LHC Computing Grid Project

Quarterly Status and Progress Reports

July – October 2010

24 November 2010

Table of Contents

Milestone Report
July – October 2010

Nov 2010

Ian Bird

The status and progress with the outstanding high level milestones is as follows:

- Support for multi-user pilot jobs. Deployment continues, with all of the Tier 1 sites now having installed glexec and either SCAS or ARGUS, or GUMS for some US sites. A significant number of Tier 2 sites have now also installed the software and are available to be used by the experiments. At the moment there is very little reason for the experiments to change their software to use glexec, particularly in the first few months of data taking where the pressure to produce physics has meant that such software changes have not been considered.
- CREAM CE deployment. More than half of the sites now have at least one instance of the CREAM CE installed. The intention is that this CE should replace the old LCG-CE as it will provide better scalability and performance. The adaptation of Condor-g required by ATLAS in order to submit to CREAM has now been verified and is undergoing stress testing. A number of stability and scalability issues have been addressed by the developers. More production work can start to move to the CREAM CEs in the next months.
- Data Management prototypes. Following the data management workshop in Amsterdam in June and the follow up in London in early July, a number of prototype projects were proposed. Progress on these will be followed in the GDB meetings during the Autumn, with a goal of deciding which of them should be pursued by the end of Q4 2010.
- Automated gathering of installed capacity data. The mechanism and process for data publication are agreed, and most sites publish data now. The published data for the Tier 0 and Tier 1 sites has now been validated (the milestone for Sep 30 for this was met). The milestone for correct Tier 2 site reporting is the end of October. The reporting is available through an online tool that allows one to view the experiment requirements, the pledges, as well as the installed capacity. Care should be taken in interpreting some of the installed capacity data as sites publish their real installed capacity which may be more than their pledges if they provide resources to other than LHC experiments.

5 November 2010

WLCG Sites Reliability (OPS Tests)

May – October 2010

Average of the 8 best sites (not always the same 8)

May 10	Jun 10	Jul 10	Aug 10	Sep 10	Oct 10
100	99	100	100	98	99

Average of ALL Tier-0 and Tier-1 sites

May 10	Jun 10	Jul 10	Aug 10	Sep 10	Oct 10
98	98	97	99	94	96

Detailed Monthly Site Reliability (OPS tests)

Site	May 10	Jun 10	Jul 10	Aug 10	Sep 10	Oct 10
CA-TRIUMF	98	97	100	100	99	100
CERN	100	100	100	100	96	100
DE-KIT	100	100	93	99	87	99
ES-PIC	96	100	99	100	99	97
FR-CCIN2P3	100	100	100	100	87	96
IT-INFN-CNAF	96	96	99	99	95	100
NDGF	99	93	98	100	97	92
NL-T1	87	94	80	98	95	87
TW-ASGC	97	97	100	96	72	83
UK-T1-RAL	100	97	100	95	99	100
US-FNAL-CMS	100	100	100	99	100	100
US-T1-BNL	100	100	97	99	100	100
Target	*97*	*97*	*97*	*97*	*97*	*97*

Colors: Green > Target Orange > 90% Target Red < 90% Target

Availability of WLCG Tier-1 Sites + CERN for OPS

May 2010 - October 2010

Data from Nagios and Gridview **Plots show Availability for last 6 Months**

Availability is calculated as uptime / (total_time - time_status_was_UNKNOWN)

Target reliability for each site is 97 % and Target for 8 best sites is 98 % from January, 2009

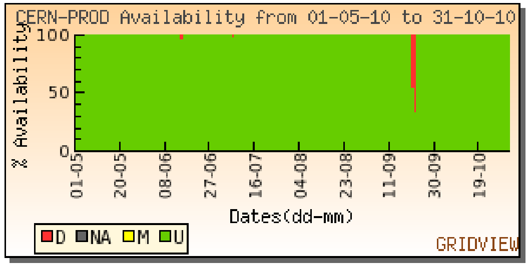

CERN **Avail : 99 %** **Unkn : 2 %**

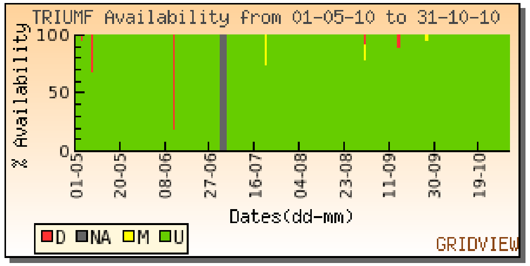

CA-TRIUMF **Avail : 99 %** **Unkn : 7 %**

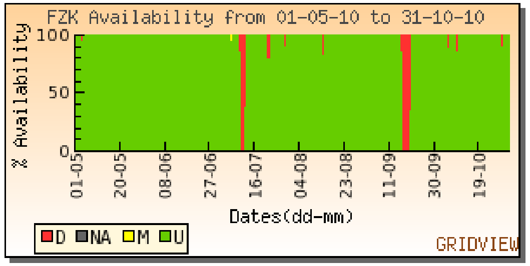

DE-KIT **Avail : 96 %** **Unkn : 1 %**

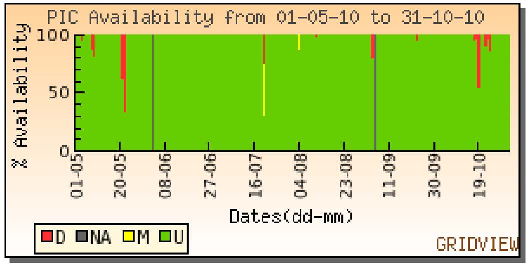

ES-PIC **Avail : 98 %** **Unkn : 6 %**

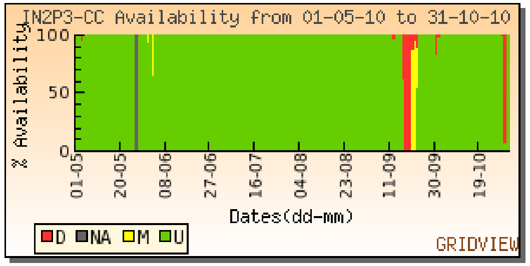

FR-CCIN2P3 **Avail : 96 %** **Unkn : 3 %**

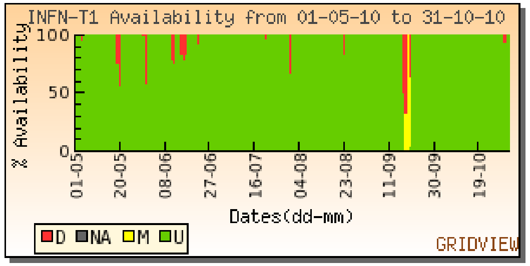

IT-INFN-CNAF **Avail : 97 %** **Unkn : 1 %**

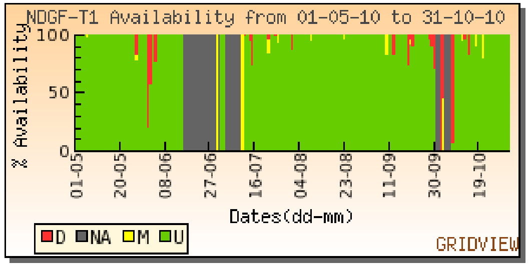

NDGF **Avail : 96 %** **Unkn : 20 %**

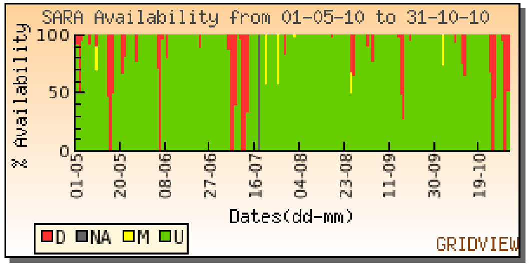

NL-T1 **Avail : 89 %** **Unkn : 4 %**

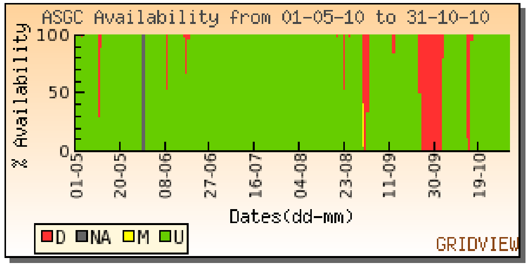

TW-ASGC **Avail : 91 %** **Unkn : 4 %**

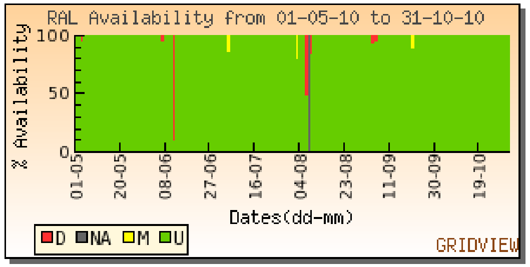

UK-T1-RAL **Avail : 99 %** **Unkn : 2 %**

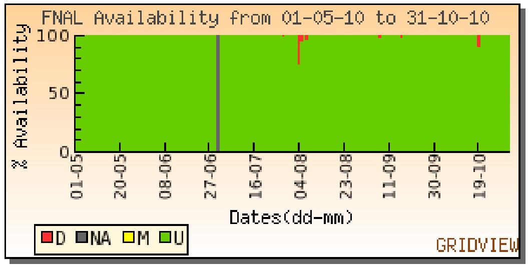

US-FNAL-CMS **Avail : 100 %** **Unkn : 2 %**

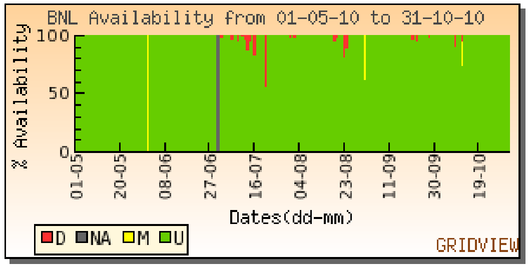

US-T1-BNL **Avail : 99 %** **Unkn : 2 %**

Availability of WLCG Tier-1 Sites + CERN for ALICE

May 2010 - October 2010

Data from SAM and Gridview **Plots show Availability for last 6 Months**

Availability is calculated as uptime / (total_time - time_status_was_UNKNOWN)

Target reliability for each site is 97 % and Target for 8 best sites is 98 % from January, 2009

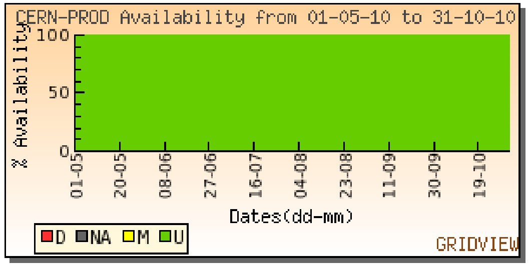

CERN **Avail : 100 %** **Unkn : 0 %**

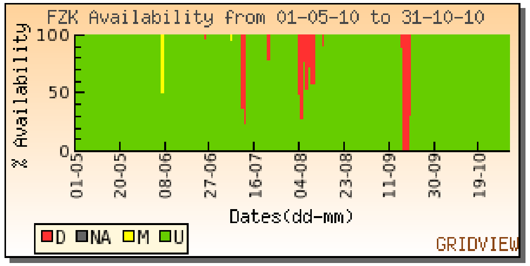

DE-KIT **Avail : 95 %** **Unkn : 1 %**

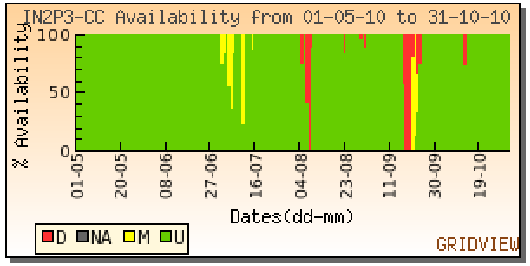

FR-CCIN2P3 **Avail : 95 %** **Unkn : 2 %**

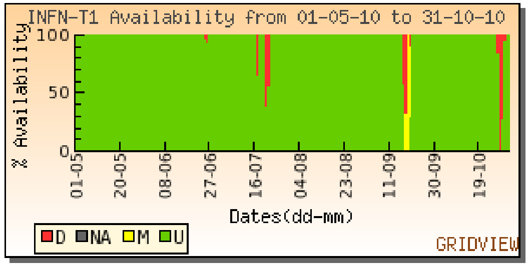

IT-INFN-CNAF **Avail : 96 %** **Unkn : 0 %**

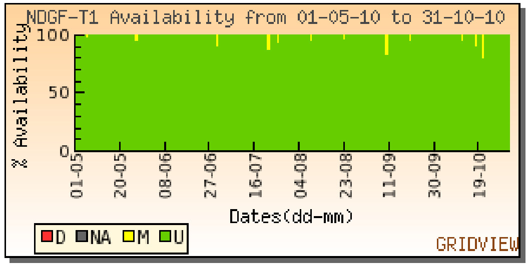

NDGF **Avail : 100 %** **Unkn : 0 %**

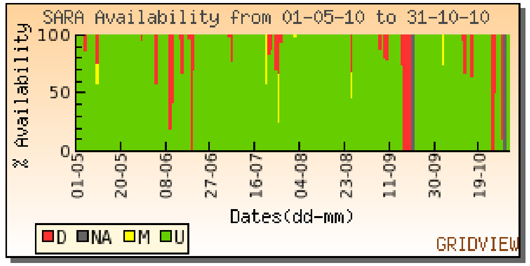

NL-T1 **Avail : 93 %** **Unkn : 5 %**

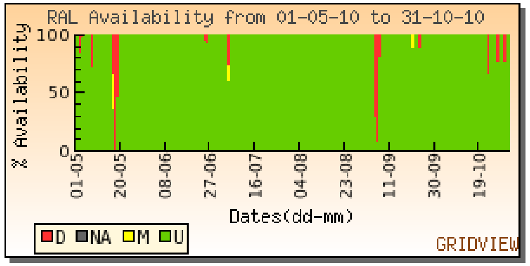

UK-T1-RAL **Avail : 97 %** **Unkn : 2 %**

Availability of WLCG Tier-1 Sites + CERN for ATLAS

May 2010 - October 2010

Data from SAM and Gridview **Plots show Availability for last 6 Months**

Availability is calculated as uptime / (total_time - time_status_was_UNKNOWN)

Target reliability for each site is 97 % and Target for 8 best sites is 98 % from January, 2009

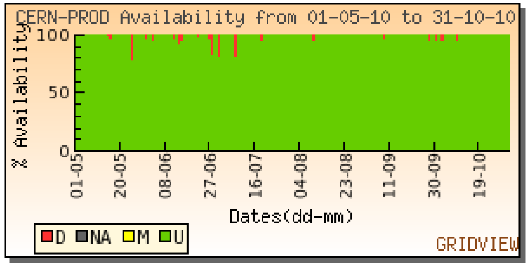

CERN **Avail : 99 %** **Unkn : 0 %**

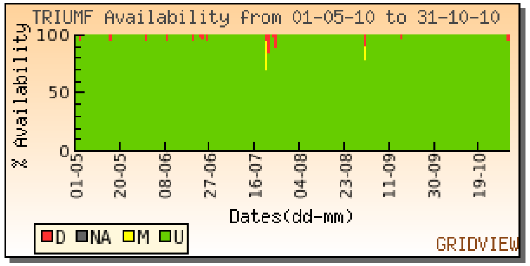

CA-TRIUMF **Avail : 99 %** **Unkn : 0 %**

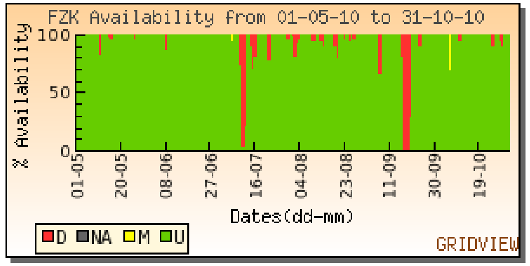

DE-KIT **Avail : 95 %** **Unkn : 0 %**

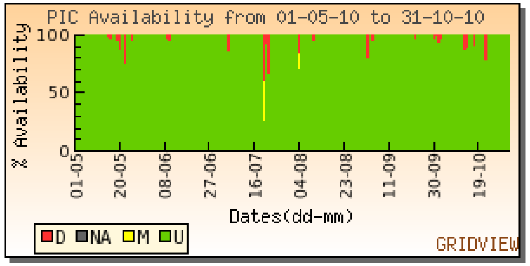

ES-PIC **Avail : 98 %** **Unkn : 1 %**

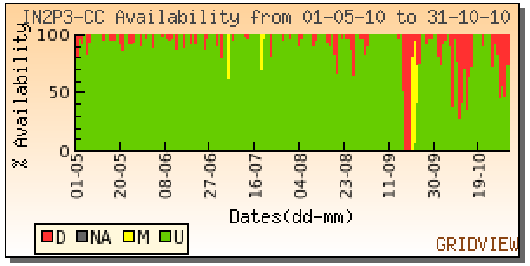

FR-CCIN2P3 **Avail : 90 %** **Unkn : 0 %**

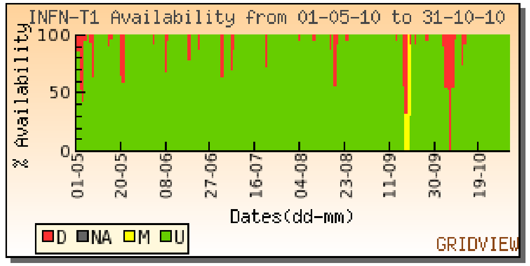

IT-INFN-CNAF **Avail : 94 %** **Unkn : 0 %**

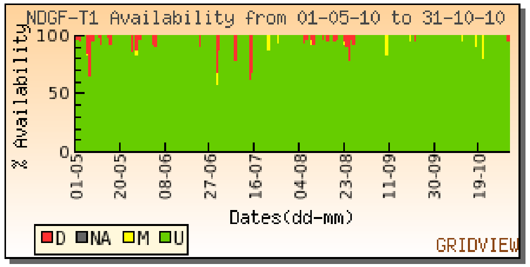

NDGF **Avail : 98 %** **Unkn : 0 %**

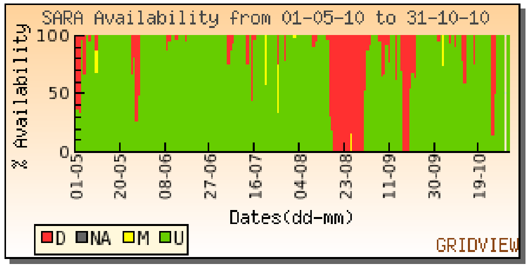

NL-T1 **Avail : 83 %** **Unkn : 0 %**

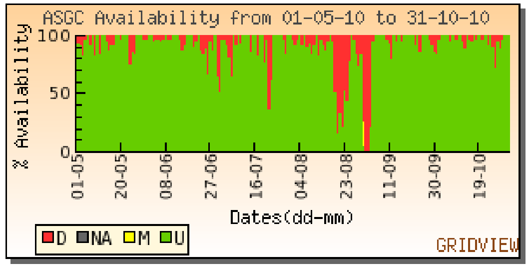

TW-ASGC **Avail : 90 %** **Unkn : 0 %**

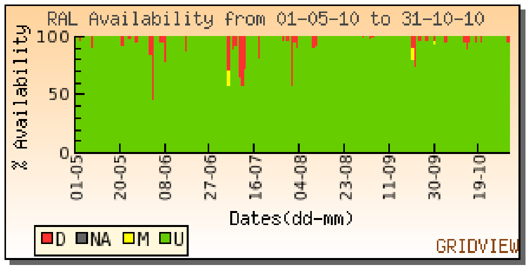

UK-T1-RAL **Avail : 97 %** **Unkn : 0 %**

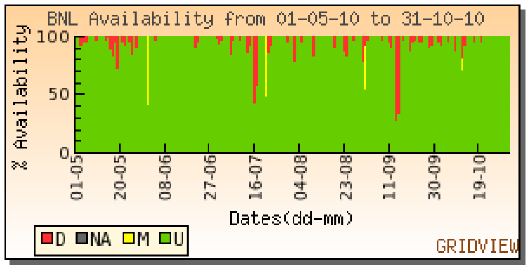

US-T1-BNL **Avail : 95 %** **Unkn : 0 %**

Availability of WLCG Tier-1 Sites + CERN for CMS

May 2010 - October 2010

Data from SAM and Gridview **Plots show Availability for last 6 Months**

Availability is calculated as uptime / (total_time - time_status_was_UNKNOWN)

Target reliability for each site is 97 % and Target for 8 best sites is 98 % from January, 2009

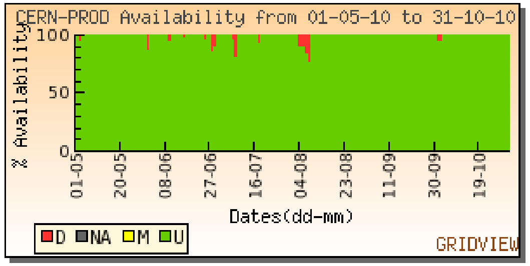

CERN **Avail : 99 %** **Unkn : 0 %**

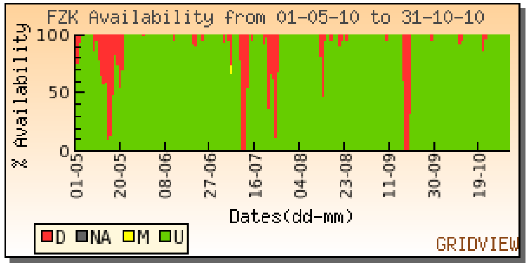

DE-KIT **Avail : 92 %** **Unkn : 1 %**

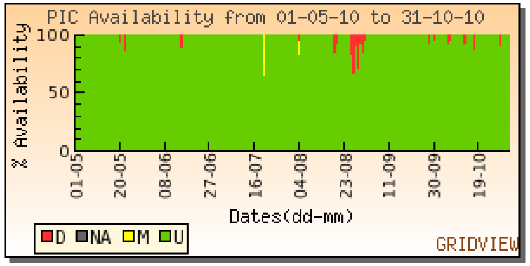

ES-PIC **Avail : 99 %** **Unkn : 0 %**

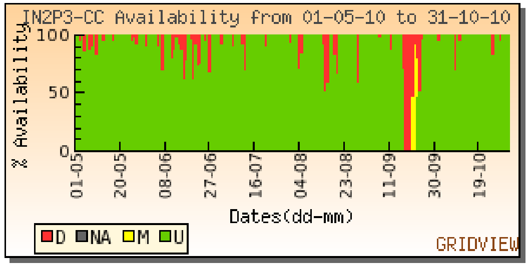

FR-CCIN2P3 **Avail : 93 %** **Unkn : 0 %**

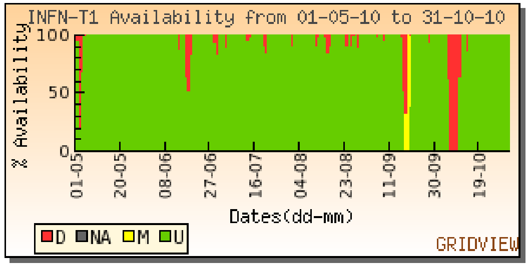

IT-INFN-CNAF **Avail : 94 %** **Unkn : 0 %**

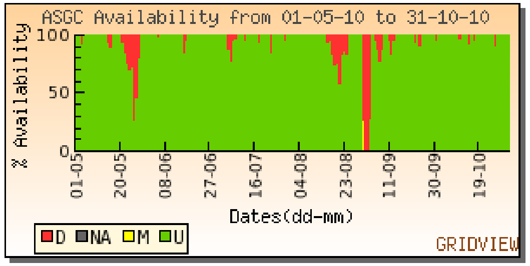

TW-ASGC **Avail : 94 %** **Unkn : 0 %**

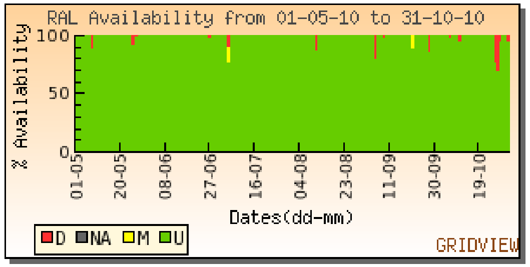

UK-T1-RAL **Avail : 99 %** **Unkn : 0 %**

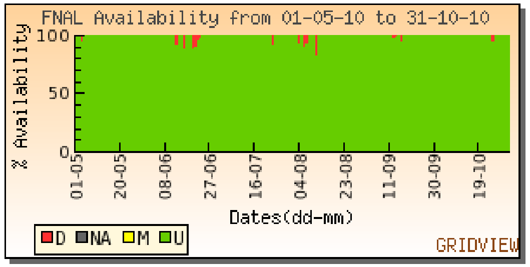

US-FNAL-CMS **Avail : 99 %** **Unkn : 0 %**

Availability of WLCG Tier-1 Sites + CERN for LHCb

May 2010 - October 2010

Data from SAM and Gridview **Plots show Availability for last 6 Months**

Availability is calculated as uptime / (total_time - time_status_was_UNKNOWN)

Target reliability for each site is 97 % and Target for 8 best sites is 98 % from January, 2009

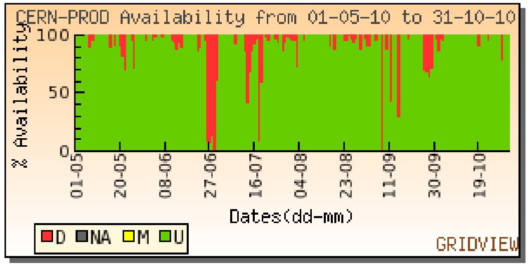

CERN **Avail : 92 %** **Unkn : 2 %**

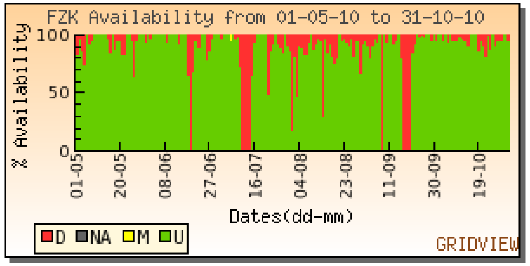

DE-KIT **Avail : 89 %** **Unkn : 3 %**

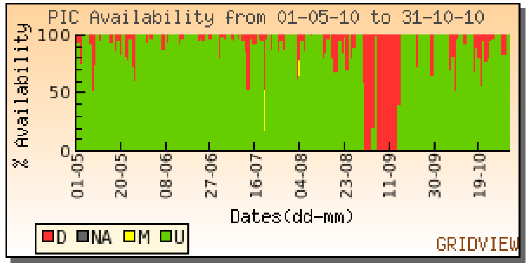

ES-PIC **Avail : 86 %** **Unkn : 1 %**

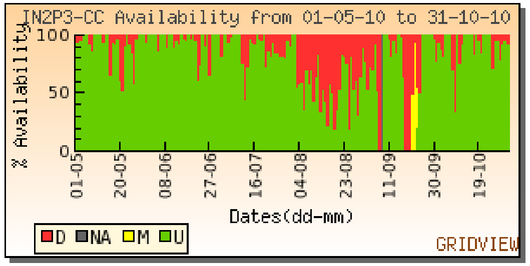

FR-CCIN2P3 **Avail : 82 %** **Unkn : 1 %**

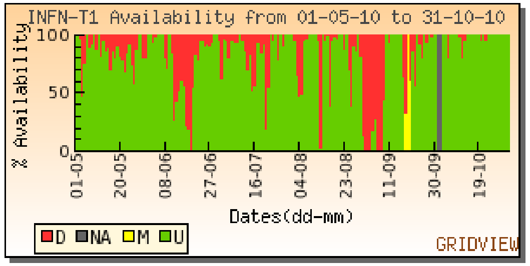

IT-INFN-CNAF **Avail : 81 %** **Unkn : 5 %**

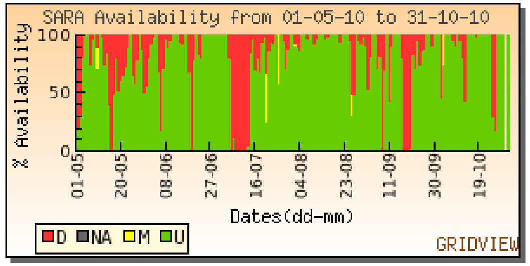

NL-T1 **Avail : 80 %** **Unkn : 2 %**

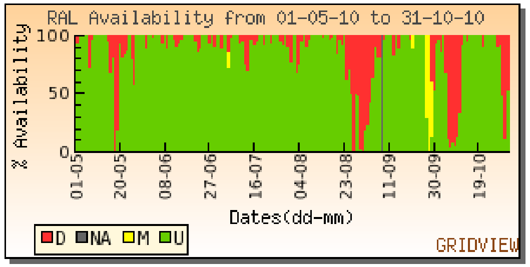

UK-T1-RAL **Avail : 85 %** **Unkn : 1 %**

Tier-2 Availability and Reliability Report

Federation Summary - Sorted by Name

October 2010

Data from Nagios and Gridview

https://twiki.cern.ch/twiki/pub/LCG/GridView/Gridview_Service_Availability_Computation.pdf

Availability = Uptime / (Total time - Time_status_was_UNKNOWN)
Reliability = Uptime / (Total time - Scheduled Downtime - Time_status_was_UNKNOWN)
HS06 : Installed capacity of the site measured in HEPSPEC06 (HS06)
Reliability and Availability for Federation - Weighted average of all sites in the Federation based on installed capacity(HS06)

Colour coding : N/A | < 30% | < 60% | < 90% | >= 90%

Federation	Reliability	Availability
AT-HEPHY-VIENNA-UIBK	100 %	100 %
AU-ATLAS	97 %	97 %
BE-TIER2	96 %	96 %
BR-SP-SPRACE	100 %	100 %
CA-EAST-T2	97 %	97 %
CA-WEST-T2	94 %	94 %
CH-CHIPP-CSCS	98 %	96 %
CN-IHEP	94 %	90 %
CZ-Prague-T2	100 %	100 %
DE-DESY-ATLAS-T2	100 %	100 %
DE-DESY-GOE-ATLAS-T2	89 %	89 %
DE-DESY-LHCB	99 %	99 %
DE-DESY-RWTH-CMS-T2	99 %	99 %
DE-FREIBURGWUPPERTAL	94 %	94 %
DE-GSI	N/A	N/A
DE-MCAT	97 %	97 %
EE-NICPB	58 %	58 %
ES-ATLAS-T2	96 %	96 %
ES-CMS-T2	95 %	95 %
ES-LHCb-T2	99 %	99 %
FI-HIP-T2	70 %	70 %
FR-GRIF	100 %	100 %
FR-IN2P3-CC-T2	96 %	96 %
FR-IN2P3-CPPM	99 %	99 %
FR-IN2P3-IPHC	90 %	90 %
FR-IN2P3-LAPP	88 %	88 %
FR-IN2P3-LPC	100 %	100 %
FR-IN2P3-SUBATECH	100 %	100 %
HU-HGCC-T2	93 %	92 %
IL-HEPTier-2	75 %	75 %
IN-DAE-KOLKATA-TIER2	99 %	86 %
IN-INDIACMS-TIFR	52 %	47 %
IT-ALICE-federation	91 %	91 %
IT-ATLAS-federation	91 %	91 %
IT-CMS-federation	91 %	91 %
IT-LHCb-federation	96 %	95 %
JP-Tokyo-ATLAS-T2	96 %	89 %
KR-KISTI-T2	100 %	100 %
KR-KNU-T2	100 %	100 %
NO-NORGRID-T2	88 %	88 %
PK-CMS-T2	94 %	93 %
PL-TIER2-WLCG	96 %	96 %
PT-LIP-LCG-Tier2	91 %	91 %
RO-LCG	89 %	89 %
RU-RDIG	98 %	98 %
SE-SNIC-T2	94 %	94 %
SI-SiGNET	84 %	84 %
T2_US_Caltech	100 %	100 %
T2_US_Florida	97 %	97 %
T2_US_MIT	99 %	98 %
T2_US_Nebraska	94 %	93 %
T2_US_Purdue	95 %	93 %
T2_US_UCSD	100 %	100 %
T2_US_Wisconsin	83 %	79 %
TR-Tier2-federation	98 %	98 %
TW-FTT-T2	80 %	80 %
UA-Tier2-Federation	N/A	N/A
UK-London-Tier2	95 %	95 %
UK-NorthGrid	77 %	77 %
UK-ScotGrid	95 %	94 %
UK-SouthGrid	98 %	98 %
US-AGLT2	100 %	100 %
US-MWT2	100 %	100 %
US-NET2	100 %	100 %
US-SWT2	100 %	97 %
US-WT2	97 %	92 %

Tier-2 Availability and Reliability Report

Federation Summary - Sorted by Name

October 2010

Data from Nagios and Gridview

https://twiki.cern.ch/twiki/pub/LCG/GridView/Gridview_Service_Availability_Computation.pdf

Availability = Uptime / (Total time - Time_status_was_UNKNOWN)
Reliability = Uptime / (Total time - Scheduled Downtime - Time_status_was_UNKNOWN)
HS06 : Installed capacity of the site measured in HEPSPEC06 (HS06)
Reliability and Availability for Federation - Weighted average of all sites in the Federation based on installed capacity(HS06)

Colour coding : N/A | < 30% | < 60% | < 90% | >= 90%

Federation	Reli-ability	Avail-ability
BR-SP-SPRACE	100 %	100 %
FR-GRIF	100 %	100 %
FR-IN2P3-LPC	100 %	100 %
KR-KISTI-T2	100 %	100 %
KR-KNU-T2	100 %	100 %
T2_US_Caltech	100 %	100 %
US-AGLT2	100 %	100 %
US-NET2	100 %	100 %
T2_US_UCSD	100 %	100 %
DE-DESY-ATLAS-T2	100 %	100 %
CZ-Prague-T2	100 %	100 %
US-MWT2	100 %	100 %
FR-IN2P3-SUBATECH	100 %	100 %
AT-HEPHY-VIENNA-UIBK	100 %	100 %
US-SWT2	100 %	97 %
ES-LHCb-T2	99 %	99 %
DE-DESY-LHCB	99 %	99 %
FR-IN2P3-CPPM	99 %	99 %
T2_US_MIT	99 %	98 %
DE-DESY-RWTH-CMS-T2	99 %	99 %
IN-DAE-KOLKATA-TIER2	99 %	86 %
TR-Tier2-federation	98 %	98 %
CH-CHIPP-CSCS	98 %	96 %
UK-SouthGrid	98 %	98 %
RU-RDIG	98 %	98 %
US-WT2	97 %	92 %
T2_US_Florida	97 %	97 %
DE-MCAT	97 %	97 %
AU-ATLAS	97 %	97 %
CA-EAST-T2	97 %	97 %
JP-Tokyo-ATLAS-T2	96 %	89 %
BE-TIER2	96 %	96 %
FR-IN2P3-CC-T2	96 %	96 %

Federation	Reli-ability	Avail-ability
IT-LHCb-federation	96 %	95 %
PL-TIER2-WLCG	96 %	96 %
ES-ATLAS-T2	96 %	96 %
UK-ScotGrid	95 %	94 %
UK-London-Tier2	95 %	95 %
ES-CMS-T2	95 %	95 %
T2_US_Purdue	95 %	93 %
DE-FREIBURGWUPPERTAL	94 %	94 %
CN-IHEP	94 %	90 %
PK-CMS-T2	94 %	93 %
T2_US_Nebraska	94 %	93 %
SE-SNIC-T2	94 %	94 %
CA-WEST-T2	94 %	94 %
HU-HGCC-T2	93 %	92 %
IT-ALICE-federation	91 %	91 %
IT-ATLAS-federation	91 %	91 %
IT-CMS-federation	91 %	91 %
PT-LIP-LCG-Tier2	91 %	91 %
FR-IN2P3-IPHC	90 %	90 %
RO-LCG	89 %	89 %
DE-DESY-GOE-ATLAS-T2	89 %	89 %
FR-IN2P3-LAPP	88 %	88 %
NO-NORGRID-T2	88 %	88 %
SI-SiGNET	84 %	84 %
T2_US_Wisconsin	83 %	79 %
TW-FTT-T2	80 %	80 %
UK-NorthGrid	77 %	77 %
IL-HEPTier-2	75 %	75 %
FI-HIP-T2	70 %	70 %
EE-NICPB	58 %	58 %
IN-INDIACMS-TIFR	52 %	47 %
DE-GSI	N/A	N/A
UA-Tier2-Federation	N/A	N/A

Tier-2 Availability and Reliability Report

October 2010

Data from Nagios and Gridview

https://twiki.cern.ch/twiki/pub/LCG/GridView/Gridview_Service_Availability_Computation.pdf

Availability = Uptime / (Total time - Time_status_was_UNKNOWN)
Reliability = Uptime / (Total time - Scheduled Downtime - Time_status_was_UNKNOWN)
HS06 : Installed capacity of the site measured in HEPSPEC06 (HS06)
Reliability and Availability for Federation - Weighted average of all sites in the Federation based on installed capacity(HS06)

Colour coding : N/A | < 30% | < 60% | < 90% | >= 90%

Federation Site	Phy. CPU	Log. CPU	HS06	Reliability	Availability	Unknown	Reliability History Jul-10	Aug-10	Sep-10
AT-HEPHY-VIENNA-UIBK (Austria, Austrian Tier-2 Federation)									
HEPHY-UIBK	39	262	1,874	100 %	100 %	0 %	100 %	79 %	82 %
Hephy-Vienna	150	665	833	99 %	99 %	0 %	99 %	99 %	97 %
AU-ATLAS (Australia, University of Melbourne)									
Australia-ATLAS	54	216	2,538	97 %	97 %	0 %	21 %	99 %	96 %
BE-TIER2 (Belgium, Belgian Tier-2 Federation)									
BEgrid-ULB-VUB	200	730	N/A	97 %	97 %	0 %	100 %	100 %	99 %
BelGrid-UCL	545	631	N/A	95 %	95 %	0 %	97 %	82 %	79 %
BR-SP-SPRACE (Brazil, SPRACE, São Paulo)									
sprace	N/A	N/A	N/A	100 %	100 %	3 %	99 %	99 %	97 %
CA-EAST-T2 (Canada-East Federation)									
CA-SCINET-T2	168	672	8,400	97 %	97 %	0 %	77 %	94 %	77 %
CA-WEST-T2 (Canada-West Federation)									
CA-ALBERTA-WESTGRID-T2	56	224	1,814	81 %	81 %	10 %	82 %	80 %	90 %
CA-VICTORIA-WESTGRID-T2	84	336	4,539	96 %	96 %	0 %	98 %	90 %	98 %
SFU-LCG2	70	280	3,376	97 %	97 %	0 %	100 %	98 %	100 %
CH-CHIPP-CSCS (Switzerland, CHIPP)									
CSCS-LCG2	192	1,152	11,520	98 %	96 %	0 %	92 %	86 %	97 %
CN-IHEP (China, IHEP, Beijing)									
BEIJING-LCG2	226	904	8,885	94 %	90 %	0 %	98 %	99 %	96 %
CZ-Prague-T2 (Czech Rep., FZU AS, Prague)									
prague_cesnet_lcg2	40	160	1,314	98 %	98 %	10 %	100 %	91 %	82 %
praguelcg2	682	2,728	21,715	100 %	100 %	10 %	96 %	100 %	100 %
DE-DESY-ATLAS-T2 (Germany ATLAS Federation, DESY)									
DESY-HH	736	4,232	33,941	100 %	100 %	0 %	98 %	96 %	97 %
DESY-ZN	178	712	6,636	99 %	99 %	0 %	100 %	99 %	100 %
DE-DESY-GOE-ATLAS-T2 (Germany, ATLAS Federation, HH/Goe)									

Federation	Site	Phy. CPU	Log. CPU	HS06	Reliability	Availability	Unknown	Reliability History Jul-10	Aug-10	Sep-10
	GoeGrid	266	1,064	N/A	89 %	89 %	0 %	96 %	100 %	81 %
DE-DESY-LHCB (Germany LHCb Federation, DESY)										
	DESY-ZN	178	712	6,636	99 %	99 %	0 %	100 %	99 %	100 %
DE-DESY-RWTH-CMS-T2 (Germany, CMS Federation)										
	DESY-HH	736	4,232	33,941	100 %	100 %	0 %	98 %	96 %	97 %
	DESY-ZN	178	712	6,636	99 %	99 %	0 %	100 %	99 %	100 %
	RWTH-Aachen	506	2,024	17,002	98 %	98 %	0 %	94 %	93 %	85 %
DE-FREIBURGWUPPERTAL (Germany, ATLAS Federation FR/W)										
	UNI-FREIBURG	204	816	26,112	95 %	95 %	0 %	96 %	100 %	75 %
	wuppertalprod	232	928	29,696	93 %	93 %	0 %	99 %	84 %	100 %
DE-GSI (Germany, GSI, Darmstadt)										
	GSI-LCG2	2	8	N/A	N/A	N/A	N/A	N/A	N/A	N/A
DE-MCAT (Germany, ATLAS Federation, Munich)										
	LRZ-LMU	600	1,200	N/A	97 %	97 %	0 %	86 %	100 %	99 %
	MPPMU	109	872	N/A	98 %	98 %	3 %	94 %	95 %	90 %
EE-NICPB (Estonia, NICPB, Tallinn)										
	T2_Estonia	63	404	N/A	58 %	58 %	1 %	78 %	93 %	44 %
ES-ATLAS-T2 (Spain, ATLAS Federation)										
	IFIC-LCG2	450	1,720	13,258	93 %	93 %	2 %	95 %	87 %	96 %
	UAM-LCG2	82	329	2,961	100 %	100 %	1 %	100 %	99 %	100 %
	ifae	30	360	4,949	99 %	99 %	2 %	97 %	100 %	100 %
ES-CMS-T2 (Spain, CMS Federation)										
	CIEMAT-LCG2	298	836	9,581	92 %	91 %	2 %	100 %	100 %	95 %
	IFCA-LCG2	N/A	N/A	14,146	97 %	97 %	1 %	95 %	99 %	98 %
ES-LHCb-T2 (Spain, LHCb Federation)										
	UB-LCG2	68	272	898	100 %	100 %	1 %	100 %	63 %	96 %
	USC-LCG2	164	354	1,985	99 %	99 %	1 %	100 %	100 %	100 %
FI-HIP-T2 (Finland, NDGF/HIP Tier2)										
	CSC	32	64	N/A	70 %	70 %	16 %	99 %	93 %	78 %
FR-GRIF (France, GRIF, Paris)										
	GRIF	1,383	5,559	44,904	100 %	100 %	1 %	100 %	100 %	98 %
FR-IN2P3-CC-T2 (France, CC-IN2P3 AF)										
	IN2P3-CC-T2	451	2,537	21,818	96 %	96 %	1 %	100 %	97 %	86 %
FR-IN2P3-CPPM (CPPM, Marseille)										
	IN2P3-CPPM	180	645	5,160	99 %	99 %	1 %	100 %	96 %	88 %
FR-IN2P3-IPHC (France, CC-IN2P3 IPHC)										
	IN2P3-IRES	256	1,216	10,944	90 %	90 %	1 %	100 %	99 %	77 %
FR-IN2P3-LAPP (France, LAPP, Annecy)										
	IN2P3-LAPP	220	752	7,391	88 %	88 %	5 %	100 %	100 %	84 %

Federation	Site	Phy. CPU	Log. CPU	HS06	Reliability	Availability	Unknown	Reliability History Jul-10	Aug-10	Sep-10
FR-IN2P3-LPC (France, LPC, Clermont-Ferrand)										
	IN2P3-LPC	452	1,824	14,296	100 %	100 %	1 %	92 %	100 %	85 %
FR-IN2P3-SUBATECH (France, SUBATECH, Nantes)										
	IN2P3-SUBATECH	104	408	3,487	100 %	100 %	1 %	100 %	100 %	87 %
HU-HGCC-T2 (Hungary, HGCC Federation)										
	BUDAPEST	103	406	3,665	94 %	94 %	14 %	100 %	95 %	99 %
	ELTE	44	56	175	66 %	61 %	14 %	81 %	45 %	99 %
IL-HEPTier-2 (Israel, HEP-IL Tier-2 Federation)										
	IL-TAU-HEP	34	272	N/A	90 %	90 %	11 %	98 %	57 %	85 %
	TECHNION-HEP	34	272	N/A	91 %	91 %	1 %	98 %	100 %	98 %
	WEIZMANN-LCG2	124	496	54,560	75 %	75 %	1 %	71 %	76 %	91 %
IN-DAE-KOLKATA-TIER2 (India, VECC/SINP, Kolkata)										
	IN-DAE-VECC-02	94	220	N/A	99 %	86 %	0 %	100 %	99 %	98 %
IN-INDIACMS-TIFR (India, TIFR, Mumbai)										
	INDIACMS-TIFR	30	240	20,400	52 %	47 %	0 %	82 %	92 %	72 %
IT-ALICE-federation (Italy, INFN ALICE Federation)										
	INFN-BARI	274	1,045	9,405	97 %	97 %	0 %	99 %	91 %	97 %
	INFN-CATANIA	156	440	2,640	27 %	27 %	0 %	66 %	99 %	98 %
	INFN-FRASCATI	36	176	1,433	96 %	96 %	0 %	96 %	96 %	76 %
	INFN-LNL-2	176	656	6,704	100 %	100 %	0 %	99 %	100 %	100 %
	INFN-MILANO-ATLASC	52	434	3,528	79 %	79 %	0 %	74 %	92 %	99 %
	INFN-NAPOLI-ATLAS	98	434	3,641	94 %	89 %	0 %	98 %	99 %	96 %
	INFN-PISA	764	3,056	30,560	93 %	92 %	0 %	96 %	79 %	92 %
	INFN-ROMA1	109	544	4,406	100 %	100 %	0 %	99 %	100 %	100 %
	INFN-ROMA1-CMS	106	356	3,072	84 %	84 %	0 %	99 %	100 %	100 %
	INFN-TORINO	128	488	3,649	99 %	99 %	0 %	99 %	100 %	100 %
IT-ATLAS-federation (Italy, INFN ATLAS Federation)										
	INFN-BARI	274	1,045	9,405	97 %	97 %	0 %	99 %	91 %	97 %
	INFN-CATANIA	156	440	2,640	27 %	27 %	0 %	66 %	99 %	98 %
	INFN-FRASCATI	36	176	1,433	96 %	96 %	0 %	96 %	96 %	76 %
	INFN-LNL-2	176	656	6,704	100 %	100 %	0 %	99 %	100 %	100 %
	INFN-MILANO-ATLASC	52	434	3,528	79 %	79 %	0 %	74 %	92 %	99 %
	INFN-NAPOLI-ATLAS	98	434	3,641	94 %	89 %	0 %	98 %	99 %	96 %
	INFN-PISA	764	3,056	30,560	93 %	92 %	0 %	96 %	79 %	92 %
	INFN-ROMA1	109	544	4,406	100 %	100 %	0 %	99 %	100 %	100 %
	INFN-ROMA1-CMS	106	356	3,072	84 %	84 %	0 %	99 %	100 %	100 %
	INFN-TORINO	128	488	3,649	99 %	99 %	0 %	99 %	100 %	100 %
IT-CMS-federation (Italy, INFN CMS Federation)										
	INFN-BARI	274	1,045	9,405	97 %	97 %	0 %	99 %	91 %	97 %
	INFN-CATANIA	156	440	2,640	27 %	27 %	0 %	66 %	99 %	98 %
	INFN-FRASCATI	36	176	1,433	96 %	96 %	0 %	96 %	96 %	76 %

Federation	Site	Phy. CPU	Log. CPU	HS06	Reliability	Availability	Unknown	Reliability History Jul-10	Aug-10	Sep-10
	INFN-LNL-2	176	656	6,704	100 %	100 %	0 %	99 %	100 %	100 %
	INFN-MILANO-ATLASC	52	434	3,528	79 %	79 %	0 %	74 %	92 %	99 %
	INFN-NAPOLI-ATLAS	98	434	3,641	94 %	89 %	0 %	98 %	99 %	96 %
	INFN-PISA	764	3,056	30,560	93 %	92 %	0 %	96 %	79 %	92 %
	INFN-ROMA1	109	544	4,406	100 %	100 %	0 %	99 %	100 %	100 %
	INFN-ROMA1-CMS	106	356	3,072	84 %	84 %	0 %	99 %	100 %	100 %
	INFN-TORINO	128	488	3,649	99 %	99 %	0 %	99 %	100 %	100 %
IT-LHCb-federation (Italy, INFN LHCb Federation)										
	INFN-BARI	274	1,045	9,405	97 %	97 %	0 %	99 %	91 %	97 %
	INFN-CATANIA	156	440	2,640	27 %	27 %	0 %	66 %	99 %	98 %
	INFN-CNAF-LHCB	2,252	8,192	85,516	99 %	99 %	0 %	99 %	99 %	89 %
	INFN-FRASCATI	36	176	1,433	96 %	96 %	0 %	96 %	96 %	76 %
	INFN-LNL-2	176	656	6,704	100 %	100 %	0 %	99 %	100 %	100 %
	INFN-MILANO-ATLASC	52	434	3,528	79 %	79 %	0 %	74 %	92 %	99 %
	INFN-NAPOLI-ATLAS	98	434	3,641	94 %	89 %	0 %	98 %	99 %	96 %
	INFN-PISA	764	3,056	30,560	93 %	92 %	0 %	96 %	79 %	92 %
	INFN-ROMA1	109	544	4,406	100 %	100 %	0 %	99 %	100 %	100 %
	INFN-ROMA1-CMS	106	356	3,072	84 %	84 %	0 %	99 %	100 %	100 %
	INFN-TORINO	128	488	3,649	99 %	99 %	0 %	99 %	100 %	100 %
JP-Tokyo-ATLAS-T2 (Japan, ICEPP, Tokyo)										
	TOKYO-LCG2	288	1,152	16,531	96 %	89 %	2 %	98 %	99 %	99 %
KR-KISTI-T2 (Republic of Korea, KISTI, Daejeon)										
	KR-KISTI-GCRT-01	134	536	8,849	100 %	100 %	0 %	87 %	80 %	100 %
KR-KNU-T2 (Republic of Korea, CHEP of KNU, Daegu)										
	LCG_KNU	140	336	N/A	100 %	100 %	0 %	95 %	100 %	100 %
NO-NORGRID-T2 (Norway, UNINETT SIGMA Tier-2)										
	NO-NORGRID-T2	N/A	N/A	N/A	88 %	88 %	27 %	96 %	96 %	85 %
PK-CMS-T2 (Pakistan, Pakistan Tier-2 Federation)										
	NCP-LCG2	60	240	2,400	94 %	93 %	0 %	99 %	98 %	99 %
	PAKGRID-LCG2	10	10	N/A	N/A	N/A	N/A	N/A	N/A	N/A
PL-TIER2-WLCG (Poland, Polish Tier-2 Federation)										
	CYFRONET-LCG2	1,020	5,104	52,571	96 %	96 %	0 %	99 %	98 %	74 %
	PSNC	532	4,720	N/A	99 %	99 %	0 %	97 %	68 %	56 %
	WARSAW-EGEE	396	1,376	N/A	90 %	50 %	0 %	100 %	95 %	92 %
PT-LIP-LCG-Tier2 (Portugal, LIP Tier-2 Federation)										
	LIP-Coimbra	46	184	1,879	97 %	97 %	1 %	86 %	93 %	99 %
	LIP-Lisbon	139	532	5,432	100 %	100 %	1 %	87 %	96 %	100 %
	NCG-INGRID-PT	312	1,248	8,524	84 %	84 %	4 %	88 %	87 %	89 %
RO-LCG (Romania, Romanian Tier-2 Federation)										
	NIHAM	1	2	N/A	58 %	58 %	8 %	0 %	48 %	91 %
	RO-02-NIPNE	84	212	N/A	N/A	N/A	100 %	0 %	N/A	N/A

Federation	Site	Phy. CPU	Log. CPU	HS06	Reliability	Availability	Unknown	Reliability History Jul-10	Aug-10	Sep-10
	RO-07-NIPNE	258	258	2,180	86 %	86 %	7 %	96 %	89 %	71 %
	RO-11-NIPNE	12	12	122	0 %	0 %	0 %	67 %	3 %	3 %
	RO-13-ISS	16	64	64,850	90 %	90 %	7 %	97 %	95 %	88 %
	RO-14-ITIM	40	160	1,304	87 %	58 %	9 %	85 %	98 %	100 %
	RO-16-UAIC	66	264	2,376	85 %	85 %	7 %	81 %	98 %	98 %
RU-RDIG (Russian Fed., Russian Data-Intensive GRID)										
	ITEP	136	240	N/A	99 %	99 %	33 %	99 %	78 %	97 %
	JINR-LCG2	566	1,132	11,364	100 %	100 %	32 %	100 %	100 %	98 %
	RRC-KI	1,784	1,784	N/A	44 %	44 %	30 %	82 %	89 %	99 %
	RU-Protvino-IHEP	120	400	N/A	100 %	100 %	33 %	96 %	100 %	99 %
	RU-SPbSU	12	48	N/A	55 %	55 %	32 %	80 %	76 %	79 %
	Ru-Troitsk-INR-LCG2	41	162	N/A	78 %	78 %	30 %	95 %	89 %	80 %
	ru-Moscow-FIAN-LCG2	30	52	N/A	12 %	12 %	10 %	77 %	99 %	95 %
	ru-Moscow-MEPHI-LCG2	54	168	N/A	N/A	N/A	100 %	N/A	N/A	N/A
	ru-Moscow-SINP-LCG2	116	232	N/A	91 %	91 %	32 %	81 %	91 %	96 %
	ru-PNPI	156	312	3,432	92 %	92 %	33 %	85 %	99 %	98 %
SE-SNIC-T2 (Sweden, SNIC Tier-2)										
	SE-SNIC-T2	N/A	N/A	N/A	94 %	94 %	24 %	60 %	94 %	95 %
SI-SiGNET (Slovenia, SiGNET)										
	SiGNET	335	1,162	11,388	84 %	84 %	27 %	98 %	100 %	72 %
T2_US_Caltech (USA, Caltech CMS T2)										
	cit_cms_t2	N/A	N/A	N/A	100 %	100 %	0 %	100 %	98 %	99 %
T2_US_Florida (USA, Florida CMS T2)										
	uflorida-hpc	N/A	N/A	N/A	95 %	95 %	8 %	96 %	100 %	100 %
	uflorida-pg	N/A	N/A	N/A	100 %	100 %	1 %	100 %	100 %	96 %
T2_US_MIT (USA, MIT CMS T2)										
	mit_cms	N/A	N/A	N/A	99 %	98 %	6 %	100 %	100 %	100 %
T2_US_Nebraska (USA, Nebraska CMS T2)										
	nebraska	N/A	N/A	N/A	94 %	93 %	4 %	100 %	100 %	100 %
T2_US_Purdue (USA, Purdue CMS T2)										
	purdue-rcac	N/A	N/A	N/A	100 %	100 %	0 %	99 %	100 %	100 %
	purdue-steele	N/A	N/A	N/A	89 %	86 %	21 %	100 %	99 %	99 %
T2_US_UCSD (USA, UC San Diego CMS T2)										
	ucsdt2	N/A	N/A	N/A	100 %	100 %	3 %	99 %	98 %	100 %
T2_US_Wisconsin (USA, U. Wisconsin CMS T2)										
	GLOW	N/A	N/A	N/A	83 %	79 %	29 %	100 %	99 %	99 %
TR-Tier2-federation (Turkey, Turkish Tier-2 Federation)										
	TR-03-METU	154	308	N/A	100 %	100 %	21 %	97 %	96 %	91 %
	TR-10-ULAKBIM	202	468	3,229	98 %	98 %	18 %	80 %	79 %	76 %
TW-FTT-T2 (Taipei, Taiwan Analysis Facility Federation)										

Federation	Site	Phy. CPU	Log. CPU	HS06	Reliability	Availability	Unknown	Reliability History Jul-10	Aug-10	Sep-10
	TW-FTT	102	404	N/A	80 %	80 %	0 %	96 %	98 %	91 %
UK-London-Tier2 (UK, London Tier 2)										
	UKI-LT2-Brunel	211	593	4,638	100 %	100 %	0 %	100 %	92 %	97 %
	UKI-LT2-IC-HEP	410	1,640	13,120	99 %	99 %	0 %	100 %	100 %	97 %
	UKI-LT2-QMUL	1,592	4,224	22,262	91 %	91 %	0 %	94 %	98 %	89 %
	UKI-LT2-RHUL	100	400	3,160	98 %	98 %	0 %	86 %	54 %	85 %
	UKI-LT2-UCL-CENTRAL	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
	UKI-LT2-UCL-HEP	104	288	2,604	100 %	100 %	0 %	100 %	100 %	96 %
UK-NorthGrid (UK, NorthGrid)										
	UKI-NORTHGRID-LANCS-HEP	88	256	3,123	97 %	97 %	1 %	N/A	79 %	67 %
	UKI-NORTHGRID-LIV-HEP	143	572	8,318	16 %	16 %	0 %	100 %	39 %	0 %
	UKI-NORTHGRID-MAN-HEP	1,010	1,810	18,482	97 %	97 %	0 %	98 %	100 %	98 %
	UKI-NORTHGRID-SHEF-HEP	300	400	3,900	100 %	100 %	0 %	100 %	100 %	97 %
UK-ScotGrid (UK, ScotGrid)										
	UKI-SCOTGRID-DURHAM	168	672	5,699	90 %	90 %	0 %	70 %	98 %	93 %
	UKI-SCOTGRID-ECDF	492	1,968	22,351	100 %	100 %	0 %	87 %	93 %	99 %
	UKI-SCOTGRID-GLASGOW	618	1,912	15,320	90 %	87 %	0 %	100 %	94 %	84 %
UK-SouthGrid (UK, SouthGrid)										
	EFDA-JET	124	248	1,714	100 %	92 %	0 %	99 %	100 %	98 %
	UKI-SOUTHGRID-BHAM-HEP	72	384	3,368	93 %	93 %	0 %	94 %	91 %	98 %
	UKI-SOUTHGRID-BRIS-HEP	80	228	1,919	100 %	100 %	0 %	75 %	100 %	94 %
	UKI-SOUTHGRID-CAM-HEP	55	220	2,181	97 %	97 %	1 %	98 %	100 %	80 %
	UKI-SOUTHGRID-OX-HEP	104	416	3,562	99 %	99 %	0 %	95 %	94 %	84 %
	UKI-SOUTHGRID-RALPP	460	1,544	12,584	99 %	99 %	0 %	96 %	90 %	84 %
US-AGLT2 (USA, Great Lakes ATLAS T2)										
	AGLT2	N/A	N/A	N/A	100 %	100 %	1 %	99 %	100 %	100 %
US-MWT2 (USA, Midwest ATLAS T2)										
	MWT2_IU	N/A	N/A	N/A	100 %	100 %	0 %	98 %	100 %	100 %
	MWT2_UC	N/A	N/A	N/A	100 %	100 %	0 %	100 %	100 %	100 %
US-NET2 (USA, Northeast ATLAS T2)										
	BU_ATLAS_Tier2	N/A	N/A	N/A	100 %	100 %	0 %	100 %	100 %	100 %
	hu_atlas_tier2	N/A	N/A	N/A	100 %	100 %	0 %	100 %	100 %	100 %
US-SWT2 (USA, Southwest ATLAS T2)										
	OU_OCHEP_SWT2	N/A	N/A	N/A	100 %	98 %	0 %	99 %	100 %	100 %
	SWT2_CPB	N/A	N/A	N/A	99 %	92 %	0 %	100 %	100 %	100 %
	UTA_SWT2	N/A	N/A	N/A	100 %	100 %	0 %	90 %	100 %	100 %
US-WT2 (USA, SLAC ATLAS T2)										
	WT2	N/A	N/A	N/A	97 %	92 %	1 %	99 %	97 %	98 %

WLCG Service Report
July – October 2010

25 October 2010

Maria Girone, Jamie Shiers

This report covers the first summer of LHC data taking – results from which were presented at major physics conferences. In particular, the WLCG service was used to process and analyze data in time for ICHEP, delivering results just a few days after data taking – an unprecedented achievement.

This success was marred by a continuing number of incidents for which a Service Incident Report was requested. Whereas in Q2 most of these incidents, with the exception of the CASTOR@CERN data loss, were solved within a matter of hours, there were several events in this quarter that took up to one month to solve. These include a prolonged database outage at the NL-T1 that made the site unusable for LHCb and the whole "cloud" (the Tier1 plus associated Tier2s) unavailable for ATLAS, a database incident at ASGC affecting the CASTOR service and a number of network degradations, one of which was only resolved after 29 days.

As a result of these incidents and leading on from work pioneered by CMS, decision points have now been proposed for major incidents as follows:

- Up to 24 hours – site internal: report expected to daily WLCG Operations meeting and to any associated tickets;
- Beyond 24 hours and up to 96 hours – first level escalation: WLCG Management Board informed; preparation for prolonged site downtime (agreement on possible backup site(s), discussion of inter-VO issues etc.);
- Beyond 96 hours and up to 2 weeks – second level escalation: continued regular reporting and priority on service recovery;
- Beyond 2 weeks – third level escalation: WLCG Overview and Collaboration Boards informed.

These targets are intended to be flexible and used as guidelines: if site / service recovery is expected shortly after one such deadline an informed decision should be made, e.g. at the WLCG Operations meeting, as to whether additional grace period should be granted. On the other hand, should it be clear earlier – e.g. in the case of a major catastrophe such as fire, when it is known that the recovery will be long – the corresponding level of escalation should be triggered immediately.

As regards network incidents: the current mechanism of assigning the problem to (one of) the sites involved is considered the correct initial approach and works well in many cases. However, there continue to be a number of incidents – such as those that occurred in this quarter – where the problem does not get correctly followed nor is resolved in a sufficiently timely manner. To address this problem, two actions are considered: *involvement* – whereby network experts are involved in the problem at an early stage and *ownership* – responsibility for following up with all relevant parties is assigned to the appropriate body. This may be the "network experts" team, a representative from a / the concerned experiment or a WLCG service representative as appropriate. This will be discussed at the upcoming LHCOPN meeting at CERN with the goal of arriving at a concrete implementation of this model in the immediate future.

As in previous quarters, the Key Performance Indicators of GGUS statistics, Site Usability plots and Service Incident Reports / Risk Assessments, continue to provide a realistic overview of the service during a given period and are used in the regular reports to the WLCG Management Board.

Summary of Main Service Incidents

Previous quarterly reports have included a table listing by date, site and service the main incidents for which a Service Incident Report was produced. These are typically characterized by a serious degradation or total loss of service of at least several hours and / or when an alarm ticket was generated.

Date	Duration	Summary
n/a	n/a	No significant incidents during this quarter (no SIRs)

Table 1 – CERN CASTOR-related Service Incident Reports

Comments: failure or degradation of the CERN CASTOR and related services can have a corresponding impact on raw data recording, first pass processing and/or data export *inter alia*. The fact that no major problems occurred during this quarter strongly suggests that the strategy to limit changes to those that are absolutely necessary has paid off. (Other factors may have also contributed to this noticeable reduction in CERN CASTOR-related SIRs).

Site	Date	Duration	Summary
CERN			LHCb LFC replication halted
CERN	13 Sep	1.5h	Spontaneous reboots of nodes 2 & 4 of CMSR
NL-T1	Aug	>3 weeks	Replication stopped due to DB problems below
NL-T1	Aug	>3 weeks	Site / cloud unavailable due to DB / recovery problems
ASGC	31 Aug	4 days	CASTOR outage due to stager DB problem
CERN	23 Aug	35h	ATLAS conditions streaming to Tier1s down
CERN	20 Aug	4h	Instability of nodes 3 & 4 of CMSR
CERN	9 Aug	16h	LHCB online DB unavailable
CERN	13 Jul	~8h	Short interruptions of online/offline streaming

Table 2 – Database Related Service Incident Reports

Comments: online – offline replication is particularly critical, followed by Tier0 services, Tier1 services and inter-site replication (where latency of some hours is normally tolerable.) During this quarter the overall service suffered from a slightly higher number of DB-related incidents (up by 50% although the statistics are thankfully low) but of which several exceeded (one by far) 24 hours. Action has been triggered by the WLCG Management Board to ensure that a similar failure to that suffered by NL-T1 does not have such major consequences. To achieve this goal, strategies for ATLAS LFCs are being studied – possibly implemented by a R/O copy of the “cloud LFC” at CERN (LHCb uses a replicated global LFC and is thus not exposed in this way.)

Site	Service	Date	Duration	Summary
RAL/NDGF	Network	27 Aug	~1 month	Data transfer errors / problems
BNL/CNAF	Network	23 Aug	~1 month	Data transfer errors / problems
PIC	CE	25 Jul	30h	Cooling problems caused 50% of WNs to be shutdown
PIC	SE	22 Jul	10h	SRM not available for ATLAS due to dCache pool costs configuration
PIC	CE	20 Jul	3h	Unavailable after scheduled downtime due to wrong gridmapdir migration
CERN	Several	19 Jul	2h	Cooling failure in vault (human error)
OSG/GOC	GOC	15 Jul	4h	GOC service outage
NDGF	SE	14 Jul	16h	SRM downtime followed by degradation
KIT	Site	10 Jul	>4h	Outage of central and local services due to cooling failure
NDGF	LFC	8 Jul	3h	lfc1.ndgf.org downtime
NL-T1	SE	5 Jul	1 week	Reduced availability caused by data corruption
KIT	SE	5 Jul	18h	CMS dCache down due to h/w failure

Table 3 – Other Service Incident Reports

Comments: cooling failures continue at roughly the same rate as in the past – these are probably unavoidable. The long delays in resolving network-related incidents will be addressed by the plan outlined earlier in this report. They often involve multiple parties, e.g. BNL,ESNET,GEANT,GARR,CNAF in the case of the BNL-CNAF transfer issues.

An analysis of Service Incidents since the beginning of 2009 shows a roughly constant number of infrastructure problems, a smaller number of network or middleware related problems, whereas data / database incidents dominate (figure 1). In addition, a high fraction of these problems are not resolved in 24 or even 96 hours (figure 2).

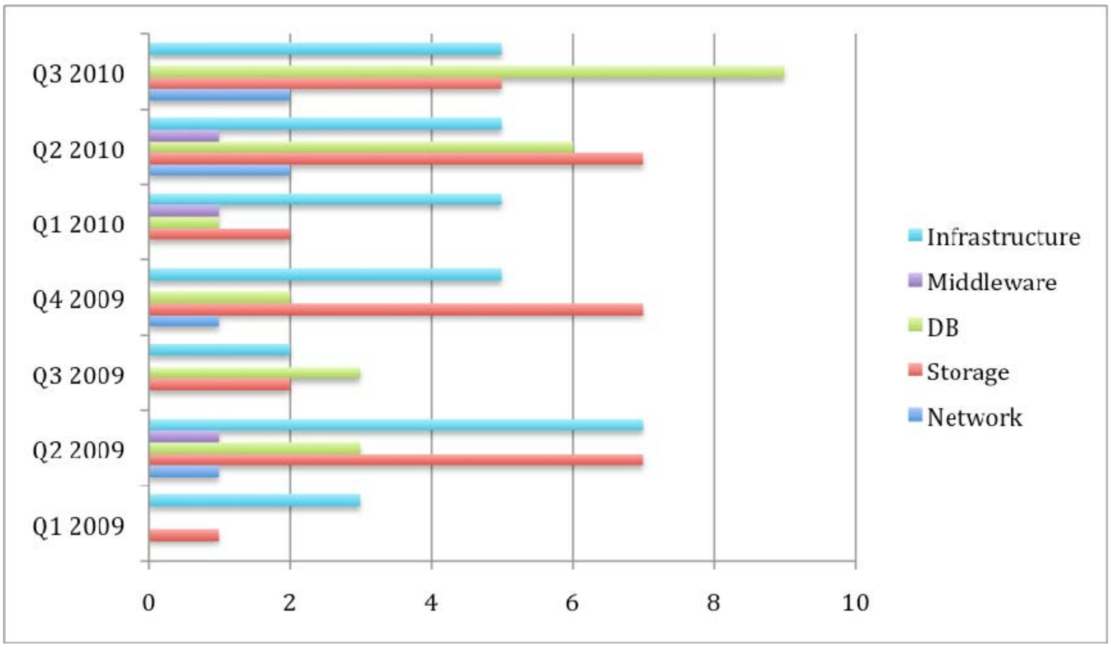

Figure 1 - Service Incidents by Area

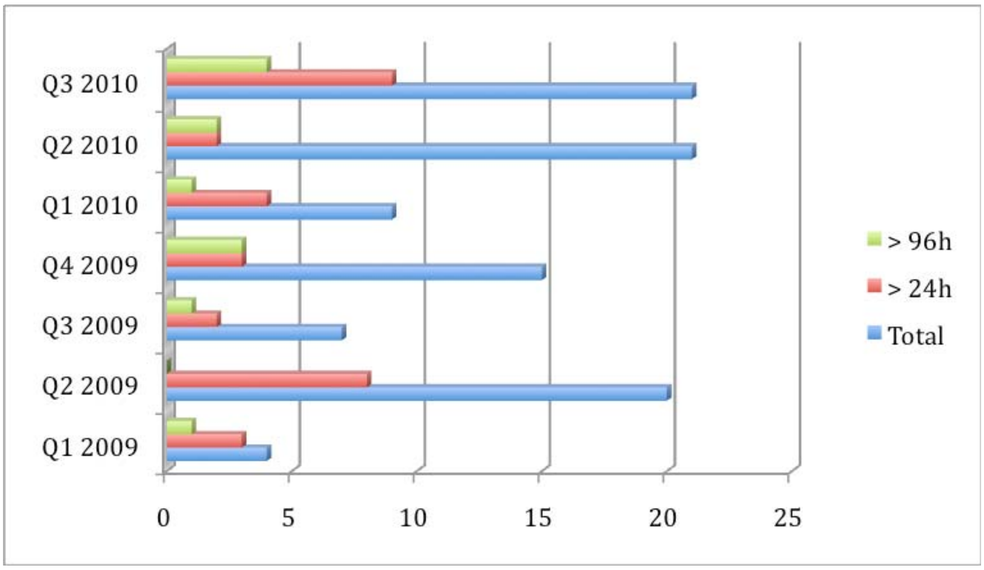

Figure 2 - Time to Resolution

Outlook for the remainder of 2010

Data taking will continue for much of the rest of the year – including a short Heavy Ion (HI) run – followed by continued reprocessing and analysis. The details of this run are still being discussed but it is expected that:

- ALICE will export data from the HI run only after the run has stopped and over several months;
- ATLAS will export data during the run to all available Tier1 sites;
- CMS will export data only to FNAL;
- (LHCb does not take HI data).

In other words, pressure on the service will continue and will most probably ramp-up with the HI run and beyond. In addition, plans for running over the Christmas period will have to be developed, most likely including also Tier2 sites (as analysis related work must be expected to continue during this period.)

In response to the various database-related service incidents in recent months, ATLAS is reviewing its strategy for database services and Tier1s and is expected revise its deployment model in the coming months. This is likely to include the use of databases for conditions data at fewer Tier1s with more reliance on FronTier / Squid.

Summary and Conclusions

Not withstanding the continued service incidents and in particular those that took many days to resolve, the service has continued to meet the challenges of LHC exploitation and usually met or exceeded targets for response to and resolution of problems. Good press was received – such as in the Economist article of July 29th – which commend the many years of preparation and on-going work on service delivery. However, there continue to be a number of major issues to address, such as improved response to network incidents, streamlined response to major site downtimes and indeed the handling of security events. Despite these problems, presentations by the experiments at CHEP 2010 in October in Taipei emphasized the significant amount of work that they had been able to perform on the grid. Improvements in the time taken to resolve problems are clearly desirable but in the meantime the experiments are able to workaround these problems, albeit with the (sometimes significant) loss of resources.

Grid Deployment Board Report
July – October 2010

Nov 2010

John Gordon

Introduction

There were only two meetings of GDB held during these four months; the July meeting being replaced by a WLCG Workshop in London and the August one cancelled due to holidays. The agendas show the topics covered and the slides and papers. http://indico.cern.ch/categoryDisplay.py?categId=3l181

Long-Running Issues

Some issues appear regularly at the GDB because progress in reaching closure has been very slow or because they are recurring themes. Among these were:

CREAM: A problem with Condor retrieving outputs held up production use for some time. Fixed in July and serious testing undertaken by ATLAS. Looking in better shape.

Multi-User Pilot Jobs: seemed to have reached a deadlock where sites are not actively pushing for this although most responses to a questionnaire were that they wanted identity changing. Thus experiments have not given much priority to the production testing of glexec etc.

Reliability of Sites: Experiments have come to realise that sites will never be 100% reliable and any computer model which cannot cope with this efficiently is doomed. There is still much that sites can do though in sharing experience in solving reliability issues and improving resilience through redundancy.

Experiment Operations: experiments report quarterly on their hottest operational issues.

Topical Issues

These are issues which have arisen approximately during the period of this report.

Handling of network problems: there seems to be limited transparency to end users when WAN network problems are investigated. Someone needs to take ownership of a problem to push information around to those concerned. Work is starting on reviewing the procedures to ensure that communication traverses the whole chain.

CERNVMFS: This distributed filesystem which pulls files from a central place on demand, was developed to build virtual machine images. It is being tested as a vehicle for experiment software distribution and the results look encouraging. It uses web caching to reduce network traffic by only downloading on demand and by removing file duplication both inside and between releases.

Data Access and Management Demonstrators: the June Jamboree which considered ways to improve the methods in which LHC experiments transfer and use data, spawned about a dozen demonstrators deserving of further consideration. These were encouraged to develop further and demonstrate some innovation by the end of the year. There was a checkpoint in October and some common themes emerged around xrootd, catalog syncing, caching and NFS4.1.

Applications Area Report
July – October 2010

November 2010

Pere Mato

In the last quarter AA has prepared 8 different releases of the common software stack for LHC experiments. The changes within these releases include major upgrades of the AA developed projects (ROOT, POOL, COOL, CORAL, RELAX) and at least one upgrade of 26 "external packages" (i.e. software developed outside AA). Currently 14 different software platforms are being supported and recently the complete software stack was ported to the latest Mac OSX 10.6 on 32 and 64 bit architectures. Python continues being one of the cornerstones of AA and LHC software with an ever-increasing number of modules. This huge success of the language required a new deployment method to be implemented for the currently 23 supported Python modules.

ROOT

In view of the coming production release 5.28 in December an overall campaign of code correction and consolidation in response to the reports generated by the Coverity static code analysis tool has taken place. This great tool is able to report very special cases that may happen with a very small probability. We have been actively working in fixing the vast majority of reports, in particular the ones flagged as "high impact".

Finalized the support for member-wise streaming of collections and made it the default, the files using the new format are between 2% and 10% smaller and can be read 12% to 30% faster. Investigation of the feasibility of a new C++ interpreter based on the industry standard C++ compiler Clang is on going and show good promises. Significant increase in the number of new groups setting up PROOF-based analysis farms, mostly from ATLAS and CMS. This also increased the request for user support. The RooStats package has been greatly improved taking into account the new requirements from the LHC data analyses. For example it has been used in a prototype analysis for combining Higgs results of ATLAS and CMS.

Persistency Framework

Three new releases of the Persistency Framework projects have been prepared motivated by fixes to problems reported by experiments. One important issue in Oracle database services was also analyzed and solved. The cause has been an incompatibility between various versions the gssapi library used by Oracle, Globus and Xerces.

Server-side process crashes triggered by COOL applications were observed on the ATLAS and LHCb databases, after the Oracle security updates were applied in June. The team prepared a COOL-based stress test suite that was successful in reproducing the issue on a test database and this has also been used to validate the lastest patch proposed by Oracle.

Simulation

New Beta release of Geant4, 9.4-beta, was announced as scheduled on June 25th. Among the developments included are a new revised Bertini cascade hadronic model, new physics-lists configurations utilizing improved modeling for anti-baryons and hyperons (using the CHIPS hadronic model instead of parameterised LEP/HEP models). Improvements in memory management for issues also reported by performance monitoring teams in ATLAS and CMS, and a new trapezoid shape, G4GenericTrap, requested by ALICE. Additionally two patch-releases, 9.2.p04 and 9.3.p02, were recently distributed, including also fixes for few issues affecting the releases currently in use by the LHC experiments.

Good progress has been made in understanding the effects of models transition on energy resolution in Geant4. The simulation quality of new physics configurations combining the Fritiof model and the Bertini intra-nuclear cascade, has improved considerably, with rather positive feedback from validation studies made by ATLAS and CMS.

The regression testing suite for MonteCarlo Generators has been renewed to be now based on distributions and using tools like Rivet and the HepMC-analysis tool. Good progress has been also achieved in the realization of a new web interface for the tuning and validation of Generators against the latest available public data from the experiments.

ALICE Report
July – October 2010

Nov 2010

Yves Schutz

Software

The main issue of excessive memory consumption by the reconstruction task is progressively coming under control. Data size and CPU power for pp reconstruction are according to the values foreseen in the Computing Model. CPU power for Monte-Carlo processing is still larger than the values assumed in the Computing Model. No substantial reduction is foreseen in the near future.

Operation

All Data collected by ALICE (800 M of minimum bias triggers and 90 M of rare triggers plus various millions of cosmic and calibration events) have been reconstructed in a first pass quasi on line at the CERN T0. Pass 2 has been processed with improved algorithms and condition parameters in external T1 sites on the raw data replicated outside of CERN. The reconstructed events are replicated four times in Tier1s and Tier2s for analysis. Several analysis trains are run regularly over the entire set of reconstructed data. End users analyses are actively pursued on the GRID and on the PROOF-enabled analysis facilities. All CPU power available to ALICE in T0, T1s and T2s is used for raw data reconstruction, Monte Carlo production and analysis tasks. About 50% of the installed disk storage is used so far. The running mode including two CASTOR disk buffers at T0 has been exercised and validated for the maximum rates of the HI run.

Monte-Carlo production

About 80% of the consumed CPU resources are used for Monte Carlo production. In the Computing Model this number is 55%.

Grid services

The efficiency of the GRID is now mainly depending on the availability of the many SE elements. This is mainly an issue related to the availability of human resources with adequate competence in the sites and centrally at CERN.

Changes

The computing resources required in 2013 will be calculated using the real parameters rather than those from the Computing Model.

Concerns

In view of the already large amount of resources used so far mainly for Monte-Carlo processing, the resources pledged for ALICE might be barely sufficient until the end of year and definitively insufficient for the AA data processing in 2011 in preparation of the QM conference in Mars 2011.

ATLAS Report
July – October 2010

November 3, 2010

Kors Bos

ATLAS took data and successfully did its calibration, alignment and first pass processing in the Tier-0. No re-processing of the data in the Tier-1's was attempted. The software release in the Tier-0 was frozen to the one that was also used for the May re-processing so all data was consistently processed through the same software and could be used for the summer conferences. There was a significant increase in the analysis activities in the Tier-2's but because there was also no attempt to generate a new sample of simulated data the Tier-1's stayed underused.

A major issue was the LFC outage at SARA that took much longer than expected: 3 weeks. Apart from the broken hardware it turned out to be impossible to rebuild the database from the backup even when other hardware was used. It showed that a procedure must be defined by which re-builds are regularly tested to still function. It also showed that the LFC is one of the most critical services of a Tier-1 and without it the whole cloud is dead. Steps have been defined to work towards a live backup database in another site.

In the whole quarter 604 ggus team tickets (normal shifter tickets) were issued, 14 of which were alarm tickets, 10 of those were related to storage issues. Out of these 14 alarms CERN had 7, 5 related to storage and 2 to LSF, of the other 7 alarms 2 were related to LFC and 5 to storage, all in different Tier1s.

A new way of data access for analysis was started in July in the US cloud: an analysis task is launched in the Tier-1 and the first jobs will start there locally. Meanwhile the Panda system realizes which other files need to be accessed and sends those to a Tier-2 site where free CPUs are available. When the data has arrived it launches the corresponding analysis jobs. This is caching on demand instead of using pre-staged data. When this was shown to work for 2 weeks the pre-placement of data from the Tier-1 to the Tier-2's was stopped without any noticeable effect on the analysis efficiency. Initially this was done for ESD and DESD data formats but it now has been extended to also other types. Moreover the Italian, the French and the Dutch cloud now also use this Panda Dynamic Data Placement PD2P within their clouds.

CMS Report
July – October 2010

Nov 2010

Ian Fisk

CMS Experiences

This quarter began with the final preparations, processing passes, and analysis activities for ICHEP and completed with nearly the end of the 2010 proton-proton running and preparation for heavy ions. CMS Computing generally performed well during this intense period. The Tier-0 infrastructure continued to run reasonably smoothly. The CMS data acquisition rate was higher than the design during much of the quarter as CMS developed a better understanding of the trigger and selection methods. The lower level of pile-up in these events means the reconstruction time is also shorter than the design parameters, which allowed the Tier-0 to accommodate the higher rates. The trigger group was conscientious to stay below rates that would overwhelm the offline system. Toward the end of the quarter the machine achieved a higher average time in stable beams and the average utilization of the Tier-0 processors increases accordingly. Even with higher instantaneous luminosity and a higher trigger rate CMS maintained close to the design rate into the Express stream. In CMS the express stream is expected to be fully reconstructed within an hour of data collection. The Tier-0 successfully maintained this rate even with higher luminosity and progressively more interesting events.

During this quarter CMS migrated to a new software version for the online, Tier-0 offline, Tier-1 reprocessing and analysis systems. The transition was reasonably smooth due to the validation of the Offline group in CMS, and the code has a variety of physics and technical improvements. The optimized IO is particularly noticeable in the computing facilities with significant improvements in the achieved CPU efficiency clearly visible in the accounting reports.

The Tier-1 centers were well exercised during this period in all the expected capacities. The Tier-1 centers performed two large-scale reprocessing passes on the data in preparation for ICHEP. One pass was to compensate for a recently identified detector effect and was prepared and completed within a few days. During the early Fall the Tier-1 centers participated heavily in Monte Carlo production and simulation reprocessing. While the Tier-2s performed the majority of the simulation the extra production capacity between data reprocessing passes from the Tier-1s allowed CMS to maintain an aggressive schedule of nearly 1.4B events to produce during the second half of the year. The average CPU utilization of the Tier-1 centers increased during July, August, and September.

In addition to producing simulated events centrally for CMS, the Tier-2 centers continued to be the primary location for analysis. The number of jobs and users ramped up in the preparation for ICHEP and continued at high levels with a small reduction during August. The average number of individuals submitting to the Tier-2s in a week is around 400 with more than 800 individuals submitting each month. CMS reaches around 100k grid submissions per day, which is around the computing model expectations.

The third quarter of 2010 was exciting due to the increasing volume of data and sustained analysis and organized processing activities. This environment gives us some insight into the challenges expected in 2011.

Nov 2010

LHCb Report
July-October 2010

Peter Clarke

Data processing context:

- In Q3 the luminosity has risen by several orders of magnitude. Our HLT accept rate has been tuned to remain essentially constant. However our event size has increased from an anticipated 35kB to 75kN for the high μ running. This resulted in larger CPU times and necessitated significant actions early in the quarter. Our reconstruction code is now at twice the anticipated level.
- Processing on the Grid has remained overall satisfactory. Data access instabilities at T1s continue to be the main cause of delays to processing.
- We continue to run a three step process on the Grid, i.e. Express data processing for quality checks ➔ Full processing ➔ Merging of the output DSTs ➔ Replication.
- Again due to the high μ, a larger DST retention rate and a larger DST event size (130 kB instead of 85 kB) has necessitated careful management of disk space. We expect to fully utilize our entire disk pledge by the end of 2011, but will need to limit MC and User space to achieve this.
- We are about to commence a complete reprocessing which we expect to complete in December. We will shortly launch a new MC production campaign.

Particular issues of note are:

- **Shared area: Unchanged:** As reported previously we have continuing software repository problems: lack of stability or scalability of the software repository (usually NFS mounted disk). This is a permanent concern and this service should be really considered seriously as a critical service. In particular it is important to isolate VOs for each other's activity.

- **LHCb running conditions:** The success of LHC in terms of luminosity, combined with the likely continued running at high μ means that our CPU and disk requirements are being stretched.

LHCb usage patterns are shown in the figures below.

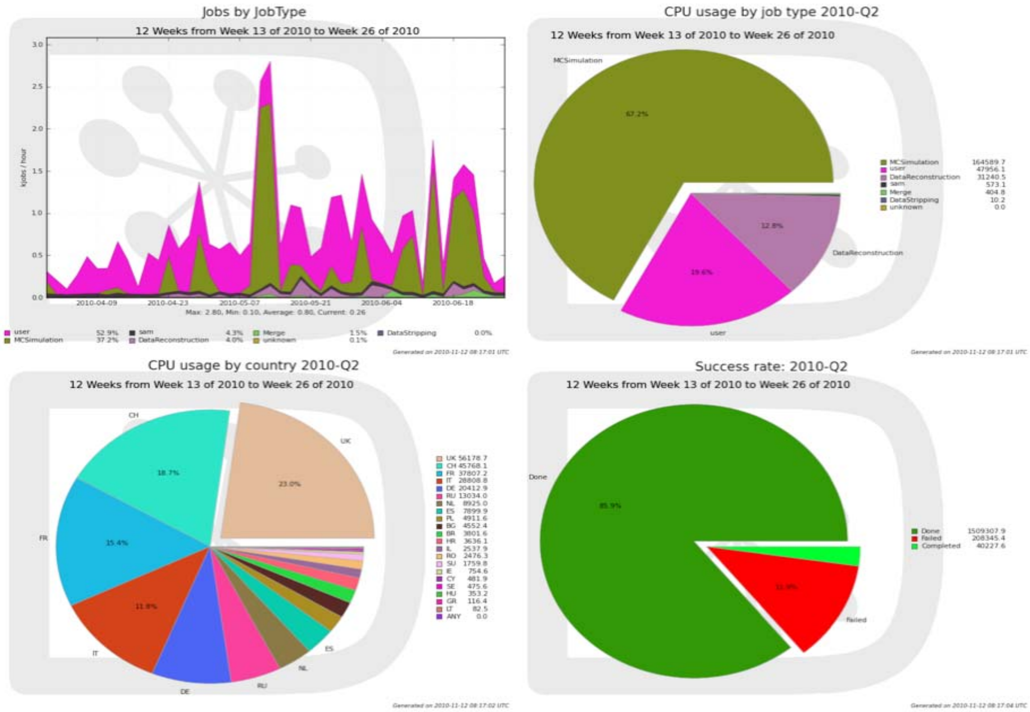
Jobs by JobType
12 Weeks from Week 13 of 2010 to Week 26 of 2010
kjobs / hour
2010-04-09
2010-04-23
2010-05-07
2010-05-21
2010-06-04
2010-06-18
Max: 2.80, Min: 0.10, Average: 0.80, Current: 0.26
user 52.9%
MCSimulation 37.2%
sam 4.3%
DataReconstruction 4.0%
Merge 1.5%
unknown 0.1%
DataStripping 0.0%
Generated on 2010-11-12 08:17:01 UTC
CPU usage by job type 2010-Q2
12 Weeks from Week 13 of 2010 to Week 26 of 2010
MCSimulation
67.2%
19.6%
12.8%
DataReconstruction
user
MCSimulation 164589.7
user 47956.1
DataReconstruction 31240.5
sam 573.1
Merge 404.8
DataStripping 10.2
unknown 0.0
Generated on 2010-11-12 08:17:01 UTC
CPU usage by country 2010-Q2
12 Weeks from Week 13 of 2010 to Week 26 of 2010
CH
UK
FR
IT
DE
RU
NL
ES
23.0%
18.7%
15.4%
11.8%
UK 56178.7
CH 45768.1
FR 37807.2
IT 28808.8
DE 20412.9
RU 13034.0
NL 8925.0
ES 7899.9
PL 4911.6
BG 4552.4
BR 3801.6
HR 3636.1
IL 2537.9
RO 2476.3
SU 1759.8
IE 754.6
CY 481.9
SE 475.6
HU 353.2
GR 116.4
LT 82.5
ANY 0.0
Generated on 2010-11-12 08:17:02 UTC
Success rate: 2010-Q2
12 Weeks from Week 13 of 2010 to Week 26 of 2010
Done
85.9%
Failed
Done 1509307.9
Failed 208345.4
Completed 40227.6
Generated on 2010-11-12 08:17:04 UTC