	
MyProxy	Procedures not fully documented, no expert piquet service
VOMS	Problem determination procedure does not cover everything, no expert piquet service
BDII	Some documentation slightly outdated, no expert piquet service
VOBOX	No expert piquet service, only sysadmin piquet , but OK for VO boxes
Dashboard	No certification process, no automatic configuration
SAM	100%
Other non-Grid services	
AFS, Kerberos	No certification process at CERN, problem determination procedure not documented, no test environment
Twiki	Relies heavily on AFS (svc, backend data, backups)



Service	Readiness
AliEn	100%
VO BOX for ALICE	Admin guide for 64-bit setup provided by ALICE as HowTo General setup and problem determination/solution procedures provided by ALICE Support methods, technical contacts and procedures defined at the T0/T1s through SLAs
Xrootd	100%
PROOF (ALICE CAF)	Support via mailing list


Service	Readiness
DDM central catalogues	Problem determination by shiffters , solving by experts Configuration must be done by an expert
Panda / Bamboo	Problem determination mostly by experts No automatic configuration
DDM site services	Certification process in place , but “preproduction” instance is a subset of the production system Hardware requirements ok for central activities, unknown for analysis Backup via Oracle Monitoring via SLS
ATLAS Metadata Interface	Lacks a rigid certification process, but has test infrastructure No admin guide Backups via Oracle Streams No special procedure for problem determination Support chain and monitoring being improved



Service	Readiness
DBS	100%
Tier-0	No admin guide, no certification, constantly evolving
PhEDEx	Problem determination procedure not complete but improving
FroNTier	100%
Production tools	Monitoring for non-experts should be improved Status of ongoing workflows not backed up
WebTools	Monitoring should be improved (via Lemon, SLS)
VOBox	100%

- All **DM/WM** services are in production since > 2 years (sometimes > 4 years)
- **Tier-0** services require **fast updates** ⇒ processes **less formal**
- Most services have **documented procedures** for installation, configuration, start-up, testing. They all have **contacts**
- Backups: managed by CERN IT for what is needed, everything else is transient

Service	Readiness
DIRAC central services	Working on admin guides No documented procedures for problem solving
T1 VOBOX services	Working on admin guides No documented procedures for problem solving
Bookkeeping	Insufficient documentation Problem determination procedure not documented (expert required)




- **24x7 support** in place everywhere
- **VO boxes** support defined by Service Level Agreements in majority of cases

15-Jan-09		WLCG High Level Milestones – 2008/09														
ID	Date	Milestone	Done (green)				Late < 1 month (orange)				Late > 1 month (red)					
			ASGC	CC IN2P3	CERN	DE-KIT	INFN CNAF	NDGF	PIC	RAL	SARA NIKHEF	TRIUMF	BNL	FNAL		
VOBoxes Support																
WLCG-07-04	Apr 2007	VOBoxes SLA Defined Sites propose and agree with the VO the level of support (upgrade, backup, restore, etc) of VOBoxes	Aug 2008	Aug 2008						Aug 2008			Jan 2009			
WLCG-07-05	May 2007	VOBoxes SLA Implemented VOBoxes service implemented at the site according to the SLA	Aug 2008	Aug 2008					Jan 2009	Aug 2008			Jan 2009			
WLCG-07-05b	Jul 2007	VOBoxes Support Accepted by the Experiments VOBoxes support level agreed by the experiments	ALICE	n/a												
			ATLAS							n/a	n/a					
			CMS							n/a				n/a	n/a	n/a
			LHCb	n/a						n/a					n/a	n/a

Almost final Final

- Almost all **CERN** central services are **fully ready**
 - A few concerns about WMS monitoring
 - Several services only “best effort” outside working hours (**no expert piquet service**, but sysadmin piquet)
 - Documentation of some Grid services not fully satisfactory
- ALICE is ok!
- ATLAS ok but relying on experts for problem solving and configuration; analysis impact is largely unknown
- CMS services basically ready
 - Tier-0 services (data processing): development/deployment process very rapid, but working satisfactorily
- LHCb has shortcomings in procedures, documentation
 - But note that DIRAC3 is still relatively new



- Service reliability is not taken in consideration here
 - Services “fully ready” might actually be rather fragile; depends on the deployment strategies
- Overall, all critical services are in a **good shape**
 - **No showstoppers** of any kind identified
 - To fix the few remaining issues by the restart of the data taking
- The computing systems of the LHC experiments rely on **services mature from the point of view of documentation, deployment and support**

