

Evolution of Atlas Computing and Data model and how this affects the IFIC Tier-2

Santiago González de la Hoz, E. Oliver, M.
Villaplana, A. Fernández, F. Fassi, M. Kaci, V.
Lacort, J. Salt, J. Sánchez, V. Sánchez

Instituto de Física Corpuscular (IFIC)
Universitat de València - CSIC



Université Mohammed V Agdal

جامعة محمد الخامس اكدال

Faculté des Sciences Rabat


كلية العلوم الرباط



Outline


- Introduction to **Atlas** and **IFIC Spanish Tier2**
- **Evolution** of the **Atlas data and computing model**
 - *Flattening the model to a Mesh*
 - *Availability and Connectivity*
 - *Job Distribution*
 - *Dataset replication*
 - *IFIC infrastructure and operations*
- Example of **Grid** and **Physics Analysis**

ATLAS CENTERS







<http://dashb-earth.cern.ch/dashboard/dashb-earth-atlas.kmz>

ATLAS Production on the Grid



ATLAS Data Distribution and Data Handling




SPANISH CONTRIBUTION TO ATLAS 2011




Computing (3900 cores & 4 PB)

- ES Tier1 at PIC Barcelona for Atlas:
1,600 cores, 1.7 PB disk and 1.6 PB tape
 - Federated ES Atlas-Tier2 (IFIC/UAM/IFAE):
1,960 cores and 2 PB disk
 - Local Tier3 farms at IFIC, UAM and IFAE:
300 cores and 250 TB disk
- fundamental for physics analysis




Maximum: 5,915, Minimum: 0.00, Average: 2,780, Current: 2,702

NEEvents Processed MEvents (Million Events) (Cumulative Graph)
7053 Hours from Week 00 of 2011 to Week 42 of 2011 UTC




Total: 37,432, Average Rate: 0.00/s

Spanish ATLAS Tier2



IFIC : is the coordination Center of the Spanish ATLAS Tier-2

- The Spanish ATLAS Tier-2 is integrated in the WLCG Project (Worldwide LHC Computing GRID) and follows the ATLAS Computing Model



Deployed Equipment:

Site Name	HEP-SPEC06	Disk (TB)
IFIC	6950	940
UAM	3314	302
IFAE	4993	470
Total	15257	1712

as of February 2012

¡ IFIC has reached 1 PB disk Storage !

Human Resources:

14 FTE

¡The Spanish ATLAS Tier-2 has Got 2 PB Disk Storage!

IFIC Computing Infrastructure Resources



as of April 2012



ATLAS Tier2:

- 48 x Xeon Quadcore E5472 3.00 GHz 16 GB RAM
- 32 x Xeon Quadcore L5520 2.77 GHz 24 GB RAM
=> 640 cores
- 20 x Sun X4500 y X4540
- 7 x Supermicro
=> 1 PB

← ATLAS

- SAS disks
- 2 TB/disk
- 10 Gbe

GRID-CSIC:

- 48 x Xeon Quadcore E5472 3.00 GHz 16 GB RAM (Infiniband)
- 106 x Xeon Quadcore E5472 3.00 GHz 16 GB RAM
=> 1232 cores
- 5 x Sun X4540
=> 180 TB

CSIC resources (not experiment resources):

- 25% IFIC users
- 25% CSIC
- 25% European Grid
- 25% Iberian Grid

← EGI

To migrate Scientific application to the GRID

IFIC Computing Infrastructure Resources (May 2012)



#	CPU	Cores	Mem/core	HEP-SPEC2006/box	HEP06 TOT
32	Xeron E5472 3.00GHz	8	2 GB	73.62	2355.8
32	Xeon L5527 2.27 GHz	8	3 GB	147.20	4710.4
					7066.2
Pledge 2011					6650



With the GRID-CSIC project we are adding
complementary resources up to:
12 KHEP-SPEC06

IFIC Storage resources and Network




SUN X4500: 5x(500GB) + 1x(1TB) 112 TB
SUN X4540: 14x(1TB) 476 TB
SuperMicro: 11x(2TB) 633 TB

TOT: 1221 TB

Pledge 2012 1175 TB

- SUN thumpers:
 - X4500 + X4540
 - Lustre v1.8
- New Supermicro servers:
 - SAS disks, capacity 2 TB per disk and 10 Gbe connectivity
- Srmv2 (STORM)
- 3 x gridftp servers
- Data Network based on gigabyte ethernet.
- 10GBe network



Storage tokens

JUNE 2012




TOKEN	ASSIGNED	USED	FREE	%
CALIB	15.0	9.8	5.2	65.0
DATA	1,008.4	690.8	317.6	68.0
PROD	20.7	3.0	17.6	14.0
GROUP	152.5	51.2	101.3	33.0
SCRATCH	51.2	25.0	26.2	48.0
HOT	7.0	1.0	6.0	14.0
*LOCAL	76.8	70.4	6.4	91.0
sum	1,254.8	780.8	474.0	62.0

LOCAL is for T3 local users, out of T2 Pledges. Numbers reflect actual usable space for Filesystem.

IFIC Summary of year 2011




- 3,579,606 Jobs
- 6,038,754 CPU consumption hours
- 13,776,655 KSi2K (CPU time normalised)
- Supporting 22 Virtual Organizations (VOs)



(C) CESSA 'EGI View': IFIC-LCG2 / normcpu / 2011:1-2011:12 / SITE-VO / all (x) / ACCBAR-LIN / i

today
2011-12-19 23:30 UTC




Prevision 2013

For Spanish Tier2

Año	2010	2011	2012	2013
CPU(HS06)	12000	13900	13300	14450
Disk (TB)	1045	1880	2350	2650

- For IFIC (50%):
 - We already fulfil the CPU requirements:
 - 2 WN (275 HS06)
 - Increasing Disk:
 - 150 TB => 3 SuperMicro x 57.6 (2 TB disks)

ATLAS Computing Model: T0 Data Flow




6

ATLAS Computing & Distributed Model

- Original model was a working starting point but Tier-2 activity was strongly linked to the associated Tier-1 reliability.
 - Some Tier-1s did not have associated Tier-2s:
 - Few tasks to process even with big storage to host a significant fraction of tasks
 - Tier-2 had computing resources to do reprocessing but was limited due to required direct access to the Tier-1 database.
 - If the LFC at Tier-1 and the Tier-1 was in downtime there would not be activity in its associated Tier-2s.
 - If Tier-1 Storage service was in downtime, the production and data distribution to/from the Tier-2s was stopped.

ATLAS Computing & Distributed Model

As a consequence, some of the sites were not used at full capacity, especially Tier-2s!!!!!!






- Operational improvements were taken by ATLAS
 - based on **existing network provides good connectivity** to many Tier-1s and/or Tier-2s.
 - To **make direct transfers from Tier-2s to Tier-1s of different clouds.**
 - Even transfers with other Tie-2s globally
 - Evolutions came with **CVMFS & Frontier/Squid**
 - CVMFS is a network file system based on HTTP
 - **ATLAS software releases** and the smaller file-based database are now installed on the server at CERN.
 - there is no more need to install them at the site.
 - This has **removed the workload in software installation and the bottlenecks with the shared file systems.**
 - Frontier/Squid is an http-based system to access database
 - with caching, avoiding a high load on the database and latency in accessing the database from remote sites.
 - Introduction of the system has removed limits with the database access, allowing the jobs running at Tier-2 sites accessing the database at Tier-1 sites.

From Tiered to a Mesh model




Data Processing Model Revised



東京大学 ATLAS Operations - Ueda I. (ISGC2012, Taipei, 2012.03.02.)

Auto-routing

- mesh transfer tests from every site to every site ('sonar' tests)
- the transfer system decides the site-to-site routing based on the measurements



Connectivity

- **Network is a key component** in the evolution of the Atlas model
- **Transfers among sites measured by Atlas through FTS log :**
 - Matrix of NxN components
 - < 1 MB or > 1 GB files to split transfer and initialization components
 - the limit is set to an overall transfer performance of at least 5MB/s to at least 80% of Tier1s in ATLAS, for very large files (>1GB).
- **IFIC is qualified as T2 well connected** to T1s and was included as **multicloud site** -> **Direct transfers from/to all ATLAS T1s and process data from many T1s**

Transfer
monitoring


Proc	Source	Cloud	Destination	Cloud	SMALL FILES			MEDIUM FILES			LARGE FILES		
					MB/s	GB/s	#Ev	MB/s	GB/s	#Ev	MB/s	GB/s	#Ev
8	IFIC	ES-T1	IFIC-LOG2	ES-T2	0.55+0.71	6.26+11.78	3746	8.07+4.26	751.56+119.13	1348	8.74+4.56	1.95+0.78	6856
8	IFIC-LOG2	ES-T2	IFIC-LOG2	ES-T2	1.09+0.47	18.85+4.61	16	6.3+1.01	196.89+10.32	11	7.09+2.58	2.33+0.5	209
7	IFIC-LOG2	DE-T1	IFIC-LOG2	ES-T2	1.22+0.77	45.56+6.62	773	4.85+3.31	757.12+287.67	1409	5.78+3.98	1.9+0.82	1999
7	IFIC-LOG2	IT-T1	IFIC-LOG2	ES-T2	1.24+1.29	35.5+35.83	152	10.99+3.6	759.4+28.41	1669	12.03+5.29	1.85+0.75	5081
7	IFIC-LOG2	NO-T1	IFIC-LOG2	ES-T2	1.47+0.33	32.76+4.43	211	8.89+5.14	580.65+179.94	127	13.14+6.21	1.78+0.69	1851
7	IFIC-LOG2	UK-T1	IFIC-LOG2	ES-T2	1.56+0.88	44.11+26.72	262	10.95+5.07	874.59+184.97	1531	11.48+6.35	1.59+0.66	3058
7	IFIC-LOG2	NL-T1	IFIC-LOG2	ES-T2	1.23+0.8	56.77+21.02	1014	4.1+4.38	533.41+357.28	702	8.85+6.32	1.75+0.54	1113
7	IFIC-LOG2	TAIWAN-LOG2	TAIWAN-LOG2	ES-T2	-	-	-	1.57+1.13	434.39+290.73	791	3.16+1.44	1.65+0.46	1558
7	IFIC-LOG2	CA-T1	IFIC-LOG2	ES-T2	0.6+0.44	31.57+30.9	15	3.09+2.42	777.23+169.43	380	3.11+2.8	1.47+0.45	1927
7	IFIC-LOG2	FR-T1	IFIC-LOG2	ES-T2	0.36+0.6	18.56+31.8	204	8.16+3.87	571.89+240.96	2244	9.74+0.3	1.57+0.45	3833
7	IFIC-LOG2	NL-T1	IFIC-LOG2	ES-T2	0.81+0.06	19.75+0.83	11	8.1+1.37	785.86+118.39	966	8.42+2.5	1.79+0.76	997
7	IFIC-LOG2	US-T1	IFIC-LOG2	ES-T2	1.08+1.0	44.82+35.16	391	8.51+3.13	626.5+176.77	6226	10.24+4.45	1.74+0.63	5912
7	IFIC-LOG2	CERN-PROD	CERN-PROD	ES-T2	1.47+1.71	28.78+34.73	280	7.12+2.76	593.03+361.17	189	8.23+3.92	1.5+0.6	1079
7	IFIC-LOG2	ES-T2	IFIC-LOG2	DE-T1	0.92+0.4	20.0+0.0	10	2.41+1.35	200.0+0.0	10	2.9+1.46	2.7+0.66	69
7	IFIC-LOG2	ES-T2	IFIC-LOG2	IT-T1	1.48+0.94	36.5+28.35	10	6.12+1.47	778.24+27.61	81	8.95+5.71	2.11+0.16	8
7	IFIC-LOG2	ES-T2	IFIC-LOG2	NO-T1	0.74+0.27	20.0+0.0	15	5.55+3.35	506.73+149.78	412	8.08+4.41	1.64+0.35	23
7	IFIC-LOG2	ES-T2	IFIC-LOG2	UK-T1	0.52+0.24	24.04+12.45	7	4.68+2.1	307.74+87.16	118	5.67+3.84	1.93+0.66	90
7	IFIC-LOG2	ES-T2	IFIC-LOG2	NL-T1	1.57+0.9	34.28+12.55	20	5.69+2.81	577.67+344.03	148	6.44+4.47	2.71+0.93	200
7	IFIC-LOG2	ES-T2	IFIC-LOG2	TAIWAN-LOG2	0.4+0.04	20.0+0.0	10	2.9+0.37	200.0+0.0	10	2.71+5.65	1.47+0.55	51
7	IFIC-LOG2	ES-T2	IFIC-LOG2	CA-T1	0.36+0.27	25.27+12.33	12	3.97+3.11	354.62+293.9	23	12.11+5.36	2.25+1.21	29
7	IFIC-LOG2	ES-T2	IFIC-LOG2	FR-T1	3.41+1.91	56.43+29.98	13	9.67+3.09	685.8+226.67	36	7.42+3.75	2.69+1.11	136
7	IFIC-LOG2	ES-T2	IFIC-LOG2	NL-T1	0.71+0.04	20.0+0.0	10	8.43+2.71	874.76+181.54	165	9.81+4.63	1.13+0.26	135
7	IFIC-LOG2	ES-T2	IFIC-LOG2	US-T1	1.36+0.37	70.61+21.03	142	8.35+3.64	337.48+322.77	22	13.82+5.28	2.5+0.82	27
7	IFIC-LOG2	ES-T2	IFIC-LOG2	CERN-PROD	0.76+0.69	20.0+0.0	10	5.21+3.36	200.0+0.0	10	4.44+3.91	2.0+0.0	10
6	IFIC	ES-T2	IFIC-LOG2	ES-T2	1.2+0.11	20.0+0.0	10	3.63+4.43	346.2+245.59	187	11.94+4.6	1.39+0.41	98
6	IFIC-LOG2	ES-T2	IFIC-LOG2	ES-T2	0.84+0.11	20.0+0.0	10	6.54+1.08	183.2+18.65	32	14.27+10.55	1.64+0.46	14
6	IFIC-LOG2	ES-T2	IFIC-LOG2	ES-T2	0.81+0.11	20.0+0.0	10	6.44+0.89	200.0+0.0	10	12.78+4.71	2.5+0.6	13
6	IFIC-LOG2	ES-T2	IFIC-LOG2	ES-T2	0.96+0.08	20.0+0.0	10	8.03+0.51	200.0+0.0	10	23.08+8.91	3.43+1.16	29
6	IFIC-LOG2	ES-T2	IFIC-LOG2	ES-T2	1.14+0.17	20.0+0.0	10	5.13+5.42	183.2+15.75	33	12.85+12.76	2.75+0.82	23
6	IFIC-LOG2	ES-T2	IFIC-LOG2	ES-T2	1.0+0.04	20.0+0.0	10	6.07+1.38	200.0+0.0	10	15.23+5.16	2.0+0.0	10
6	IFIC-LOG2	ES-T2	IFIC-LOG2	ES-T2	0.91+0.09	20.0+0.0	10	5.93+0.93	200.0+0.0	10	14.89+7.84	2.0+0.0	10
6	IFIC-LOG2	ES-T2	IFIC-LOG2	ES-T2	0.97+0.05	20.0+0.0	10	6.25+0.82	200.0+0.0	5	5.38+7.46	2.0+0.0	5
2	IFIC-LOG2	FR-T2	IFIC-LOG2	ES-T2	0.61+0.2	20.0+0.0	10	4.12+0.51	200.0+0.0	5	-	-	-

Affects IFIC
transfers
changed to
several sources

Connectivity




T2D performance from Tier1s to IFIC



Connectivity


T2D performance from IFIC to Tier-1s

FTS transfer rates



- IFIC-LCG2_DATADISK - CERN-PROD_DATADISK (100 files)
- IFIC-LCG2_DATADISK - BNL-OSG2_DATADISK (1120 files)
- IFIC-LCG2_DATADISK - TRIUMF-LCG2_DATADISK (337 files)
- IFIC-LCG2_DATADISK - TAIWAN-LCG2_DATADISK (844 files)
- IFIC-LCG2_DATADISK - SARA-MATRIX_DATADISK (634 files)
- IFIC-LCG2_DATADISK - NIKHEF-ELPROD_DATADISK (331 files)
- IFIC-LCG2_DATADISK - FZK-LCG2_DATADISK (261 files)
- IFIC-LCG2_DATADISK - RAL-LCG2_DATADISK (2095 files)
- IFIC-LCG2_DATADISK - IN2P3-CC_DATADISK (1030 files)
- IFIC-LCG2_DATADISK - PIC_DATADISK (2450 files)
- IFIC-LCG2_DATADISK - INFN-T1_DATADISK (381 files)
- IFIC-LCG2_DATADISK - NDGF-T1_DATADISK (718 files)

Connectivity: Multi-cloud site & Production



Availability




Hammercloud: for ATLAS

- Sites are tested with 'typical' analysis jobs
- Problematic sites are identified online and automatically blacklisted
- Protect user jobs from problematic sites
- Last month **IFIC had > 90% availability** according to HC (ALPHA SITE)

EGI has own different availability tools:

- Based on ops VO
- Can have different results



IFIC LAST MONTH


Historic view for "panda_queues_all"
from 00:00 01.01.2012 to 00:00 29.02.2012

Show 100 entries

PANDA queue	SITE Name	TIER	CLOUD	History plot time bin = 118 hours	offline %	offline count	brokeroff %	brokeroff count	online %	online count	%
ANALY_IFIC	IFIC-LCG2	T2D	ES		2.58	3	4.72	6	91.5	12	0


T1 LFC catalog
Movement to T0

History Legend



<http://dashb-atlas-ssb.cern.ch/dashboard/request.py/siteviewhistorywithstatistics?columnid=564&view=shifter%20>

ES Availability 2012



Historic view for "panda_queues_all"
from 00:00 01.01.2012 to 00:00 31.01.2012

Show 100 entries


PANDA queue	SITE Name	TIER	CLOUD	History plot time bin = 60 hours	offline % count	brokeroff % count	online % count	NoQueue % count	test % count
ANALY_IFAE	ifae	T2D	ES		1.32 2	0.44 1	98.24 4	0 0	0 0
ANALY_IFIC	IFIC-LCG2	T2D	ES		1.32 2	8.94 5	89.74 7	0 0	0 0
ANALY_LIP-Coimbra	LIP-Coimbra	T2	ES		81.25 1	0 0	18.75 1	0 0	0 0
ANALY_LIP-Lisbon	LIP-Lisbon	T2D	ES		1.32 2	1.71 2	96.97 5	0 0	0 0
ANALY_NCG-INGRID-PT	NCG-INGRID-PT	T2D	ES		1.32 2	5.07 4	93.61 7	0 0	0 0
ANALY_UAM	UAM-LCG2	T2D	ES		1.32 2	54.27 2	44.41 4	0 0	0 0
ifae-ifaece02-at2ifae-pbs	ifae	T2D	ES		1.32 2	0 0	79.46 3	0 0	0 0
ifae-ifaece03-at2ifae-pbs	ifae	T2D	ES		1.32 2	0 0	79.46 3	0 0	0 0
IFIC-LCG2-ce03-atlasL-pbs	IFIC-LCG2	T2D	ES		6.04 3	0 0	93.17 3	0 0	0.79 1
IFIC-LCG2-ce05-atlasL-pbs	IFIC-LCG2	T2D	ES		6.04 3	0 0	93.17 3	0 0	0.79 1
UAM-LCG2-grid001-atlas-pbs	UAM-LCG2	T2D	ES		1.71 3	0 0	78.09 3	0 0	1.78 2

Showing 1 to 11 of 11 entries



First Previous 1 Next Last

Job Distribution at ES Tier-2

- Tier 2 usage in data processing
 - More job slots used at Tier2s than at Tier1
 - Large part of MC production and analysis done at Tier2s
- More Analysis jobs and “group production” at Tier-2s
 - less weight at Tier-1s
- Production shares to limit “group production” jobs at Tier-1, and run at Tier-2s
- Analysis share reduced at Tier-1s





More Analysis jobs at IFIC TIER 2



Number of jobs running at Tier0 plus all Tier1s (left) and at all Tier2s (right) from October 2010 to March 2012.

Job Distribution at ATLAS Tier-2s

Increase of analysis activities in Spring 2011 is due to the preparation for summer conferences.




Amount of analysis jobs splitted among different Tier types since January 2011

Therefore, a large part of the analysis and the MC production is done at Tier-2s.

All activities restricted to Tier-2s from October 2010 to May 2012.

Data Distribution: Replication of datasets

- New data distribution policy and the new **dynamic data placement algorithm** were deployed in Sep11.
 - Significantly changed the data volume transferred to Tier-2s.
- Using the new **Dynamic Data Placement** algorithm data distributed to:
 - Tier-1s for redundancy
 - **Tier-2s for analysis**
- Dynamic Data placement of data replicas at Tier-2s **based on usage as well as on demand.**





Replication of datasets in ES-Cloud





- ES-Cloud Tier 2 sites getting more datasets
- Bit more than Tier-1 PIC

IFIC is alpha T2D site (good connectivity and availability) -> candidate for more datasets replication



Data transfer volume in the ES-Cloud sites from January to March 2011 (left) and 2012 (right)

Replication of datasets in ATLAS



Tier-2s now get more datasets than Tier-1s because the disk size in Tier-2s has increased significantly while Tier-1 total size was not increasing as much

IFIC Lustre Filesystems and Pools

- With the [Lustre release 1.8.1](#), we have added [pool capabilities](#) to the installation.
- Allows us to partition the HW inside a given [Filesystem](#)
 - Better data management
 - Can separate heterogeneous disks in the future

4 Filesystems with various pools:

/lustre/ific.uv.es Read Only on WNs and UI.
RW on GridFTP + SRM


/lustre/ific.uv.es/sw. Software: ReadWrite on
WNs, UI (**ATLAS USES CVMFS**)

CVMFS FOR ATLAS

/lustre/ific.uv.es/grid/atlas/t3 Space for T3
users: ReadWrite on WNs and UI

xxx.ific.uv.es@tcp:/homefs on /rhome type
lustre. Shared Home for users and mpi
applications: ReadWrite on WNs and UI

- Different T2 / T3 ATLAS pools , and separated from other Vos
- Better Management and Performance




Example of Grid & Physics Analysis



Distributed Computing
and Data Management
Tools based on GRID
technologies have
been used by the IFIC
physicists to obtain
their results

Boosted top candidate




Anti-kT jets with $R = 0.4$ are indicated as red wedges in the R - ϕ view and red circles in the η - ϕ view. When the event is reclustered with $R=1$ (green) the three jets identified with the hadronic top quark decay merge into a single "fat" jet with $m_j = 197 \text{ GeV}$, $\sqrt{d_{T2}} = 110$, $\sqrt{d_{33}} = 40$. (see ATLAS-CONF-2011-073)

M. Villaplana

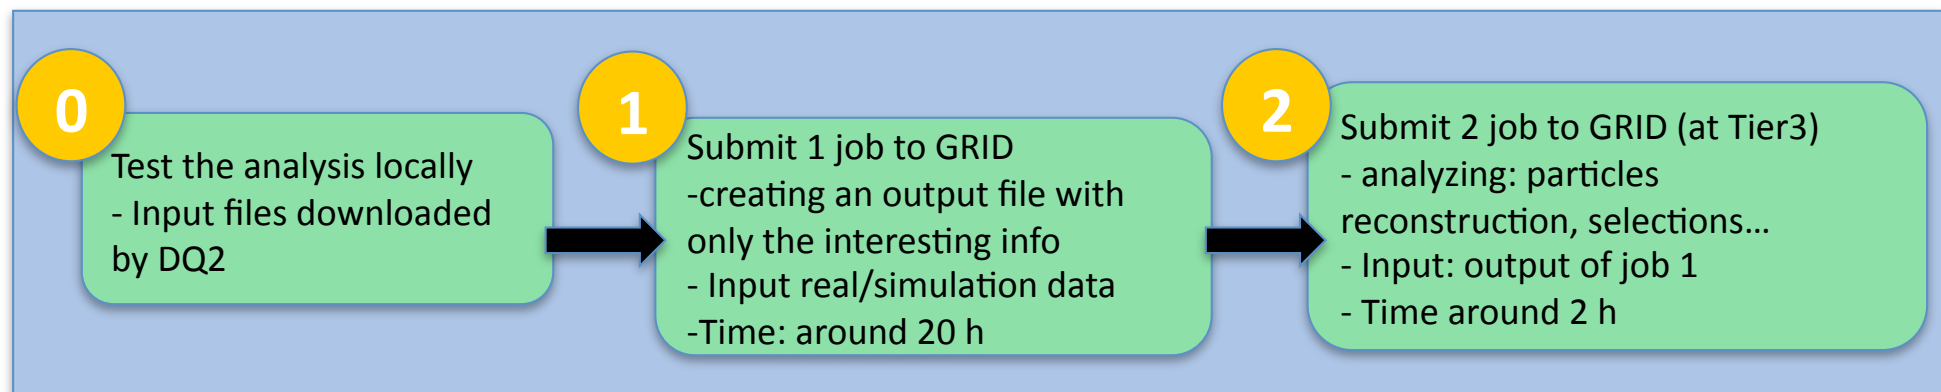
Daily user activity in Distributed Analysis

An example of Distributed Analysis

– Input files →




– Work flow:



Tier3 infrastructure at IFIC



- At IFIC the Tier3 resources are being split into two parts:
 - Resources **coupled to IFIC Tier2**
 - Grid environment
 - Use by IFIC-ATLAS users
 - Resources are idle, used by the ATLAS community
 - A **computer farm** to perform **interactive analysis** (proof)
 - outside the grid framework
- References:
 - ATL-SOFT-PROC-2011-018
 - ES Cloud CHEP contributions: ATL-SOFT-PROC-2012-017, ATL-SOFT-PROC-2012-010, ATL-SOFT-PROC-2012-006



Tools for Distributed Analysis

- For ATLAS users, **GRID tools** have been developed:
 - For **Data management**
 - **Don Quijote 2 (DQ2)**
 - Data info: name, files, sites, number,...
 - Download and register files on GRID,...
 - **Data Transfer Request (DaTri)**
 - Extension of Dq2 including quota management
 - Users make request a set of data (datasets) to create replicas in other sites (under restrictions)
 - **ATLAS Metadata Interface (AMI)**
 - Data info: events number, availability
 - For simulation: generation parameter, ...
 - For **Grid jobs**
 - **PanDa Client**
 - Tools from PanDa team for sending jobs in a easy way for user
 - **Ganga (Gaudi/Athena and Grid alliance)**
 - A job management tool for local, batch system and the grid




Daily user activity in Distributed Analysis

- 1) A **python script is created** where requirements are defined
 - Application address,
 - Input, Output
 - A replica request to IFIC
 - Splitting
- 2) Script executed with **Ganga/Panda**
 - Grid job is sent
- 3) Job finished successfully, **output files are copied** in the **IFIC Tier3**
 - Easy access for the user

Just in two weeks, 6 users for this analysis sent:

- 35728 jobs,
- 64 sites,
- 1032 jobs ran in T2-ES (2.89%),
- Input: 815 datasets
- Output: 1270 datasets



User experience

- Users are mostly satisfied with infrastructure and tools
- Functionality of dataset replication is transparent to users, but eventually they learn that it is in action:
 - “a job was finally sent to a destination that did not originally had the requested dataset”
 - “I saw later my original dataset has been copied”
 - “I just realized because I used dq2 previously to check where the dataset was”
- Replication is effective for official datasets, not for user generated ones:
 - “We see that is was not replicating our homemade data”

Summary

- Evolution of Atlas data model was presented, and how this affects the IFIC Tier-2.
 - Improvement of the Tier-2 resources usage
 - Changes in operations and infrastructure
 - Tier-2 activities are now less dependent to Tier-1 → Tier-2s participate to more critical activities in ATLAS.
 - Computing and data distribution model continues to evolve and improve beyond the original data processing model
- It was presented a real example on how users are working and their experiences:
 - Dynamic data replication useful and transparent