


# Status of critical services

Andrea Sciabà

WLCG Data taking readiness  
planning workshop  
November 13-14, 2008  
CERN

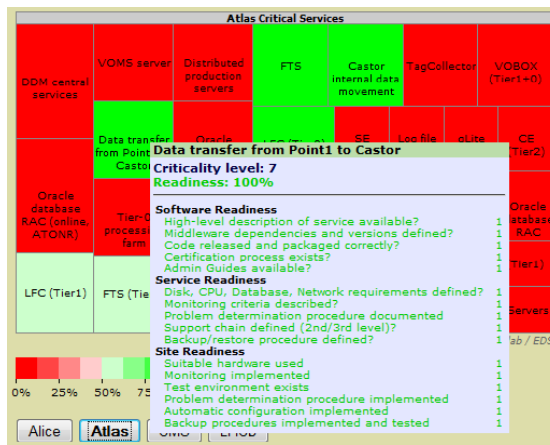
- Review the critical services lists
  - Experiment services
  - WLCG services
  - CERN services
- Assess the readiness of the services
  - Criteria defined by parameters
    - Software, service, site readiness
    - Questions addressed to the service managers
- Identify areas needing further effort



Questions chosen to determine the level of “readiness”


Answers will eventually be published in the service Gridmaps

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?



Reliability, redundancy etc. not part of this definition!





### Critical services

AliEN

WMS


Tier-0 FTS

Tier-0 + Tier-1 SRM (via FTS)


Tier-0 CASTOR + xrootd

Tier-1 dCache/CASTOR + xrootd

PROOF at Tier-0 CAF



Very high	interruption of these services <u>affects online data-taking operations or stops any offline operations</u>
High	interruption of these services <u>perturbs seriously offline computing operations</u>
Moderate	interruption of these services perturbs software development and part of computing operations



Rank	Services at Tier-0
Very high	Oracle (online), <b>DDM central catalogues</b>
High	P1→T0 transfers, online-offline DB connectivity, CASTOR internal data movement, T0 processing farm, Oracle (offline), LFC, FTS, VOMS, Dashboard, <b>Panda/Bamboo</b> , <b>DDM site services</b>
Moderate	3D streaming, WMS, SRM/SE, CAF, CVS, AFS, build system

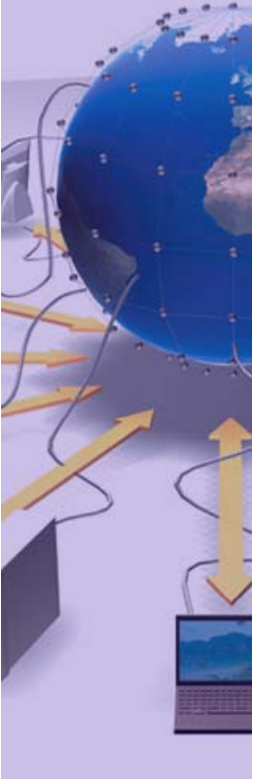
Rank	Services at Tier-1
High	LFC, FTS
Moderate	3D streaming, Oracle, SRM/SE, CE

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

Rank	Services elsewhere
High	<b>AMI database</b>

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, LXBATCH, Kerberos, Cessy-T0 transfer+processing, Web “back-ends”
9	CERN FTS, PhEDEx, FroNTier launchpad, AFS, CAF
8	WMS, VOMS, Myproxy, BDII, WAN, ProdMgr
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	<b>CERN VOBoxes (DIRAC3 central services)</b> , Tier-0 LFC, VOMS
7	Tier-0 SE, <b>T1 VOBoxes</b> , SE access from WN, FTS, WN misconfiguration, CE, Conditions DB, <b>LHCb bookkeeping service</b> , Oracle streaming, SAM
5	RB/WMS
3	T1 LFC, Dashboard

Service	Readiness
<b>Data services</b>	
CASTOR, SRM, FTS, LFC	Admin guide for SRM under development
Oracle	No piquet service
<b>Computing Services</b>	
CE, batch services	No expert piquet service, some scalability issues (when no. jobs > 50K)
WMS+LB	Insufficient monitoring and problem detection, no expert piquet service, no backup, doubts on needed resources
<b>Other Grid services</b>	
MyProxy	Procedures not fully documented, no expert piquet service
VOMS	Support chain not fully defined (no expert piquet service), problem determination procedure not fully implemented
BDII	Some documentation obsolete, no expert piquet service
VOBOX	No expert piquet service, only sysadmin piquet, but OK for VOBOXes
Dashboard	No certification process, no automatic configuration
SAM	?
<b>Other non-Grid services</b>	
Lemon	?
AFS, Kerberos	No certification process at CERN, problem determination procedure not documented, no test environment, no automatic configuration
Twiki	Relies heavily on AFS (svc, backend data, backups)



Service	Readiness
AliEn	100%
VOBOX for ALICE	Middleware dependencies from gLite UI not fully defined Admin guide not updated to 64-bit setup Problem determination procedure provided by ALICE Support chain not clearly defined
Xrootd	100%
PROOF	Support via mailing list only No automatic configuration No backup (but configuration files on AFS)



Service	Readiness
DDM central catalogues	Software dependencies being clarified Problem determination mostly by experts Configuration must be done by an expert
Panda / Bamboo	No full description, no admin guide Migrating to SVN Certification needs improving Problem determination mostly by experts No automatic configuration Backup needs developers involvement
DDM site services	Software dependencies being clarified Certification problem to be improved (now done "in production") HW requirements ok for central activities, unknown for analysis Backup via Oracle Monitoring via SLS

Service	Readiness
DBS	Certification needs improvement
Tier-0	No admin guide, no certification, constantly evolving
PhEDEx	Problem determination procedure not complete but improving
FroNTier	100%
Production tools	Problem determination procedure not complete but improving
WebTools	Monitoring could be improved (via Lemon, SLS)
VOBox	No automatic configuration

- All DM/WM services are in production since > 1 year (sometimes > 4 years)
- Most services have documented procedures for installation, configuration, startup, testing. They all have contacts
- MC production: contacts and procedures defined, services managed by DataOps
- CERN IT manages the VOBOXes and can contact the CMS service managers via dedicated mailing lists
- Backups: managed by IT for what is needed, everything else is transitory

Service	Readiness
DIRAC central services	No high level description No documented procedures for problem solving
T1 VOBOX services	No test environment No documented procedures for problem solving
Bookkeeping	No high level description, no middleware dependencies No monitoring, no automatic configuration


Note that DIRAC3 is a relatively new system!

- 24x7 support in place everywhere
- VO boxes support defined by SLAs in majority of cases

7-Nov-08		WLCG High Level Milestones - 2007													
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 NIKHEK	TRIUMF	BNL	FNAL	
<b>24x7 Support</b>															
WLCG-07-01	Feb 2007	<b>24x7 Support Definition</b> Definition of the levels of support and rules to follow, depending on the issue/alarm													
WLCG-07-02	Apr 2007	<b>24x7 Support Tested</b> Support and operation scenarios tested via realistic alarms and situations				Apr 2008	June 2008								
WLCG-07-03	Jun 2007	<b>24x7 Support in Operations</b> The sites provides 24x7 support to users as standard operations				July 2008	June 2008		Apr 2008		July 2008				
<b>VOBoxes Support</b>															
WLCG-07-04	Apr 2007	<b>VOBoxes SLA Defined</b> Sites propose and agree with the VO the level of support (upgrade, backup, restore, etc) of VOBoxes	Aug 2008	Aug 2008					Aug 2008						
WLCG-07-05	May 2007	<b>VOBoxes SLA Implemented</b> VOBoxes service implemented at the site according to the SLA	Aug 2008	Aug 2008				Mar 2008	Aug 2008		Apr 2008				
WLCG-07-05b	Jul 2007	<b>VOBoxes Support Accepted by the Experiments</b> VOBoxes support level agreed by the experiments	ALICE	n/a						n/a			n/a	n/a	n/a
			ATLAS						n/a	n/a					n/a
			CMS						n/a			n/a	n/a	n/a	
			LHCb	n/a					n/a				n/a	n/a	n/a


Almost final      Final

- Almost all CERN central services seem to be 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 a lot on experts for problem solving and configuration; analysis impact is largely unknown; Panda needs work (being done with migration to CERN)
- CMS services basically ready
  - Tier-0 services (data processing) still too “fluid”, being debugged in production during cosmic runs
- LHCb a bit short of procedures documentation
  - But note that DIRAC3 was put in production very recently






- Service reliability is not taken in consideration here
  - See yesterday's talks by John and Miguel
  - Services “fully ready” might actually be rather fragile; depends on the deployment strategies
  - Should add fault tolerance, redundancy, etc. to the readiness criteria?
- Overall, all critical services are in a good shape
  - No showstoppers of any kind identified
  - Not surprising, given that in principle we should be taking data!
  - The goal is clearly to fix the remaining issues by the restart of the data taking




- Thanks to all those who gave me the input for this talk!
  - J. Andreeva, L. Betev, P. Bittencourt, S. Campana, M. Coelho Dos Santos, M. Girone, P. Jones, B. Köblitz, P. Kreuzer, M. Lamanna, S. Lemaitre, M. Litmaath, Z. Mathe, G. McCance, P. Méndez, M. Meoni, R. Nandakumar, E. Roche, S. Metson, P. Saiz, R. Santinelli, M. Seco, U. Schwickerath, R. Silva, R. Toebbicke, S. Traylen, J. Van Eldik




# Backup slides


Detailed answers




- BDII (R. Dominques da Silva)
  - Software stable, some install guides outdated
  - HW requirements understood, monitoring criteria defined, ops/sysadmin procedures in place, support chain well defined
  - Fully quattorized installation, HW adequate
  - Expert coverage outside working hours is best effort
- CE (U. Schwickerath)
  - Reinstallation procedure defined, log files backed up
  - Reliable nodes, test environment in PPS, automatic configuration via CDB, remote syslog and TSM
  - Expert coverage outside working hours is best effort
- LXBATCH (U. Schwickerath)
  - Service description pages need updating, preprod system for Linux updates
  - Batch system status held in LSF master, scalability issues to be addressed when >50K jobs
  - Separate LSF test instance
  - Expert coverage outside working hours is best effort




- WMS + LB (E. Roche)
  - No real admin guide but plenty of docs exist
  - Doubts about the VO load estimates – HW adequate for them
  - Problem determination procedures being documented with experience
  - No backup procedure – all services stateful
  - Monitoring only partial
  - Expert coverage outside working hours is best effort
- MyProxy (M. Litmaath)
  - Docs on home page and gLite User Guide
  - HW defined at CERN
  - partial documentation of procedures
  - A testing procedure exists
- VOMS (S. Traylen)
  - No database reconnects (next version)
  - Problem determination does not cover everything
  - No out of hours coverage




- Dashboard (J. Andreeva)
  - No MW dependencies
  - No certification process
  - HW requirements not fully defined
  - Problem determination procedures not fully defined



- AliEn (P. Saiz)
  - All monitoring via MonaLisa
  - Support goes via mailing list
  - Suitable HW used at most sites
- VOBOX (P. Méndez)
  - No user or admin guide
  - Had some problems with the UI on 64-bit nodes
  - No admin guide for 64-bit installations
  - Problem determination procedure written by ALICE
  - Automatic configuration via YAIM




- xrootd (L. Betev)
  - Documentation on WWW and Savannah
  - MW dependencies on: dCache (emulation), DPM & CASTOR2, (plug-in), xrootd (native)
  - xroot releases: RPM, tarball, in-place compilation; other interfaces: packaged with corresponding storage solution code
  - Certification depends on storage solution
  - Admin guides by ALICE as How-to's
  - Monitoring: specific to storage solution; via MonaLISA for SE functionality and traffic
  - Problem determination procedures for all implementations
  - Support done via site support structure
- PROOF (L. Betev, M. Meoni)
  - SW description with dependencies and Admin Guide at
    - <http://root.cern.ch/twiki/bin/view/ROOT/PROOF>
    - <http://aliceinfo/Offline/Activities/Analysis/CAF/index.html>
  - Code managed via SVN and distributed as RPMs
  - There is a developer partition for certification
  - HW provided by IT: 15 nodes, 45 TB of disk, 120 cores, no DB, 1 Gb/s links
  - Monitoring via MonaLISA and Lemon
  - FAQ at <http://aliceinfo/Offline/Activities/Analysis/CAF/index.html>
  - No person on call, support via mailing list
  - Critical configuration files on AFS
  - No automatic configuration, no backup







- DDM Central Catalogues (B. Köblitz)
  - Backups via Oracle
  - Problem determination and procedures via experts
- DDM site services (S. Campana)
  - Certification is sort of done in a production environment
  - Hardware requirements are fully satisfied for central activities but not necessarily for analysis (difficult to quantify at the moment)



- DM/WM tools (S. Metson)
  - PhEDEx, CRAB, ProdAgent not heavily tied to any MW version
  - All software is packaged using RPM
  - All tools are properly tested, but
  - Tier-0 tools constantly evolving (and no Admin Guide)
  - DBS certification needs improving
  - Core monitoring and problem determination procedures are there, could improve
  - All that needs back-ups is on Oracle



- VOBOX (P. Bittencourt)
  - Full documentation on Twiki (including Admin Guide), dependencies always defined, using CVS and RPM
  - Certification testbed + preprod testbed for all WebTools services
    - Services tested on VM running on testbed nodes
  - Extending Lemon usage for alarms
  - Some services have operators for maintenance and support, with testing procedures
  - Requirements specified for architecture (64-bit), disk space, memory
  - No automatic configuration
  - Support goes via Patricia and – if needed – to [vobox.support@cern.ch](mailto:vobox.support@cern.ch)



- DIRAC 3 central services (M. Seco, R. Santinelli)
  - Only access procedures documented, no admin guide
  - Problem determination procedure not fully documented
  - Backup procedures defined in SLA
  - DIRAC3 test environment
- DIRAC3 site services (R. Nandakumar, R. Santinelli)
  - No admin guide
  - No test environment
- Bookkeeping (Z. Mathe, R. Santinelli)
  - Support chain defined: 1) user support, 2) Grid expert on duty or DIRAC developers, 3) BK expert
  - Service runs on a mid-range machine and has an integration instance for testing
  - Back-end is on Oracle LHCb RAC
  - Is monitored

