

ATLAS Spanish Tier2 experiences during STEP09 period

Presenter: Santiago González de la Hoz

IFIC (Instituto de Física Corpuscular), Centro mixto CSIC-Universitat de València

*¹GONZALEZ, S., ¹SALT, J., ¹AMOROS, G, ¹FERNANDEZ, A., ¹KACI, M.,
¹LAMAS, A., ¹OLIVER, E., ¹SANCHEZ, J., ¹VILLAPLANA, M., ¹VIVES, R.,
²BORREGO, C., ²CAMPOS, M., ²NADAL, J., ²PACHECO, A., ³DEL PESO, J.,
³PARDO, J., ³GILA, M., ³FERNANDEZ, P., ⁴ESPINAL, X.*

¹IFIC-Valencia, ²IFAE-Barcelona, ³UAM-Madrid, ⁴PIC-Barcelona

- **Spanish Cloud**
 - Spanish Resources
- **PIC-Tier1 (Report thanks to X. Espinal)**
 - ATLAS Services
 - Data transfers
 - Reprocessing
 - User Analysis
- **UAM-Universidad Autónoma de Madrid (Report thanks to J. Pardo & P. Fernández)**
 - Computing: Faire share, cpu load
 - Storage and Bandwidth
- **IFAE-Institut de Física d'Altes Energies, Barcelona (Report thanks to J. Nadal)**
 - Computing
 - MC production
 - Difficulties
 - Distributed Data Management
- **IFIC- Institut de Física Corpuscular, València (Report thanks to J. Sánchez)**
 - Job processing
 - Data Transfer
- **STEP09: towards data taking**
 - Issues achieved and learned
 - Status and plans



SWE Cloud:

Spain-Portugal

Tier1:

PIC-Barcelona

Tier2's:

UAM, IFAE & IFIC

LIP & Coimbra

- Tier1 at PIC Barcelona

- Offers storage and processing resources for three LHC experiments: ATLAS, CMS and LHCb.
- LHC experiments will store a copy of the collected data from the accelerator at CERN and dispatch a secondary copy to the Tier-1s centres in order to guarantee the conservation and integrity of the data.
- ~10% of the raw data from the LHC accelerator will be stored at PIC.
- Optical Private Network (OPN) Tier0 (CERN) ↔ Tier1's.
- More than 9 PetaBytes in/out PIC in 2008.

- Tier1 will provide the infrastructure for **data re-processing**, as the **raw data stored** will be reprocessed several times per year with new parameters, as calibration and alignment constants improve.

		2007	2008	2009	2010	2011	2012	2013
CPU (kSI2K) required	ATLAS	172	865	1226	1960	2687	3417	4872
	CMS	289	477	1058	2516	3292	4099	6201
	LHCb	37	167	307	633	962	1215	1263
	TOTAL	498	1509	2591	5109	6941	8731	12336
Disk (Tbytes) required	ATLAS	114	512	902	1595	2168	2743	4176
	CMS	79	358	630	1113	1513	1915	2915
	LHCb	21	97	170	301	409	518	788
	TOTAL	214	967	1702	3009	4090	5176	7880
Tape (Tbytes) required	ATLAS	68	385	681	1182	1767	2439	2819
	CMS	140	487	974	1677	2519	3358	5186
	LHCb	18	81	189	543	963	1456	2981
	TOTAL	226	953	1844	3402	5249	7253	10986

Installed

Planned

September 09

- Data Storage:
 - Experiments do need **large, reliable and scalable storage services**.
 - To server the data at the required speed in order to maximize the efficiency of the cluster.
 - **Multi-Gigabit Ethernet network architecture**, specially designed to enhance high **speed data movement between WAN** (Tier0, Tier1s, Tier2s) and **LAN** (CPU farm) .
 - **dCache storage system**.

• ATLAS Spanish Federated Tier2

- IFIC: Valencia (coordinator)
- IFAE: Barcelona
- UAM: Madrid


Ramp-up of Tier-2 Resources (after LHC rescheduling) numbers are cumulative

Evolution of ALL ATLAS T-2 resources according to the estimations made by ATLAS CB (Oct.06)

Year	2006	2007	2008	2009	2010	2011	2012
CPU(KSI2k)	925	2336.11	17494.51	26972.76	51544.64	69128.42	86712.2
Disk (TB)	289	1259.04	7744.37	13112.04	22132.3	31091.45	40050.92

Spanish ATLAS T-2 assuming a contribution of a 5% to the whole effort

Year	2006	2007	2008	2009	2010	2011	2012
CPU(KSI2k)	46	117	875	1349	2577	3456	4336
Disk (TB)	14	63	387	656	1107	1555	2003



Strong increase of resources

Present resources of the Spanish ATLAS T-2 (August'09)

Hardware is purchased. Working next week to get the pledged resources

	IFAE	UAM	IFIC	TOTAL
CPU (ksi2k) (HEP-SPEC06)	201 800	338	96	635
Disk (TB)	215	165	34	414

Accounting values are normalized according to WLCG recommendations

- Storage Element System**

	SE (Disk Storage)
IFIC	Lustre+StoRM
IFAE	dCache/disk+SRM posix
UAM	dCache





- StoRM: Posix SRM v2
- Lustre: High performance standard file system
- Shares: 50% IFIC, 25% IFAE and 25% UAM

Data (AOD) distribution and DDM FT continuously running from Tier1 to Tier2

- A Tier needs a **reliable and scalable storage system** that can hold the users data, and serve it in an efficient way to users.
- A first sketch of a Storage system matrix (evaluation of **different systems on going at CERN**):



Storage System	Local Protocol	Load Balancing	Externally Secure	POSIX Access	Single Namespace	Installation Load	Maint Load	Quotas	Cost
NFS	bad	N	N	Y	N	low	high	Y	\$0
Lustre	Y	Y	w/SRM	Y	Y	medium	medium	Y	\$0
GPFS	Y	Y	w/SRM	Y	Y	high	medium	Y	\$\$\$
xrootd	Y	Y	w/SRM	mkdir/rmdir do nothing	Y	medium	low	partitions	\$0
DPM	Y	Y	Y	special commands	Y	medium-high	low- medium	partitions	\$0
dCache	Y	Y	Y	metadata	Y	high	low- medium	partitions	\$0

- ATLAS Services Status: LFC**




Step09 period

- **ATLAS Services Status: pilot factories**
 - Two pilot factories running at PIC VObox feeding our T1s and T2s
 - MC Production (production role)
 - User Analysis (pilot role)
 - Increased pilot pressure to constantly fill all nodes doesn't affect performance:



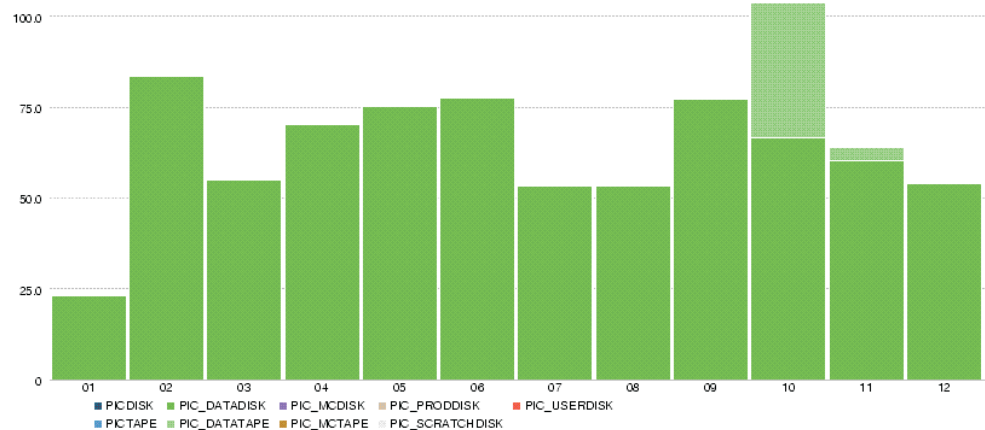
Step09 period



- Constantly queuing 30 pilots at all the sites during STEP09.
- Lowered to 15-20 during low pressure dates.

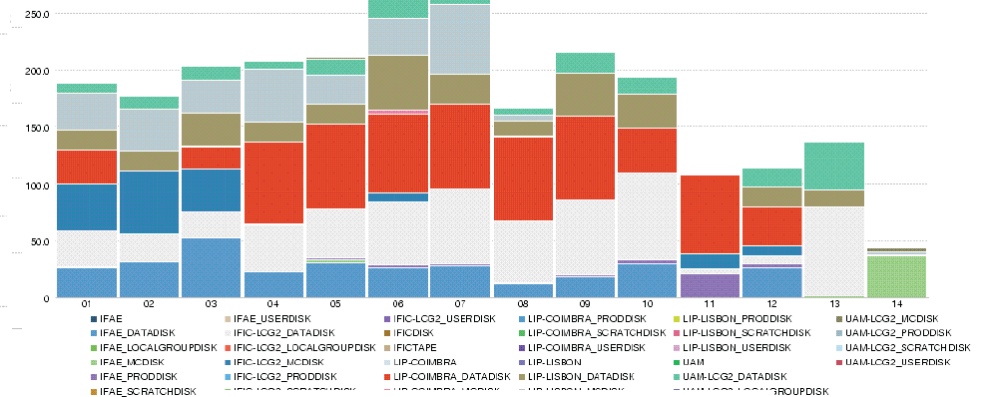
STEP09 Data Transfers

Tier-0 to PIC throughput (MB/s)




Reached > 250 MB/s with excellent efficiencies (95% to 100%)


PIC-Tier-2s data transfers (MB/s)



Tier1-Tier2s Data Aggregation



Tier1-Tier1 Data Aggregation




STEP09 reprocessing

- Main STEP target and dominant activity at the Tier-1s
 - Read ALL files from tape and store outputs to tape
- Repro workflow fully steered by PanDA (pilots) and DDM
- Metric of success during 10 days:
 - Reprocess at x5 data taking (assuming 40% LHC effc.)
 - PIC (5%T1): complete 16 TB and 10k files
- Gold metric:
 - Reprocess at x5 data taking (assuming 100% LHC eff.)
- PIC completed almost two full repro cycles: gold star ★

Objectives was to reprocess faster nominal data taking rate of 200Hz


- Baseline metric at 400 Hz (x2)
- Enhanced metric at 1000 Hz (x5) ★




TI	Base Target	Result	Comment
ASGC	10 000	4 782	Many batch system and basic setup problems
BNL + SLAC	50 000	99 276	Also ran high priority validation and other tasks
CNAF	10 000	29 997 ★	
FZK	20 000	17 954	Big tape system problems pre-STEP; no CMS
LYON	30 000	29 187	Very late start due to tape system upgrade, then good
NDGF	10 000	28 571 ★	
PIC	10 000	47 262 ★	
RAL	20 000	77 017 ★	
SARA	30 000	28 729	Tape system performance very patchy
TRIUMF	10 000	32 481 ★	Also ran high priority validation and other tasks

- User Analysis (WMS job submission and Pilot jobs)
 - Stats was corrected for Upstream Panda Failure:
 - Build job failing at sites were not accounted, includes:
 - 32/64 bits compact issues
 - Releases not/wrong installed

CLOUD	SUBMITTED	RUNNING	COMPLETED	FAILED	Efficiency	# Files
CA	0	0	5395	2344	0.70	24361
CA_PANDA	0	0	26715	7504	0.78	62756
DE	0	0	52478	19051	0.73	342093
DE_PANDA	0	0	79625	25802	0.76	215302
ES	0	0	14038	3317	0.81	95347
ES_PANDA	0	0	48075	7334	0.87	141343
FR	0	0	26694	5411	0.83	186601
FR_PANDA	0	0	116934	17516	0.87	375377
IT	0	0	38428	4758	0.89	266703
IT_PANDA	0	0	14036	2343	0.86	44358
NG	0	0	14551	2919	0.83	20179
NL	0	0	10703	20257	0.35	64016
NL_PANDA	0	0	25624	10119	0.72	90436
T0_PANDA	0	0	0	1151	0.00	0
TW	0	0	784	770	0.50	7663
TW_PANDA	0	0	18675	4140	0.82	78630
UK	0	0	21695	15190	0.59	120972
UK_PANDA	0	0	121975	23000	0.84	318112
US	0	0	153609	9770	0.94	467732
TOTAL	0	0	790034	182696	0.81	2921981




CLOUD	SUBMITTED	RUNNING	COMPLETED	FAILED	Efficiency	# Files
CA	0	0	32110	9848	0.77	87117
DE	0	0	132103	44853	0.75	557395
ES	0	0	62113	10651	0.85	236690
FR	0	0	143628	22927	0.86	561978
IT	0	0	52464	7101	0.88	311061
NG	0	0	14551	2919	0.83	20179
NL	0	0	36327	30376	0.54	154452
T0	0	0	0	1151	0.00	0
TW	0	0	19459	4910	0.80	86293
UK	0	0	143670	38190	0.79	439084
US	0	0	153609	9770	0.94	467732
TOTAL	0	0	790034	182696	0.81	2921981



- Computing: Faire share (FS)

Queue on CE



5th June
(Wrong FS)

Group	Running Jobs	Queued Jobs	%FS
Production	31	1	40.42
User analysis (wms)	11	199	41.81
User analysis (pilot)	144	0	17.72

8th June
(Right FS)





Group	Running Jobs	Queued Jobs	%FS
Production	48	29	49.76
User analysis (wms)	110	57	25.27
User analysis (pilot)	26	90	24.92

- Computing: Faire share

Production


Analysis (Pilot based)

Analysis (WMS submission)




- Computing: CPU Load

CE



All WNs had a huge load
but the CE worked
without problems

WN




- Storage and Bandwidth: Previous transfers**

25 TB from 22nd to 25th May

During the test, MCDISK
was increased barely 500
GB on the fly

PNFS load was high,
causing some
random SAM checks
to fail (timeouts)

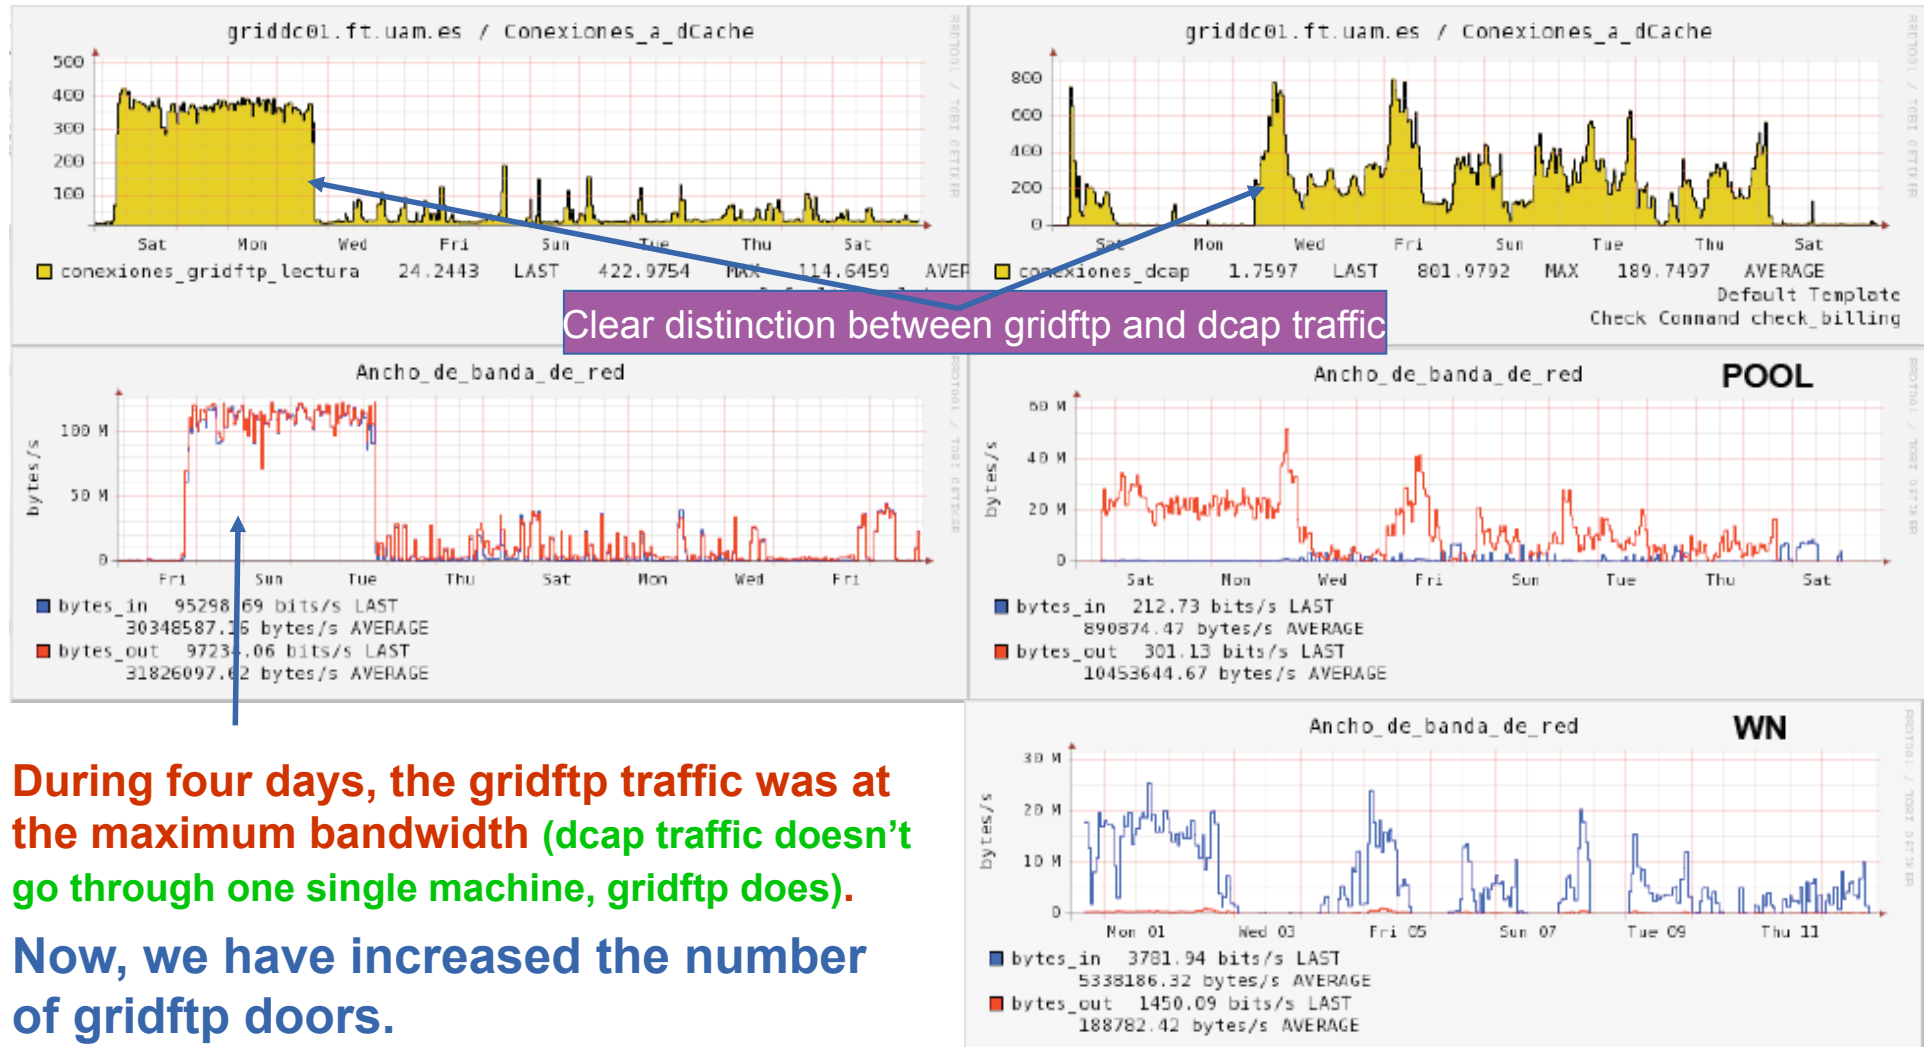
External bandwidth
was very busy



- Storage and Bandwidth: During Step09 test

gridftp


dcap



During four days, the gridftp traffic was at the maximum bandwidth (dcap traffic doesn't go through one single machine, gridftp does).

Now, we have increased the number of gridftp doors.

- **Computing**



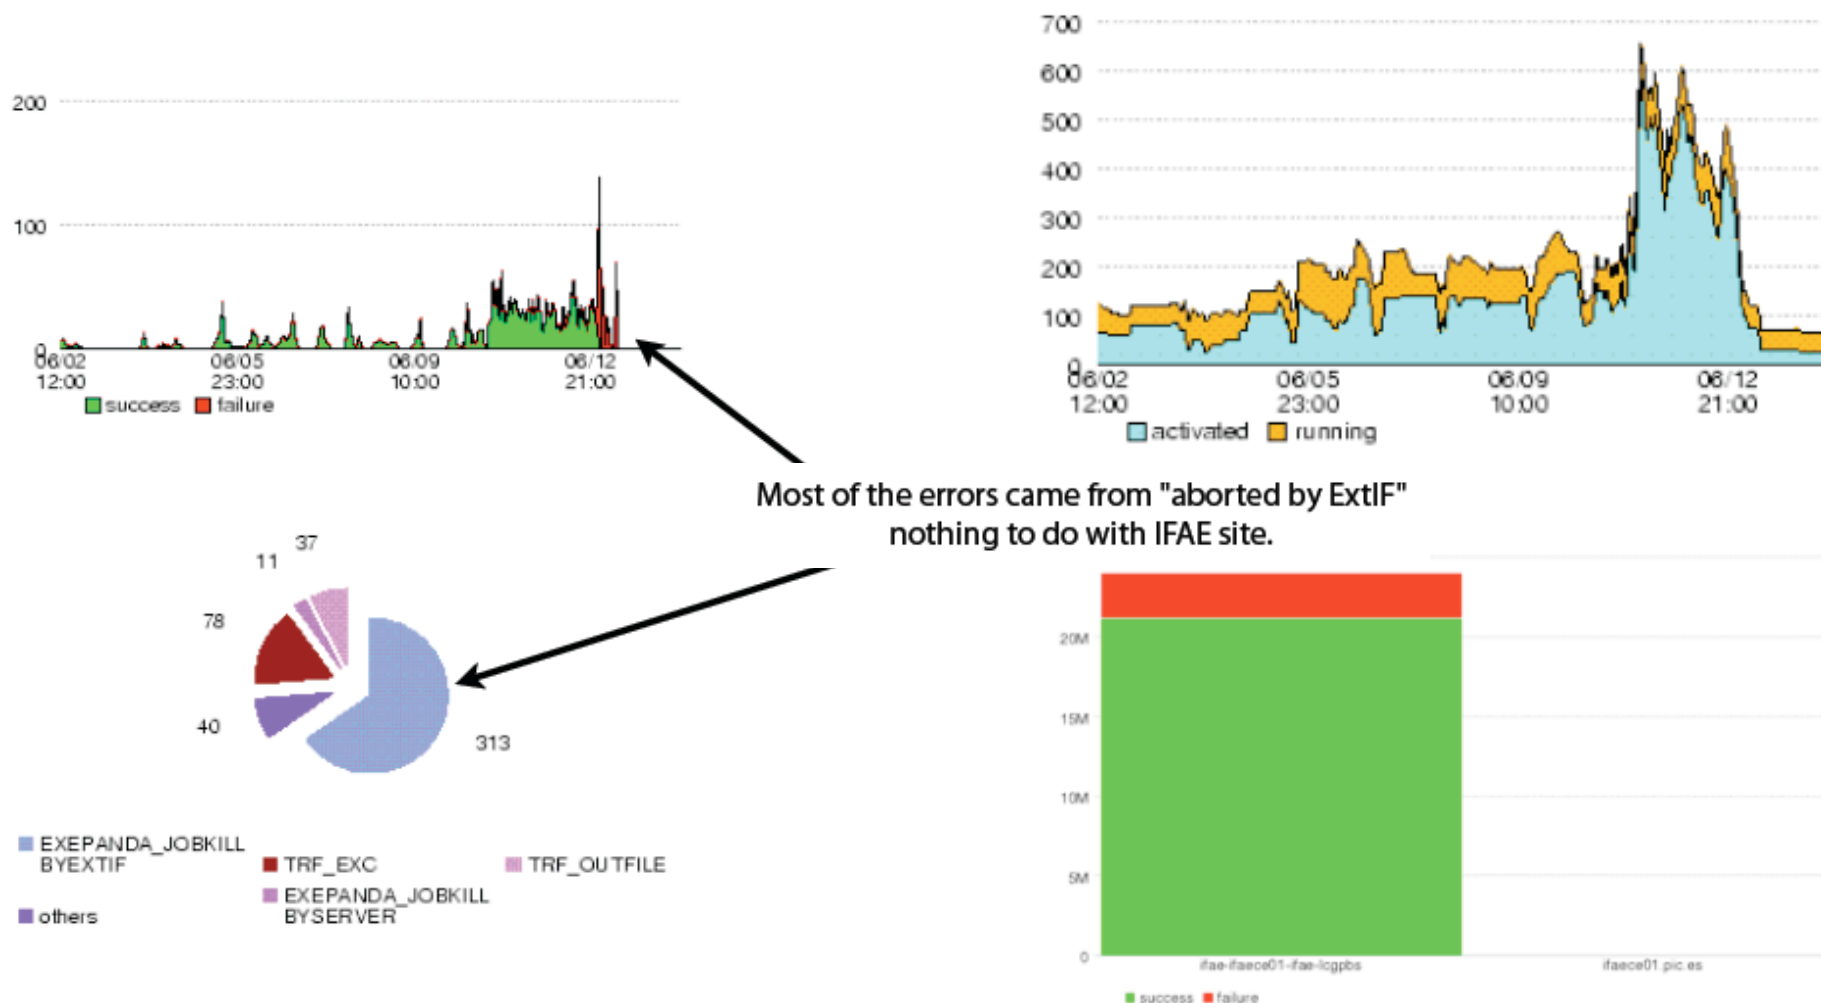
Blue: role value for user analysis (25% analysis pilot based)

Red: WMS submission (25% analysis)

Green: Official MC production


The share imposed in the IFAE's pbs was being accomplished.

- MC production during STEP09:
 - Good achievement in the efficiency 92.5%




Difficulties:

Number of jobs per day prd/user/pilot during STEP09



Issue with the faire share during the first days;
Scheduler's historical information should have been removed:

- The farm was mainly dominated by Panda jobs (in blue), queuing all others.
- Some error from the WMS jobs were found due to the time outs; This jobs were waiting for a long time




Some errors with some very “weird” jobs were experienced. They were not able to ran but just blocking the queues.


But at the end IFAE got a pilot analysis efficiency around 90% and 76% for the WMS jobs

• Distributed Data Management:


- 20 TB were placed.
- Green lines shows the available space of the MCDISK token



9th June more than 400 Mbytes/s were reached



9th June DATADISK token was running out of space. Increased with 10 TB. This error lasted during 3 hours and was detected from Atlas DDM monitoring



- **Distributed Data Management:**
 - Data moved from/to IFAE these days and also the efficiency which was affected by the errors:
 - Space problem caused some transfer errors.

	Cloud	Transfers			Registrations		Transfer
		Efficiency	Throughput	Successes	Datasets	Files	
-	IFAE_DATADISK	65%	23 MB/s	10425	2595	8401	5540
	PIC_DATADISK	65%	23 MB/s	10425			5540
+	IFAE_LOCALGROUPDISK	0%	0 MB/s	0	0	0	0
-	IFAE_MCDISK	100%	3 MB/s	2334	19	2332	0
	PIC_MCDISK	100%	3 MB/s	2334			0
-	IFAE_PRODDISK	87%	3 MB/s	2225	255	2200	332
	PIC_DATADISK	100%	0 MB/s	415			0
	PIC_DATATAPE	84%	3 MB/s	1697			332
	PIC_MCDISK	100%	0 MB/s	113			0


- Job processing

- Pilot role was added in VOMS mapping
 - But **no proper scheduling policy could be setup in time**
 - In order to achieve this, **a new account pool had to be setup but it was not easy with the present configuration tool (*quattor*).**

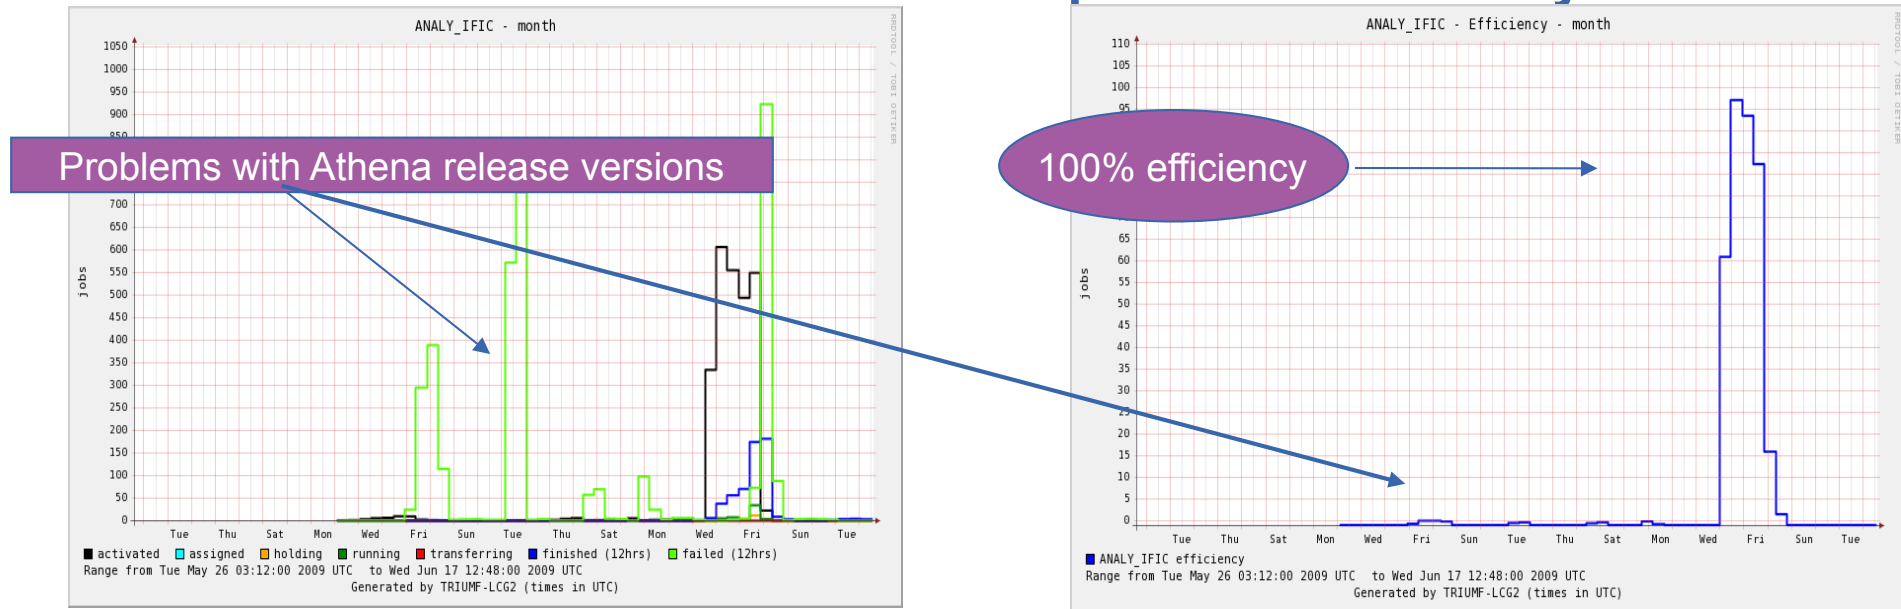
Group	%	Target
Atlas users(WMS and pilots)	43.23	50
Atlas production	56.49	50

- Farm occupancy was high and we added 40+40 cores more in last week (24)

The number of cores we added stressed the CPU power at the CE (Pentium D CPU 3.20 GHz) and we decided to upgrade it in that coming weeks.



- Job statistics from ANALY and IFIC queue from 17th May to 17th June.



- In the first week we had problems with **panda jobs** trying to use an **non existing Athena version**
- For that reason, our **efficiency for pilot jobs was around 21% (ANALY_IFIC)** while for **WMS jobs (IFIC-LCG2_MCDISK)** was around **82%** on 17th June

ES	SUBMITTED	RUNNING	COMPLETED	FAILED	Efficiency	#Files Expected	#Files Processed
IFIC-LCG2_MCDISK	62	96	1130	299	0.82	8105	8105
ANALY_IFIC	443	109	455	1691	0.21	1506	?

Statistics taken from: <http://gangarobot.cern.ch/st/status.html>

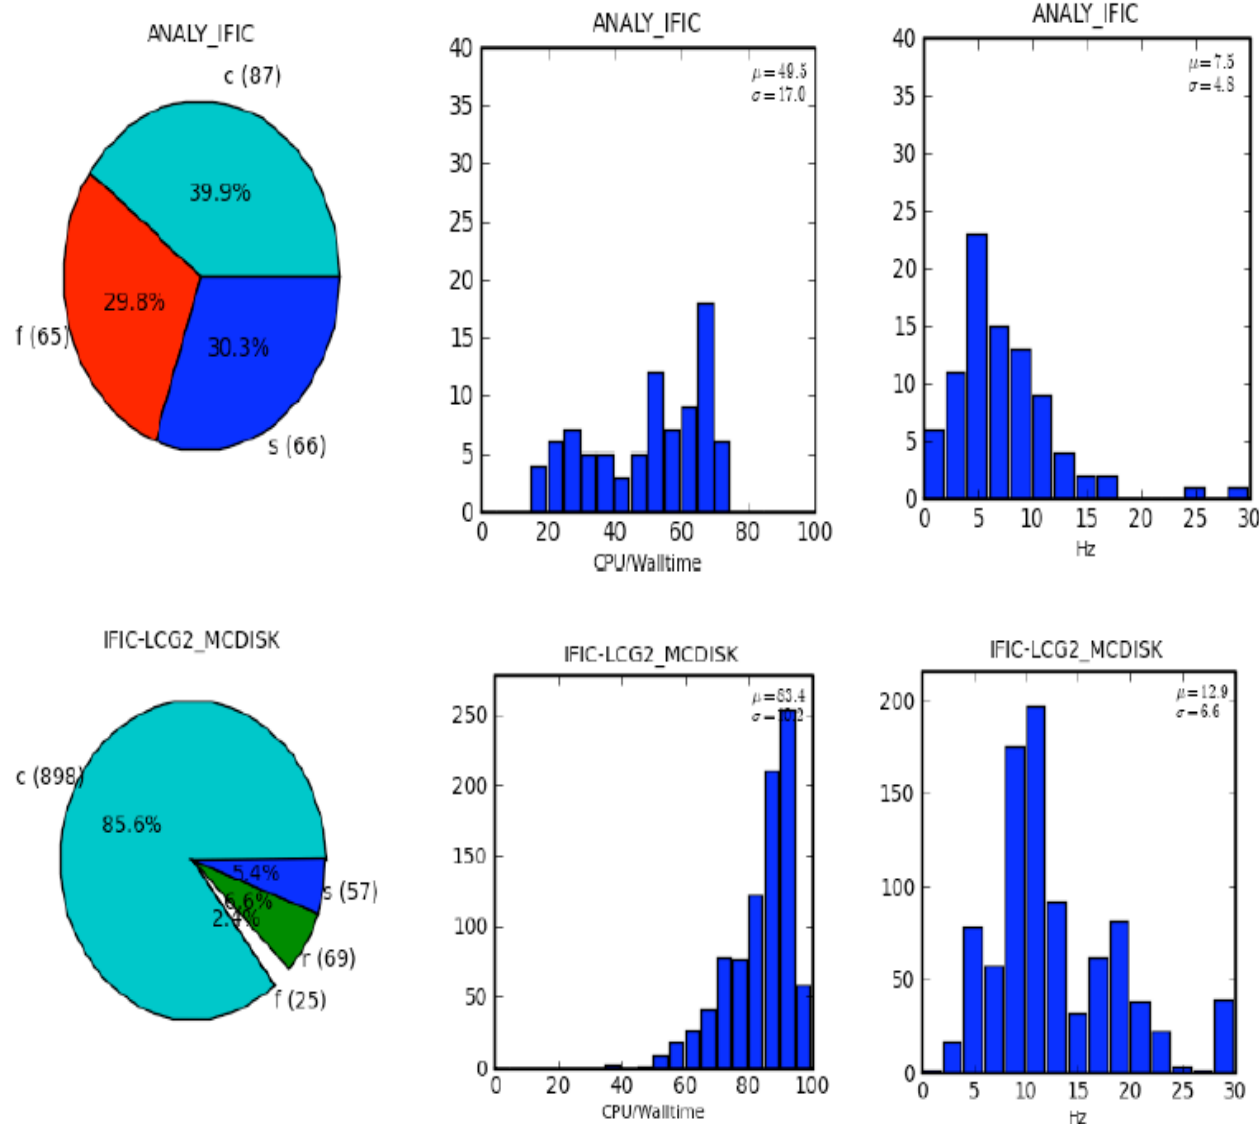
After installing the Athena versions used for panda our pilot job efficiency raised to 17% in a few minutes being around **100% in the last days.**

Influence on the **input data protocol** on **CPU/Wall time efficiency** was tested in previous UA tests.

Sites with **POSIX-like access** to the file system (Lustre) perform better **without triggering copies** in the background using lcg tools (file stager)


- The same results was found in STEP09 test for WMS jobs but not for PANDA pilot jobs, CPU/Wall time efficiency was reduced from 83.4% to 49.5%.

— Panda pilot based analysis used your site's ANALY queue local mover to copy data to the worker node; the WMS based analysis used a mixture of the file stager (copy to worker node) and local LAN protocol (known as DQ2_LOCAL, resolved to dcap, rfiio or xroot).



- **Data transfer**

- Data distribution in the week previous to the STEP09 showed a **bottleneck in the gridftp server**.
 - We deployed a new machine and the bandwidth increased **from 200 Mbps to 600 Mbps**.
 - During the test **data transfer rate was steady during long periods**. No problems were observed in the network infrastructure neither local or remote.
- WN accessed to the data using local posix I/O (read only) and gridftp protocols.



WNs and disk servers are connected to Cisco switch at 1Gbps.

In the test configuration:

WNs share 1 Gbps in groups of 8

Disk servers share 1 Gbps in groups of 2.

Present configuration:


1 Gbps for every WN

2 Gbps for every disk server

- From June 1st to 14th data size transferred to our SE was:
 - 57 TBytes incoming (112TB total, 50% to IFIC) in our ATLASDATADISK space token
 - 1TBytes outgoing in our ATLASPRODDISK space token.
- During the test **disk went to full two time**.
 - More space was added on the fly with the deployment of two new disk servers (17TB + 34 TB) to reach 102 TBytes
 - Quick reaction and the flexibility to add more servers easily to Lustre file system lead us to absorb the data.

Space tokens were updated to reflect the used and available space but:

- ATLAS tools seems to ignore the values and continue transfers even when no space available was published.



- **Targets achieved:**
 - Data Transfers:
 - Reach >250MB/s in T1-T2 distribution
 - Excellent efficiencies (95% to 100%)
 - No major issues for IFAE, IFIC and UAM rather than disk space shortage
 - MC production
 - Nice behavior not only during STEP but since more than 1 year (efficiencies ~95%)
 - User Analysis
 - First bulk experience together with many other demanding activities
 - First days fair-share inconsistency for the scheduler problem
 - Good efficiencies overall:
 - 85% for WMS jobs and 96% efficiency for the PanDA based analysis pilot jobs
 - Minor (and not site related) issues
 - Reprocessing:
 - Passed successfully even though we'll continue to tune the tape by system by installing more tape drives and optimizing the recall pool schema.
 - ATLAS thinking about using T2s for repro with FronTier technology

- **Issues learned and to improve:**
 - Storage still the most critical services all across sites
 - Minor issues at our cloud:
 - *SE instabilities: gridftp doors overload, disk space shortages*
 - Optimization for User Analysis (UA):
 - Take advantage of site's architecture:
 - *Distinguish and enforce DQ2Local or FS to target best site performance*
 - Enhance for using native and optimized protocols
 - File:// for POSIX like SE (IFIC and LIPs)
 - Tune RA for dCache sites (currently 32kb is the default in UA system)
 - Need periodic checking of the activities
 - Good implication of site people during STEP. Helped a lot to be a good ATLAS clouds !!!

- **Status and plans:**
 - ES sites are ready, showed robustness, stability and performed good.
 - ATLAS computing system is ready:
 - DDM improved during the last two years
 - PanDA MC and UA system increased global efficiencies and running stable
 - *Now centralized at CERN*
 - Last updates should be thought well in advance to have the sites ready for November and avoid big interventions after the end of the month.