ATLAS Distributed Analysis tests in the Spanish Cloud

Presenter: Santiago González de la Hoz
IFIC (Instituto de Física Corpuscular), Centro mixto CSIC-Universitat de València

[1]GONZALEZ, S.,[1]SALT, J., [1]AMOROS, G., [1]FASSI, F., [1]FERNANDEZ, A., [1]KACI, M., [1]LAMAS, A., [3]MARCH, L., [1]OLIVER, E., [1]SANCHEZ, J., [1]VILLAPLANA, M., [1]VIVES, R., [2]BORREGO, C., [2]CAMPOS, M., [2]NADAL, J., [2]PACHECO, A., [3]DEL PESO, J., [3]PARDO, J., [3]MUñOZ, L.,[3]FERNANDEZ, P., [3]DEL CANO, L., [4]ESPINAL, X.

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

www.eu-egee.org

e-infrastructure

Outline

- **The ATLAS experiment**
- **Spanish Cloud**
  - Spanish Resources
- **Distributed Analysis Test**
  - What tests have been performed in the Spanish Cloud?
  - Lessons learned
- **Conclusions**

The ATLAS Experiment

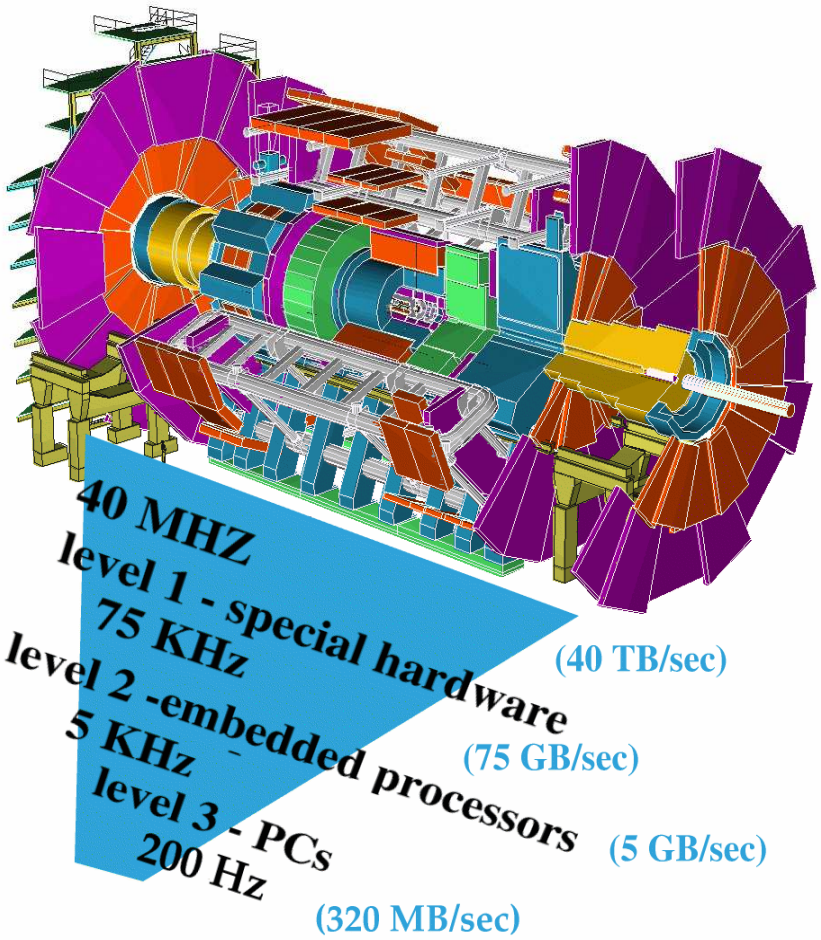

The offline computing:

- Output event rate: 200 Hz ~ 10^9 events/year
- Average event size (raw data): 1.6 MB/event

Processing:

- 40,000 of today's fastest PCs

Storage:

- Raw data recording rate 320 MB/sec
- Accumulating at 5-8 PB/year

A solution: Grid technologies

GRID computing

GRID is used to solve problems of data simulation, storage and analysis.

Data per year: ≈ Petabytes

- event generation
- simulation of what happens in the detector
- reconstruction of an event from what happened in the detector

Spanish-Iberian Cloud for ATLAS

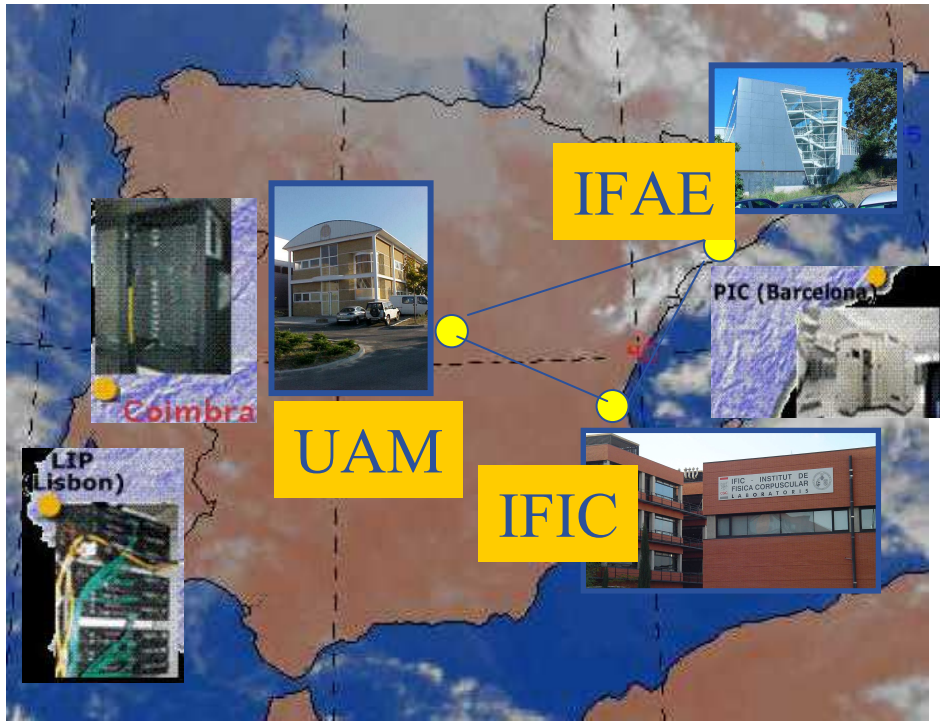

SWE Cloud:
Spain-Portugal
Tier1:
PIC-Barcelona
Tier2's:
UAM, IFAE & IFIC
LIP & Coimbra

- **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.

Spanish Resources

– It 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**.

Spanish Resources

- **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

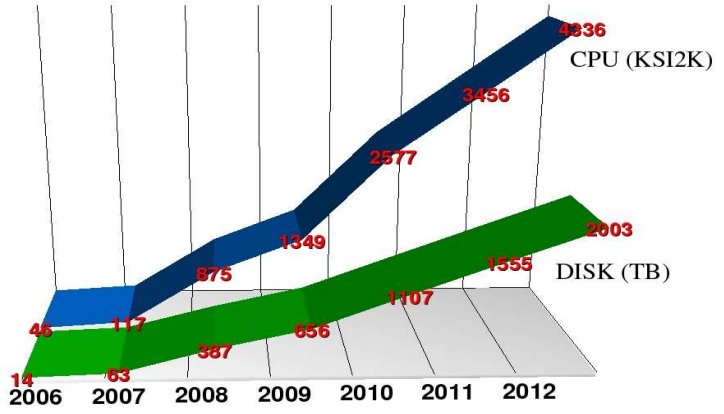

Strong increase of resources

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

	IFAE	UAM	IFIC	TOTAL
CPU (ksi2k)	201	276	438	915
Disk (TB)	104	147	198	449

New acquisitions in progress to get the pledged resources

Accounting values are normalized according to WLCG recommendations

Spanish Resources

- **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

Distributed Analysis Tests

- **Goals**
  - **Distributed Analysis Challenges** need to be performed in order to validate site and cloud readiness for the full-scale user load
    - The needs of analysis jobs differ from those of production, so though a site may function for Monte Carlo simulation, may perform poorly for analysis
  - We are trying to **identify breaking points and bottlenecks which result from the site/cloud design or configuration**

- **ATLAS testing framework**
  - An automated DA Challenge framework based on:
    - **Ganga** submits a bulk of jobs –user like– to the target sites, keep tracking of the status and report once finished.
    - Applications is a **real analysis code** from physicists, not “hello word”
    - (presently) the **Ganga EGEE/LCG and NDGF backend**, using its data-based brokering and splitting. **Bulk submit jobs to the WMS** (direct submission to the CE)
    - Both **Posix I/O and “Copy mode”** to read/access the input data
    - The jobs are running with “**user role**” (a real case)

Distributed Analysis Tests

- A single master "job" on many datasets is submitted
  - *Ganga splits this into one subjob per dataset*
  - *Ganga EGEE/LCG jobs are instrumented to collect statistics*
  - *Ganga uses by default Posix I/O to access/read the input files on the WN from the SE:*
    - dcap://.. is used for dCache
    - rfio://… is used for Castor or DPM
    - file://.. is used for Storm SEs (Lustre)

    Used in Spanish Cloud
  - *"coy mode" uses dq2-get from the local SE to the WN (has been tested after the posix I/O tests).*

– ATLAS distributed analysis team is **working with the Panda team** to make these test run on Panda
  - Need to understand how to submit/split similar to above
  - Need to instrument the Panda analysis jobs

– Metric of each individual tests are collected automatically (http://gangarobot.cern.ch/st/):
  - **Performance metrics**:
    - *Success/failure rate, CPU/Walltime, Events per second, etc..*
  - **Error classification**:
    - *Different I/O errors*

What tests have we already performed in the Spanish Cloud

- **In these challenges, we were coordinated with ATLAS Distributed Analysis team (Dan van der Ster and Johannes Elmsheuser)**

- **Presently we are interested in submitting large numbers of jobs intensive on data input**
  - Until a maximum of 200 jobs has been submitted to each site, thus up to ~1000 total. We ran different tests with different number of jobs.
  - The jobs read directly from the SE using posix I/O
  - Focusing on finding site limitations, e.g. heavy loads on storage and network resources, wms, client, etc.

- **Specifically, the tests runs an AOD muon analysis everywhere right now (from M. Biglietti):**
  - UserAnalysis pkg, Athena 14.2.20 (ATLAS software version)
  - mc08*AOD*e*s*r5* + some muon datasets replicated to the sites
    - **~40M of events, ~140K files**
  - Each job processes an entire dataset, and the system submits one job per dataset

What tests have we already performed in the Spanish Cloud

- **Example Plot: Overall Eficiency, CPU/Walltime (%)**

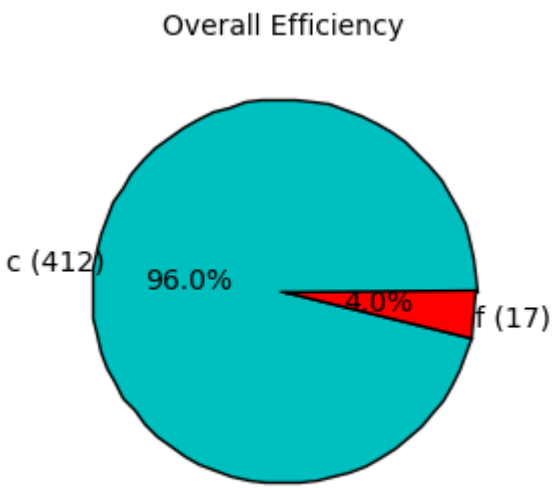

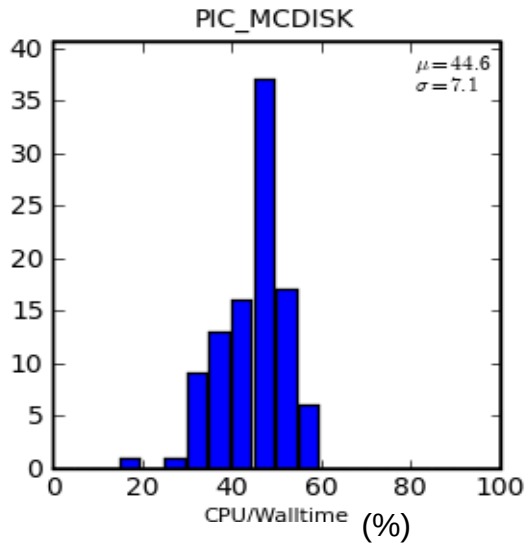

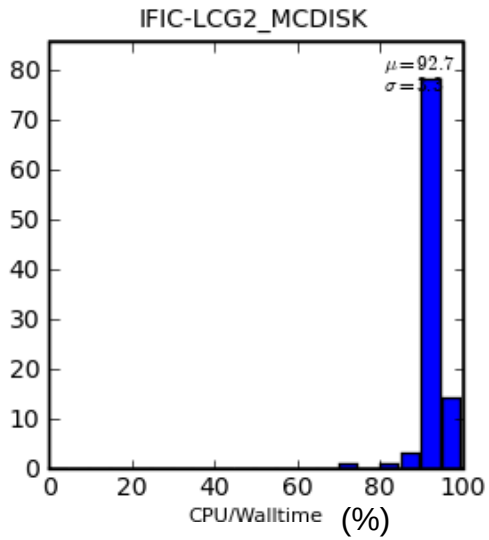

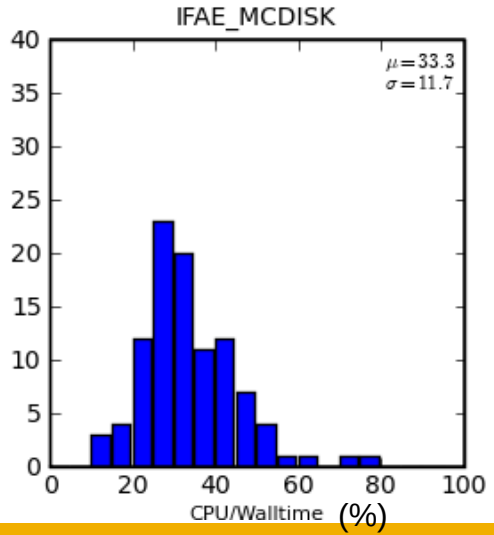

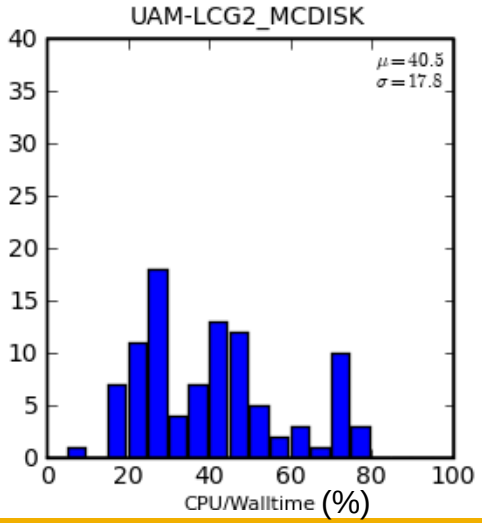

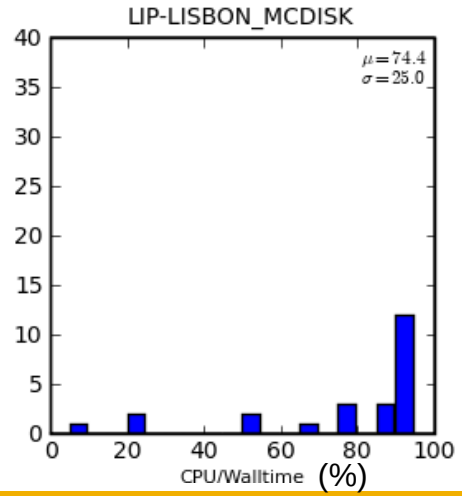

CPU/Walltime is the CPU Percent Utilization

What tests have we already performed in the Spanish Cloud

- **Example Plot: Event/second**

(This is the event rate for the athena execution only, i.e. the denominator is the time from athena start to finish)

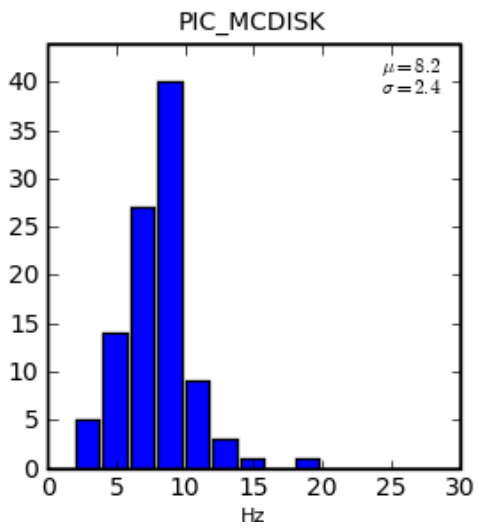

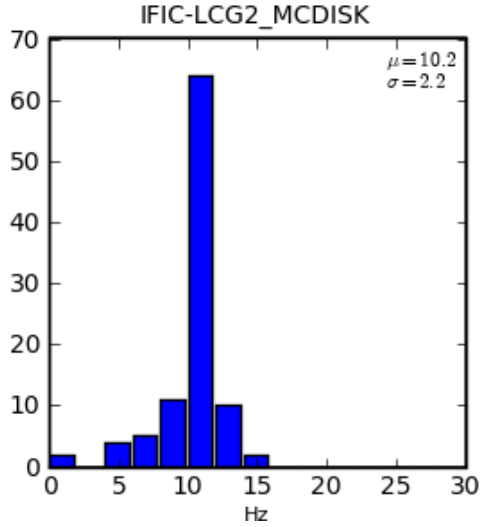

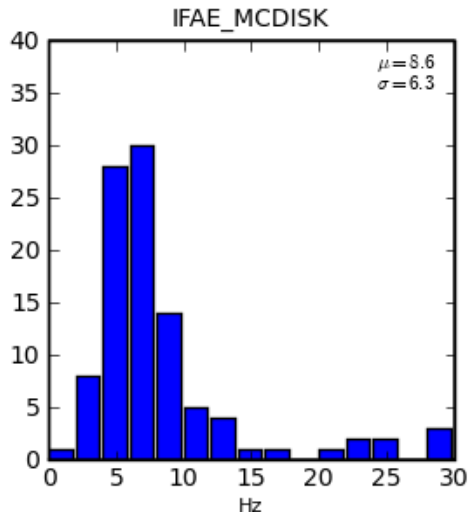

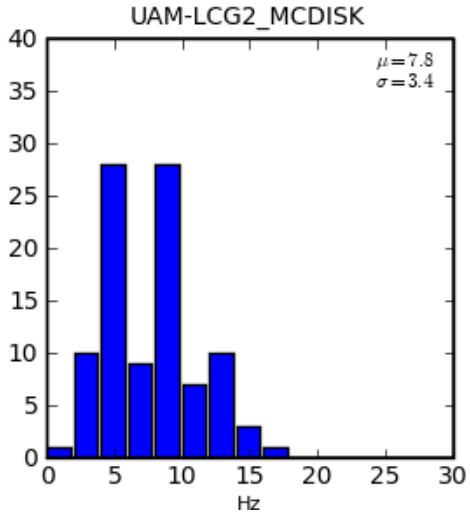

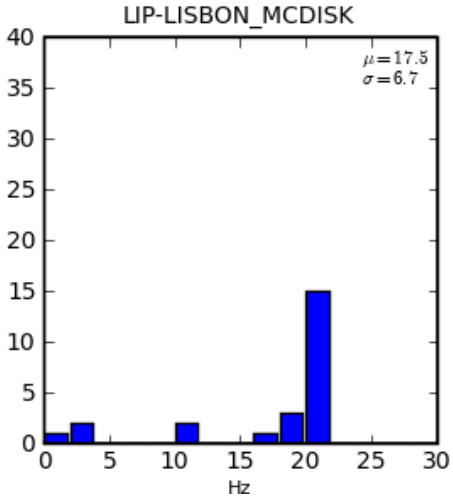

What tests have we already performed in the Spanish Cloud

- **Example Plot: All sites**

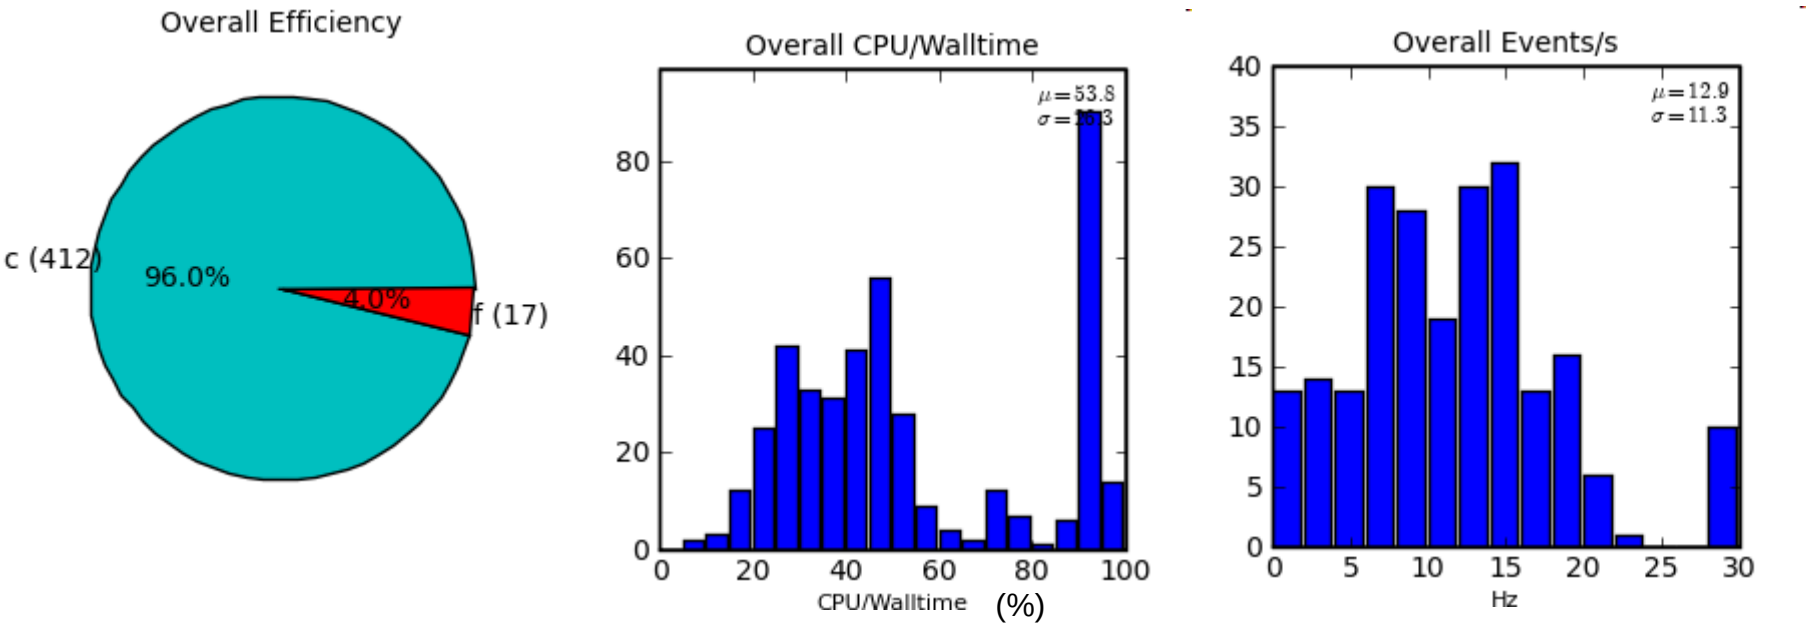

- **Under these load (500 jobs and ~40M of events), most jobs were fast (average 54% of CPU and 13 Hz)**
- **Specific goal for analysis is to work toward 85% of CPU and 15 Hz per job**
  - We are going to try new methods like FileStager

What tests have we already performed in the Spanish Cloud

- **Local posix I/O vs. FileStager:**
  - In 2008, we studied local posix I/O:
    - Why? This is how LCG/EGEE users access the data now.
      - *Why? Because of the ACM, and, in theory, posix I/O should minimize the I/O*
  - In practice, rfio/dcap/etc… is rarely tuned to Athena's access pattern.
    - Often, much more data is transferred than needed.
    - Generates high load on network, SE, and disk pools.
    - Tuning the readahead buffers is difficult
    - But there are exceptional sites.
- **Now we are looking at FileStager as an alternative:**
    - Pre-copies the next input file with lcg-cp (or anything else) in a background thread.
    - It seems to improve many, but not all, sites.
    - Probably not ideal for TAG (or other random access) analyses.

What tests have we already performed in the Spanish Cloud

- **Simultaneous tests of posix I/O and FileStager in ES cloud:**
  - 100 (posix)+50(FS) jobs per site.
- **Results of posix vs. F.S.:**
  - Success rates are similar.
  - % CPU Utilisation
    - F.S. improved
      - ***From 54% to 65%***
  - Events/s similarly increased.

Posix

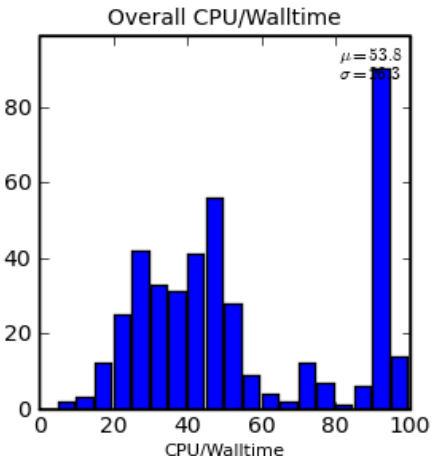

FS

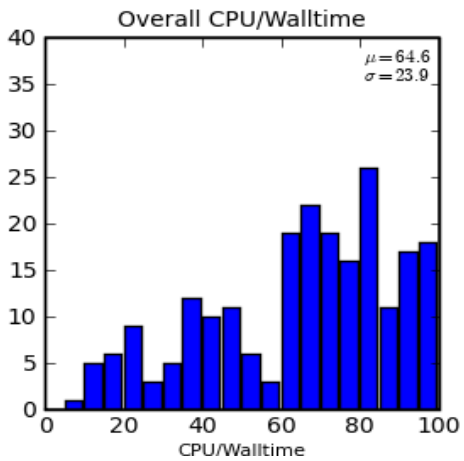

  - But, **looking at sites individually, we see a different story.**

What tests have we already performed in the Spanish Cloud

Left: Posix I/O, Right: FileStager

Using dCache

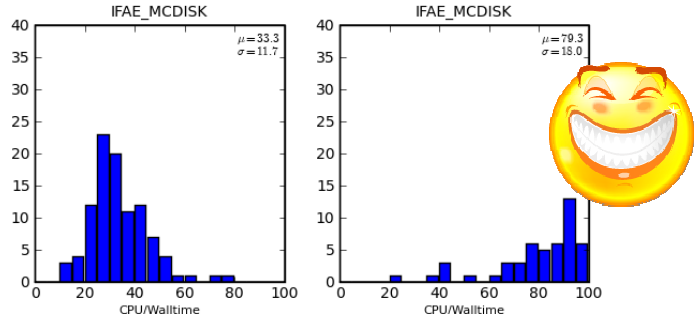

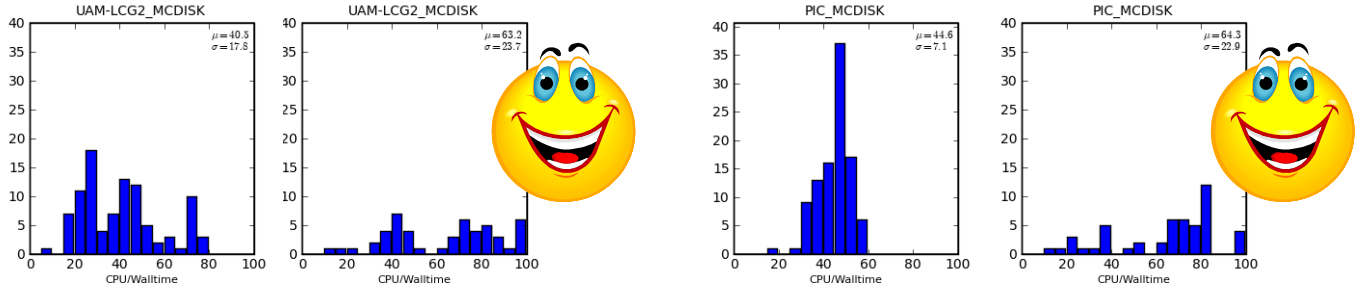

Using Lustre

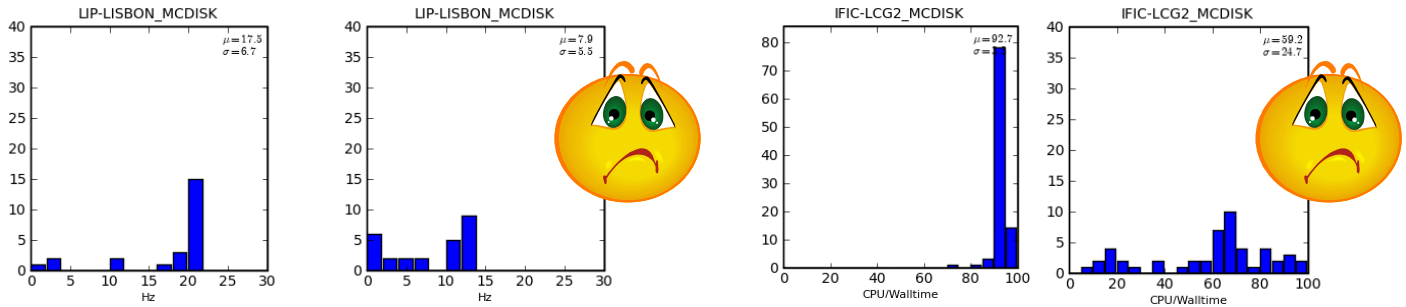

Lessons learned

- **The results and tests led to the discovery of:**
  - Target **dataset are not present** there (all Tier2s must have all AOD data). (LIP case)
  - Athena **releases wrong installed** or missing (UAM case).
  - **Inconsistency** between **mappings and permissions** between SRM and WNs, so file can't be accessed locally (IFIC case)
  - **Problem** with the **creation of directories** due to a STORM/GFAL (IFIC-LIP case)
  - **Wrong access to the files**: used rfio:/ instead of file:/ suitable for Lustre (LIP case)
  - **Wrong site identification** by Ganga (IFAE-PIC case)
  - **Bad information published** by the site in the information system (LIP Case)

- **The File stager shouldn't be used for sites like LIP and IFIC which are using Lustre as File System**
  - In this case it is faster read the files from the SE than copies them to the WN.
    - *Maybe better results with faster CPUs, local disks, etc..*

Conclusions

- **Distributed analysis tests are necessary to stress the facilities at a simulated full user load.**
- **Very useful first test that allows to identify “big” problems:**
  - missing software, missing/bad information published and already to point out some limitation (e.g.: network connection in some sites).
- **Test new parameters (FileStager) to improve access to the data using dCache and Lustre**
- **Breaking points and bottlenecks which result from the site/cloud design or configuration has been identified.**