

WLCG Service Availability

Talk Overview

- MoU Resource and Availability Targets
- Services per Tier and WLCG Service definitions
- Service Availability Monitoring (SAM)

'The' MoU – what is it

- CERN resource agreements with collaborating experiments and sites are formally described in a 'Memorandum of Understanding' that both parties sign.
- The one we refer to is that 'for Collaboration in the Deployment and Exploitation of the Worldwide LHC Computing Grid'
- Last version dated 1 June 2006 is CERN-C-RRB-2005-01/Rev. at <http://lcg.web.cern.ch/LCG/documents.html>

Tier 0 Resource plan from MoU

Annex 6.1. Tier0 Computing Capacities

CERN Tier0	Pledged	Planned to be pledged				Comment
	2006	2007	2008	2009	2010	
CPU (kSI2K)	2400	4800	12500	15900	26200	
Disk (Tbytes)	230	450	1300	1300	1800	
Tape (Tbytes)	1500	3400	13600	23600	33900	
Nominal WAN (Mbits/sec)	80000	100000	120000	140000	160000	

Annex 6.2. CERN Analysis Facility Computing Capacities

CERN Analysis Facility	Pledged	Planned to be pledged				Comment
	2006	2007	2008	2009	2010	
CPU (kSI2K)	1000	4320	10000	14600	15000	
Disk (Tbytes)	540	1600	4200	5200	5300	
Tape (Tbytes)		1480	3400	5700	5900	

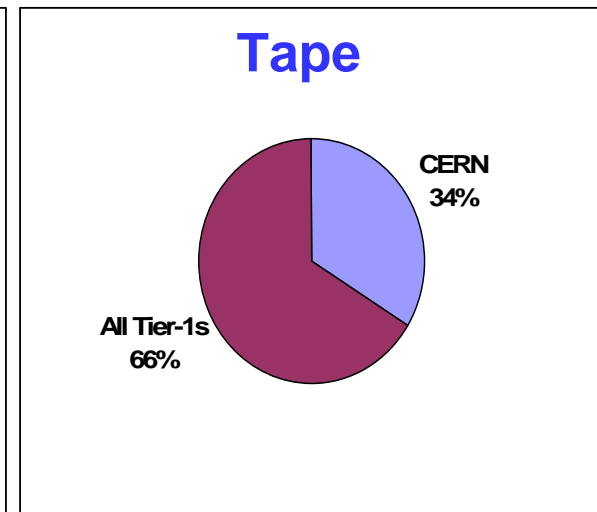


Similar tables exist for all Tier 1 and Tier 2 sites.

Summary of Computing Resource Requirements

All experiments - 2008

From LCG TDR - June 2005

	<i>CERN</i>	<i>All Tier-1s</i>	<i>All Tier-2s</i>	<i>Total</i>
CPU (MSPECint2000s)	25	56	61	142
Disk (PetaBytes)	7	31	19	57
Tape (PetaBytes)	18	35		53



The MoU Service Targets

- These define the (high level) services that must be provided by the different Tiers
- They also define average availability targets and intervention / resolution times for downtime & degradation
- These differ from Tier to Tier (less stringent as N increases) but refer to the 'compound services', such as "acceptance of raw data from the Tier0 during accelerator operation"
- Thus they depend on the availability of specific components – managed storage, reliable file transfer service, database services, ...
- An objective is to measure the services we deliver against the MoU targets
 - Data transfer rates
 - Service availability and time to resolve problems
 - Resources available at a site (as well as measured usages)
 - Resources specified in the MoU are cpu capacity (in KSi2K), tbytes of disk storage, tbytes of tape storage and nominal WAN data rates (Mbytes/sec)

CERN (Tier0) MoU Commitments

Service	Maximum delay in responding to operational problems			Average availability ^[1] on an annual basis	
	DOWN	Degradation > 50%	Degradation > 20%	BEAM ON	BEAM OFF
Raw data recording	4 hours	6 hours	6 hours	99%	n/a
Event reconstruction/data distribution (beam ON)	6 hours	6 hours	12 hours	99%	n/a
Networking service to Tier-1 Centres (beam ON)	6 hours	6 hours	12 hours	99%	n/a
All other Tier-0 services	12 hours	24 hours	48 hours	98%	98%
All other services ^[2] - prime service hours ^[3]	1 hour	1 hour	4 hours	98%	98%
All other services - outside prime service hours	12 hours	24 hours	48 hours	97%	97%

Tier 1 MoU commitments

Service	Maximum delay in responding to operational problems			Average availability measured on an annual basis	
	Service stoppage	Degradation ... by more than 50%	Degradation ... by more than 20%	During accelerator operation	At all other times
Acceptance of data from the Tier-0 Centre during accelerator operation	12 hours	12 hours	24 hours	99%	n/a
Networking service to the Tier-0 Centre during accelerator operation	12 hours	24 hours	48 hours	98%	n/a
Data-intensive analysis services, including networking to Tier-0, Tier-1 Centres outside accelerator operation	24 hours	48 hours	48 hours	n/a	98%
All other services – prime service hours	2 hour	2 hour	4 hours	98%	98%
All other services – outside prime service hours	24 hours	48 hours	48 hours	97%	97%

Some of these imply weekend/overnight staff presence or at least availability.

Availability= time running/scheduled up-time

Prime time= 08:00-18:00 weekday in time zone of host laboratory

Problem Response Time and Availability targets Tier-2 Centres

Service	<i>Maximum delay in responding to operational problems</i>		<i>availability</i>
	<i>Prime time</i>	<i>Other periods</i>	
End-user analysis facility	2 hours	72 hours	95%
Other services	12 hours	72 hours	95%

No requirement for weekend/overnight staff presence

Nominal MoU pp running data rates CERN to Tier 1 per VO

<i>Centre</i>	<i>ALICE</i>	<i>ATLAS</i>	<i>CMS</i>	<i>LHCb</i>	<i>Rate into T1 (pp) MB/s</i>
ASGC, Taipei	-	8%	10%	-	100
CNAF, Italy	7%	7%	13%	11%	200
PIC, Spain	-	5%	5%	6.5%	100
IN2P3, Lyon	9%	13%	10%	27%	200
GridKA, Germany	20%	10%	8%	10%	200
RAL, UK	-	7%	3%	15%	150
BNL, USA	-	22%	-	-	200
FNAL, USA	-	-	28%	-	200
TRIUMF, Canada	-	4%	-	-	50
NIKHEF/SARA, NL	3%	13%	-	23%	150
Nordic Data Grid Facility	6%	6%	-	-	50
Totals	-	-	-	-	1,600

These rates must be sustained to tape 24 hours a day, 100 days a year.

Extra capacity is required to cater for backlogs / peaks.

WLCG Internal Component Service Level Definitions and Availability Targets

Class	Description	Downtime	Reduced	Degraded	Availability
C	Critical	1 hour	1 hour	4 hours	99%
H	High	4 hours	6 hours	6 hours	99%
M	Medium	6 hours	6 hours	12 hours	99%
L	Low	12 hours	24 hours	48 hours	98%
U	Unmanaged	None	None	None	None

- Reduced defines the time between the start of the problem and the restoration of a reduced capacity service (i.e. >50%)
- Degraded defines the time between the start of the problem and the restoration of a degraded capacity service (i.e. >80%)
- Downtime defines the time between the start of a problem and restoration of service at minimal capacity (i.e. basic function but capacity < 50%)
- Availability defines the sum of the time that the service is down compared with the total time during the calendar period for the service. Site wide failures are not considered as part of the availability calculations. 99% means a service can be down up to 3.6 days a year in total. 98% means up to a week in total.
- None means the service is running unattended

Tier0 Services

Service	VOs	Class
SRM 2.1	All VOs	C
LFC global copy	LHCb	C
LFC local copy	ALICE, ATLAS	H
FTS	ALICE, ATLAS, LHCb, (CMS)	H
CE	All VOs	C
RB		C
Global BDII		C
Site BDII		H
Myproxy		C
VOMS		H→C
R-GMA		H

Required Tier1 Services

Service	VOs	Class
SRM 2.1	All VOs	H/M
LFC	ALICE, ATLAS	H/M
FTS	ALICE, ATLAS, LHCb, (CMS)	H/M
CE		H/M
Site BDII		H/M
R-GMA		H/M

Many also run e.g. an RB etc.

Required Tier2 Services

Service	VOs	Class
SRM 2.1	All VOs	H/M
LFC	ALICE	H/M
CE		H/M
Site BDII		H/M

To be checked with individual experiments

WLCG Service Monitoring framework

- Service Availability Monitoring (SAM) - uniform platform for monitoring all core services based on SFT (Site Functional Tests) experience. Other data sources can be used.
- Three main end users (and use cases):
 - project management - overall metrics
 - Operators (Local and Core infrastructure Centres) - alarms, detailed information for debugging, problem tracking
 - VO administrators - VO specific SFT tests, VO resource usage
- A lot of work already done:
 - SFT and GStat are monitoring CEs and Site-BDIs and will feed data into SAM
 - GRIDVIEW is monitoring data transfer rates and job statistics
 - A SAM Data schema has been established. There is a common Oracle database with that of GRIDVIEW
 - Basic displays in place (SFT reports, CIC-on-duty dashboard, GStat) (but with inconsistent presentation and points of access)
 - Basic framework for metric visualization in GRIDVIEW is ready

Current Site Functional Tests (SFT)

- Submits a short-living batch job to all sites to test various aspects of functionality: site CE uses local WMS to schedule them on an available local worker node
 - **Security tests:** CA certificates, CRLs, ... (work in progress)
 - **Data management tests:** basic operations using *lcg-utils* on the default Storage Element (SE) and chosen “central” SE (usually at CERN - 3rd party replication tests)
 - **Basic environment tests:** software version, *BrokerInfo*, CSH scripts
 - **VO environment tests:** tag management, software installation directory + VO specific job submission and tests (maintained by VOs, example: LHCb dirac environment test)
 - **Job submission - CE availability:** *Can I submit a job and retrieve the output?*
- What is not covered?
 - Only CEs and batch farms are really tested - other services are tested **indirectly**: SEs, RB, top-level BDII, RLS/LFC, R-GMA registry/schema tested by using default servers whereas sites may have several
- Maintained by Operations Team at CERN (+several external developers) and run as a cron job on CERN LXPLUS (using AFS based LCG UI)
- Jobs submitted at least every 3 hours (+ on demand resubmissions)
- Tests have a criticality attribute. These can be defined by VOs for their tests using the FCR (freedom of choice for resources) tool. The production SFT levels are set by the dteam VO. Attributes are stored in the SAM database.

Description of current SFT tests (1 of 2)

- If a test is defined as critical the failure causes the site to be marked as bad on the SFT report page. If the test is not critical only a warning message is displayed.
- **Job Submission**
 - **Symbolic name:** sft-job **Critical:** YES
 - **Description:** The result of test job submission and output retrieval. Succeeds only if the job finished successfully and the output was retrieved.
- **WN hostname**
 - **Symbolic name:** sft-wn **Critical:** NO
 - **Description:** Host name of the Worker Node where the test job was executed. This test is only for information.
- **Software Version**
 - **Symbolic name:** sft-softver **Critical:** YES
 - **Description:** Detect the version of software which is really installed on the WN. To detect the version *lcg-version* command is used and if the command is not available (very old versions of LCG) the test script checks only the version number of GFAL-client RPM.
- **CA RPMs Version (Certificate Authority RPMs)**
 - **Symbolic name:** sft-caver **Critical:** YES
 - **Description:** Check the version of CA RPMs which are installed on the WN and compare them with the reference ones. If for any reason RPM check fails (other installation method) fall back to physical files test (MD5 checksum comparison for all CE certs with the reference list).

Description of current SFT Tests (2 of 2)

- **. BrokerInfo**
 - **Symbolic name:** sft-brokerinfo **Critical:** YES
 - **Description:** Try to execute *edg-brokerinfo -v getCE*.
- **R-GMA client** (Relational Grid Monitoring Architecture)
 - **Symbolic name:** sft-rgma **Critical:** NO
 - **Description:** Test R-GMA client software configuration by executing *rgma-client-check* utility script.
- **CSH test**
 - **Symbolic name:** sft-csh **Critical:** YES
 - **Description:** Try to create and execute a very simple CSH script which dumps environment to a file. Fails if CSH script is unable to execute and the dump file is missing.
- **Apel test** (Accounting log parser and publisher)
 - **Symbolic name:** sft-apel **Critical:** NO
 - **Description:** Check if Apel is publishing accounting data for the site by using the command: *rgma -c "select count(*) from LcgRecords"*

SFT Tests can be aggregates

- **Test:** Replication management using LCG tools
- **Test symbolic name:** sft-lcg-rm **Critical:** YES
- **Description:** This is a super-test that succeeds only if all of the following tests succeed.
- GFAL infosys
 - **Description:** Check if *LCG_GFAL_INFOSYS* variable is set and if the top-level BDII can be reached and queried.
- lcg-cr to defaultSE
 - **Description:** Copy and register a short text file to the default SE using *lcg-cr* commamnd.
- lcg-cp defaultSE -> WN
 - **Description:** Copy the file registered in test 8-2 back to the WN using *lcg-cp* commamnd.
- lcg-rep defaultSE -> central SE
 - **Description:** Replicate the file registered in test 8-2 to the chosen "central" SE using *lcg-rep* commamnd.
- 3rd party lcg-cr -> central SE
 - **Description:** Copy and register a short text file to the chosen "central" SE using *lcg-cr* commamnd.
- 3rd Party lcg-cp central SE to WN
 - **Description:** Copy the file registered in test 8-5 to the WN using *lcg-cp* commamnd.
- 3rd Party lcg-rep central SE to defaultSE
 - **Description:** Replicate the file registered in test 8-5 to the default SE using *lcg-rep* commamnd.
- Replication Management using lcg-tools
- lcg-del from defaultSE
 - **Description:** Delete replicas of all the files registered in previous tests using *lcg-del* command.

SFT results display

Site Functional Tests report - Microsoft Internet Explorer

Address: <https://lcg-sft.cern.ch:9443/sft/lastreport.cgi>

Test summary

	SD	JL	JS	CT	OK	total
dteam	25	15	8	8	144	200

Colours definition

SD	Scheduled downtime	#a3a3a3
JL	Job list match failed	#aab3ff
JS	Job submission failed	#f4876b
CT	Critical tests failed	#f9d48e
NT	Non-critical tests failed	#f2f98e
OK	OK	#b2f98e

Test abbreviations

csch	CSH test
lfcrm	Replica Management using LFC
rgmasc	R-GMA Secure Connector
swdir	VO software directory
rgma	R-GMA
ver	Software Version (WN)
wn	WN host name
ca	CA certs version
rm	Replica Management
votag	VO Tag management
js	Job submission
bi	BrokerInfo
apel	Apel test

	St.	Region	Site Name	Site CE	VO dteam																			
					St.	js	ver	wn	ca	rgma	bi	csch	rm	lfcrm	votag	swdir	rgmasc	apel						
1.	SD	SouthEasternEurope	AEGIS01-PHY-SCL	ce.phy.bg.ac.yu	SD	○	2	6	0	I	○	○	○	○	○	○	○	○	○	○	○	○	○	○
2.	OK	CERN	ALBERTA-LCG2	logce01.nic.ualberta.ca	OK	○	2	6	0	I	W	○	○	○	○	○	○	○	○	○	X	X	X	X
3.	OK	NorthernEurope	BEgrid-KULeuven	kg-ce01.cc.kuleuven.ac.be	OK	○	2	6	0	I	○	○	○	○	○	○	○	○	○	○	X	X	X	X
4.	CT	NorthernEurope	BEgrid-UGent	gridce.atlantis.ugent.be	CT	○	2	6	0	I	○	○	○	○	X	X	W	○	○	X	X	X	X	X
5.	OK	NorthernEurope	BEgrid-ULB-VUB	gridce.ihe.ac.be	OK	○	2	6	0	I	○	○	○	○	○	X	○	○	○	○	○	○	○	X
6.	OK	CERN	BEIJING-CNIC-LCG2-IA64	ce-lcg.sdg.ac.cn	OK	○	2	6	0	I	W	X	○	○	○	X	W	○	○	X	X	X	X	X
7.	OK	CERN	BEIJING-LCG2	log002.ihep.ac.cn	OK	○	2	6	0	I	○	X	○	○	○	○	○	○	○	○	X	X	X	X
8.	OK	NorthernEurope	BelGrid-UCL	ingrid.cism.ucl.ac.be	OK	○	2	6	0	I	○	○	○	○	○	○	W	○	○	○	○	○	X	X
9.	OK	SouthEasternEurope	BG-INRNE	ce1.inrne.bas.bg	OK	○	2	6	0	I	○	○	○	○	○	X	○	○	○	○	○	○	○	○


See <https://lcg-sft.cern.ch/sft/lastreport.cgi>

GRIDVIEW: A Visualisation Tool for LCG

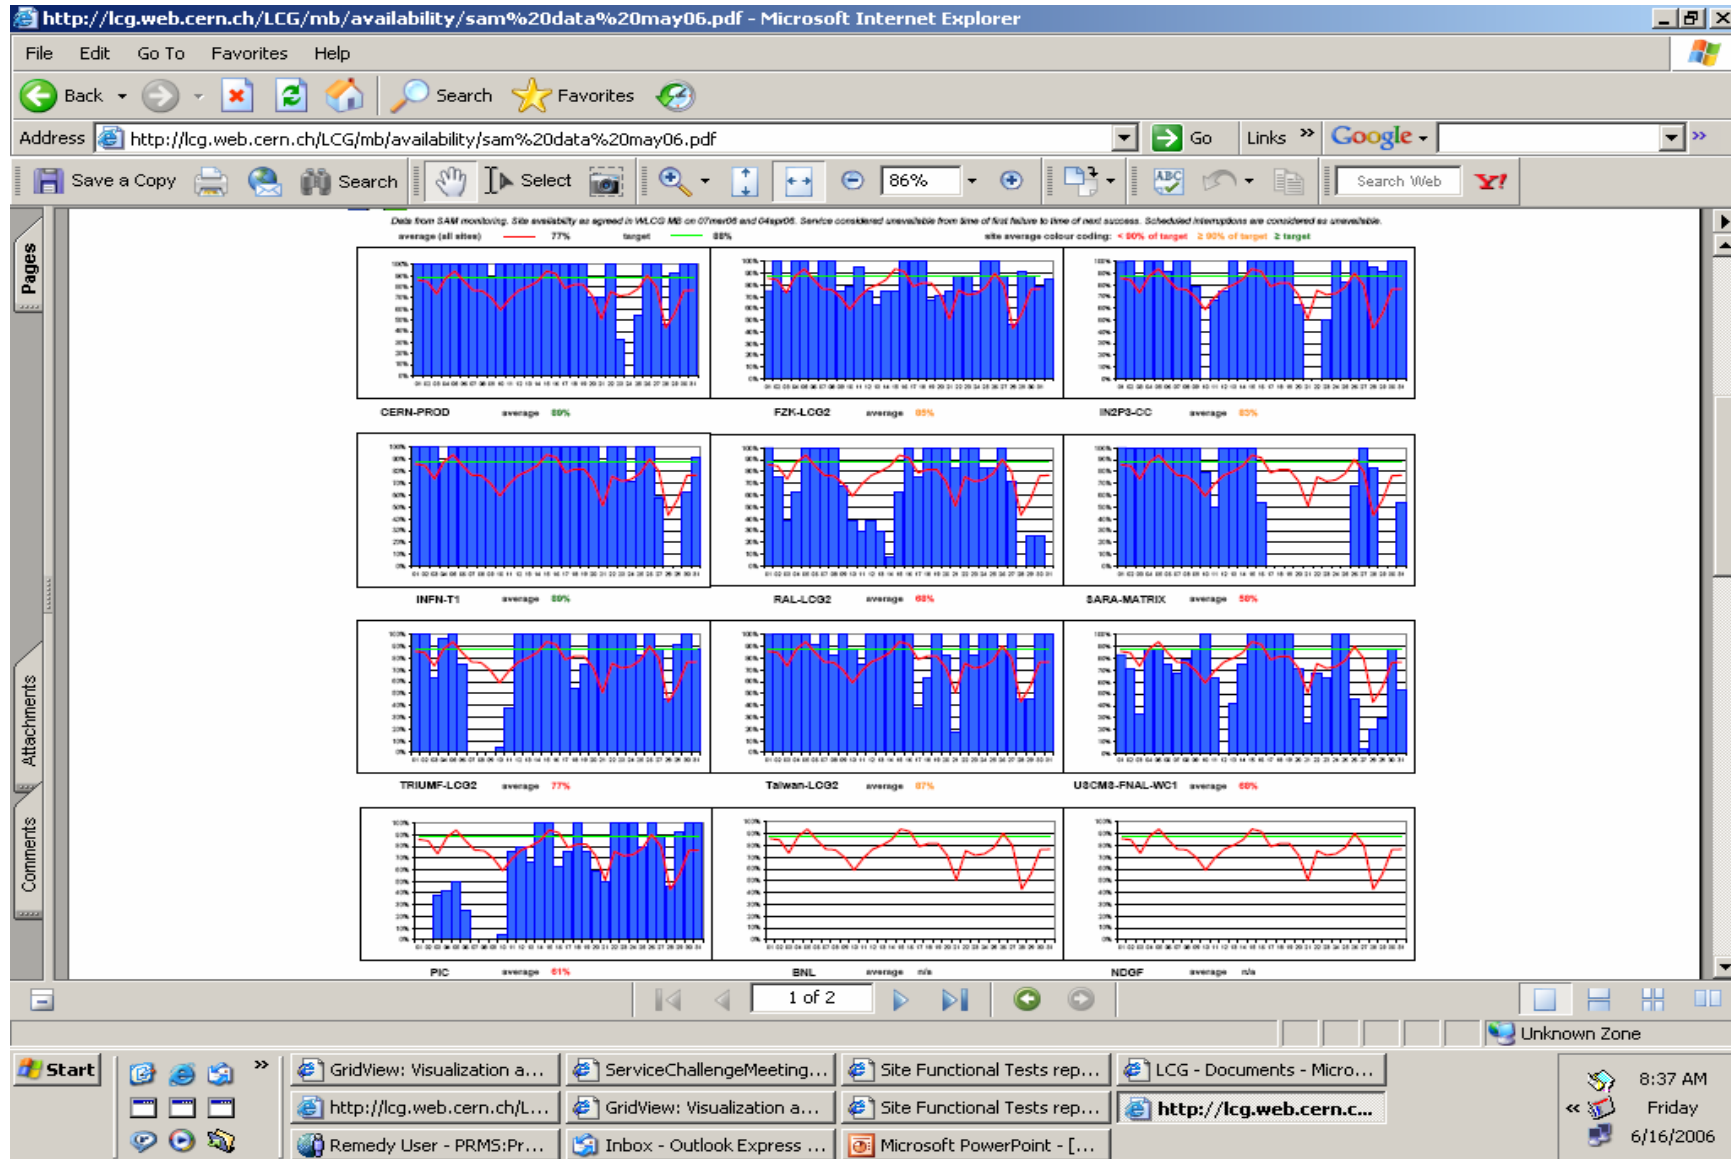
- An umbrella tool for visualisation which can provide a high level view of various grid resources and functional aspects of the LCG.
- To display a dash-board for the entire grid and provide a summary of various metrics monitored by different tools at various sites.
- To be used to notify fault situations to grid operations, user defined events to VOs, by site and network administrators to view metrics for their sites and VO administrators to see resource availability and usage for their VOs.
- GRIDVIEW is currently monitoring gridftp data transfer rates (used extensively during SC3) and job statistics
- The Basic framework for metric visualization by representing grid sites on a world map is ready
- We are extending GRIDVIEW to satisfy our service metrics requirements. Starting with simple service status displays of the services required at each Tier 0 and Tier 1 site. Extend to service quality metrics, including availability and down times, and quantitative metrics that allow comparison with LCG site MoU values.

See <http://gridview.cern.ch/GRIDVIEW/>

GRIDVIEW gridftp display: data rates from CERN to Tier 1 sites



Daily Availability of CERN and Tier 1 Sites – May 2006



6/16/2006

WLCG Service Availability - Tier2
Tutorial: H.Renshall

23