ATLAS: From STEP09 towards first beams

Graeme Stewart

2009-10-16

University of Glasgow | Department of Physics & Astronomy

GridPP
UK Computing for Particle Physics

Overview

- What we learned during STEP09
- Improvements we made
- Current priorities and preparations

Journées ‘Grilles France’, Octobre 2009

STEP09 Overview

- STEP09 was *Scale Testing for the Experiment Program 09*, i.e., Offline computing systems commissioning test
- It involved all major offline activities:
  - Monte Carlo Production
  - Full Chain Data Distribution
  - Reprocessing at Tier-1s
  - User Analysis Challenge: Hammercloud
  - ATLAS Central Services Infrastructure
- Done in conjunction with other LHC experiments

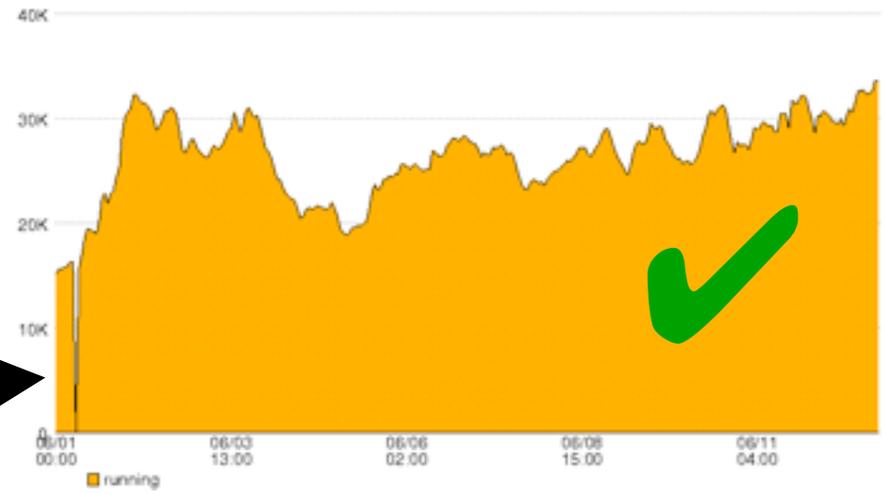

3GB/s

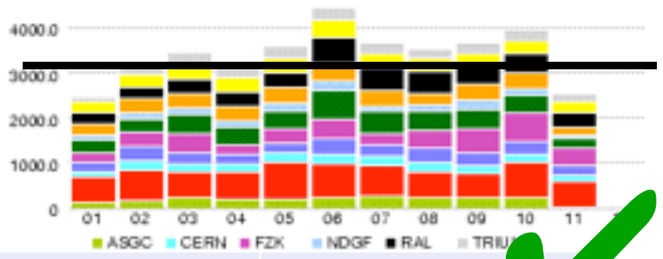

Cloud	Efficiency	Throug
ASGC	99%	397 MB/s
BNL	84%	1128 MB/s
CERN	100%	334 MB/s
CNAF	98%	561 MB/s
FZK	85%	556 MB/s
LYON	96%	620 MB/s
NDGF	84%	137 MB/s
PIC	93%	429 MB/s
RAL	99%	838 MB/s
SARA	53%	262 MB/s
TRIUMF	100%	297 MB/s

Peaks of 5.5GB/s

Though there are still many operational challenges with data distribution!

T1 Reprocessing Exercise

- Exercise designed to validate reprocessing from tape at T1s
  - Pre-stage files from tape via SRM
  - Run RAW to ESD reconstruction
  - Write ESD and log files back to tape
- Objective was to reprocess faster nominal data taking rate of 200Hz
  - Baseline metric at 400Hz (x2)
  - Enhanced metric at 1000Hz (x5) ☆
- Baseline metric translates to 20 000 RAW files reprocessed at a 10% T1

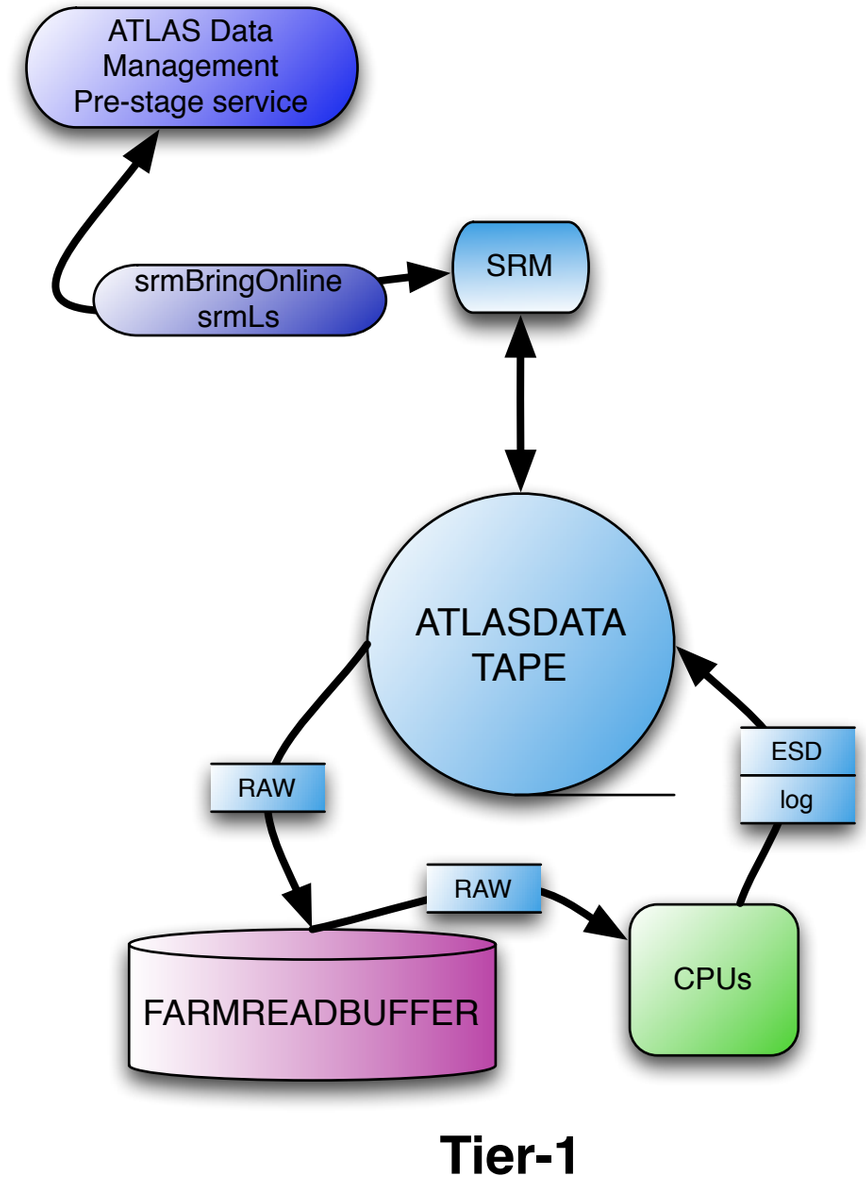

Journées 'Grilles France', Octobre 2009

Reprocessing Results

T1	*Base Target*	*Result*	*Comment*
ASGC	10 000	4 782	Many batch system and basic setup problems
BNL + SLAC	50 000	99 276	
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 ☆	

- Reprocessing from tape now validated in 6/10 T1s, with 3 more very close - improvements and retests planned

Journées 'Grilles France', Octobre 2009

Recent Results from IN2P3

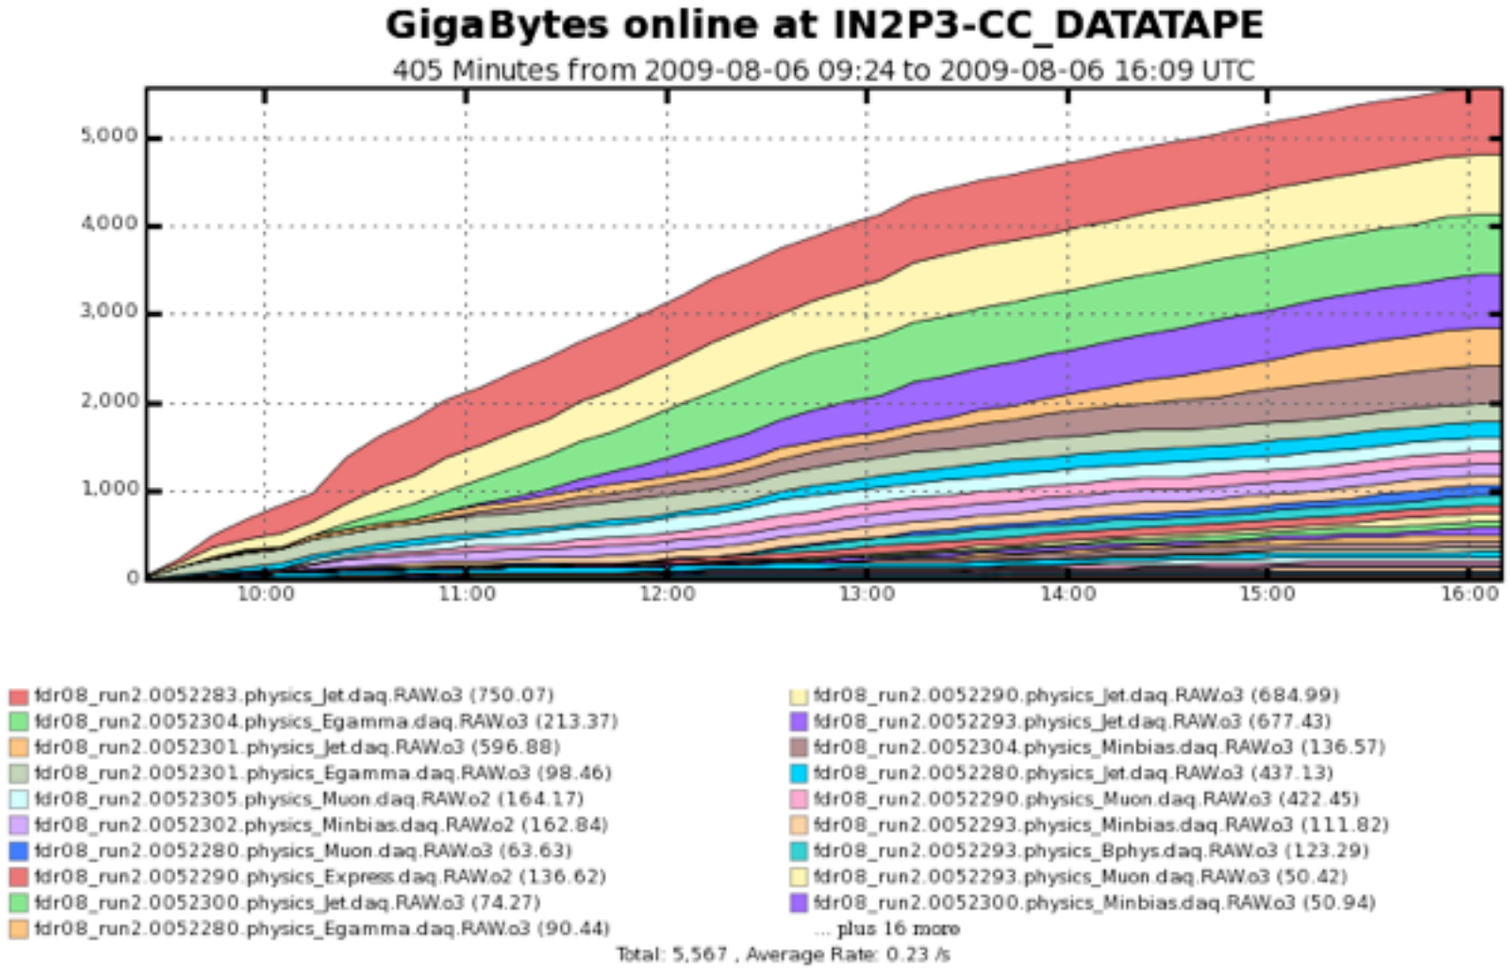

- 230MB/s staging speed

(Thanks to Catherine Biscarat)

Analysis: Global Results

- ~1M jobs submitted, 83.4% success rate
- 26.3B events processed, 28.6kHz across grid
  - N.B. This sounds impressive, but is actually only a few 10s of power users
- Average job's event rate: 7.7Hz
- Average job's cpu efficiency: 0.39

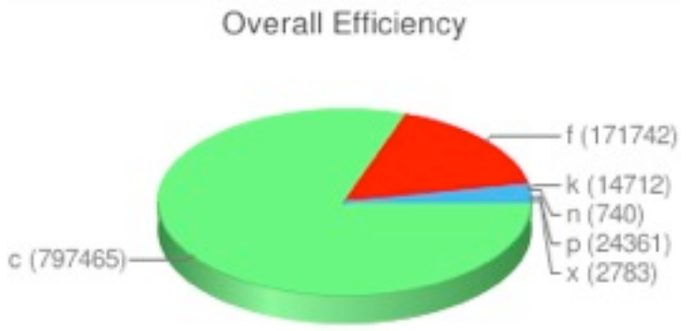

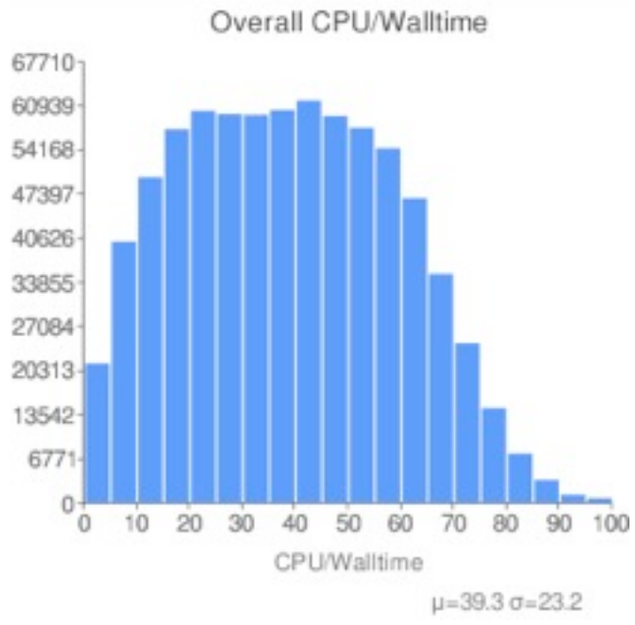

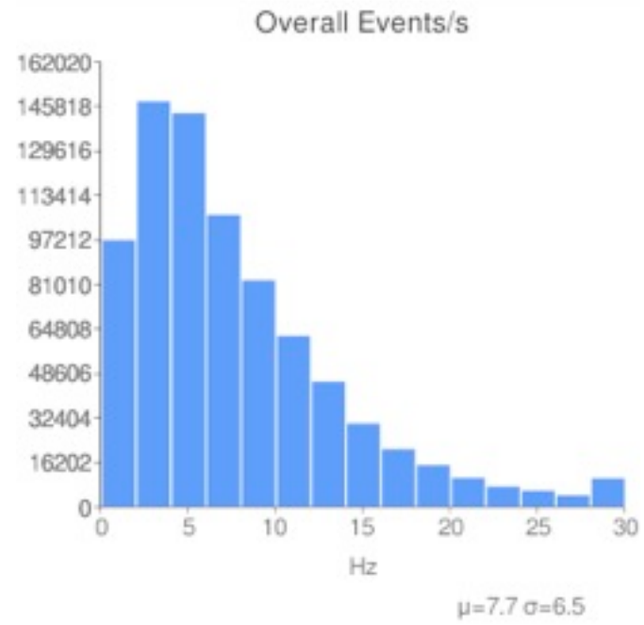

Journées 'Grilles France', Octobre 2009

Cumulative Events per Site

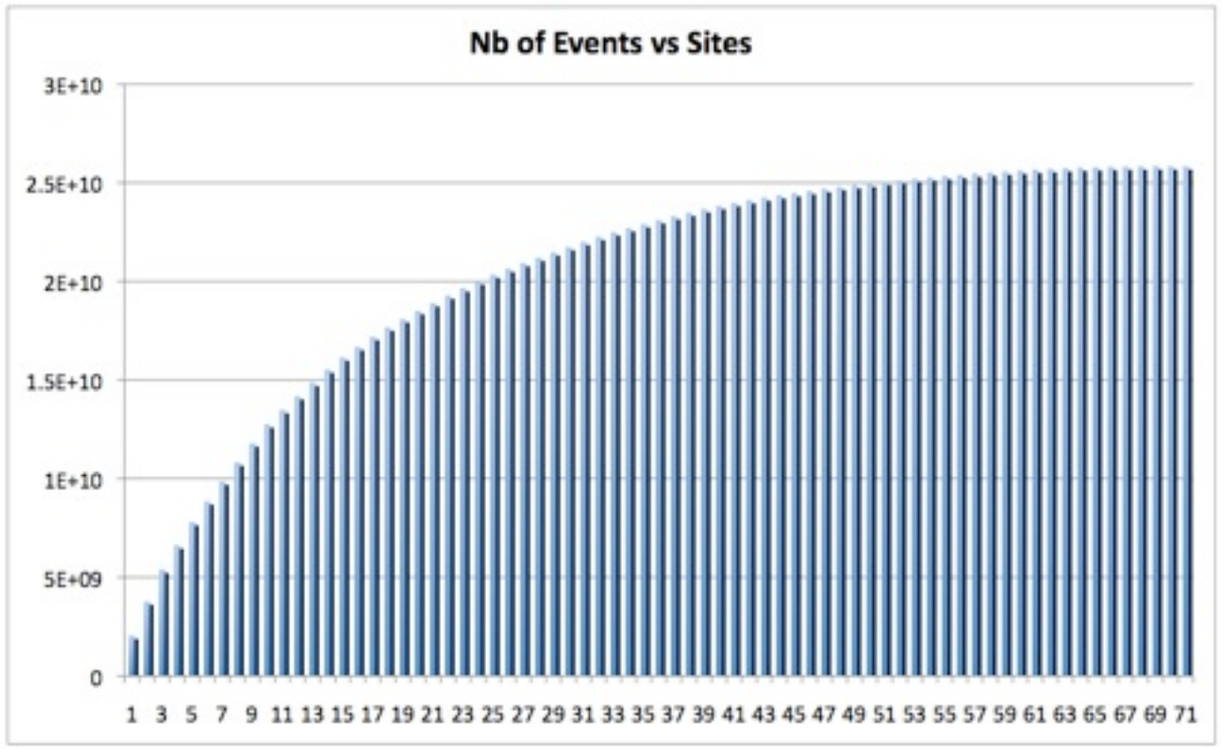

- 50% of all events are processed in 11 sites
- 90% of all events are processed in 37 sites

Clearly many sites had substantial improvements to be made

Hammercloud: Get Fit!

- UK has run 30 HammerCloud tests since STEP09
  - "Sites are tweaking RFIO buffers, balancing disk servers, LAN architecture, batch system limits, WN disk, etc."
- DE has run even more
  - "June and July tests on specific sites show great improvement eg in Wuppertal: Run jobs at ≈95% efficiency."
- NIKHEF redesigned their network to cope with the very high i/o required by analysis jobs

What works...

- Many disk servers - spread load
- No tight bottlenecks - good network infrastructure
- You need good, but not stunning, hardware
- Target: 5MB/s per job (320Mb/s per 8 cores) so 1Gb is ok
  - 10Gb is probably worthwhile on the disk servers
  - Channel bonding also can work
- dCache, DPM and Storm/Lustre can all give good results

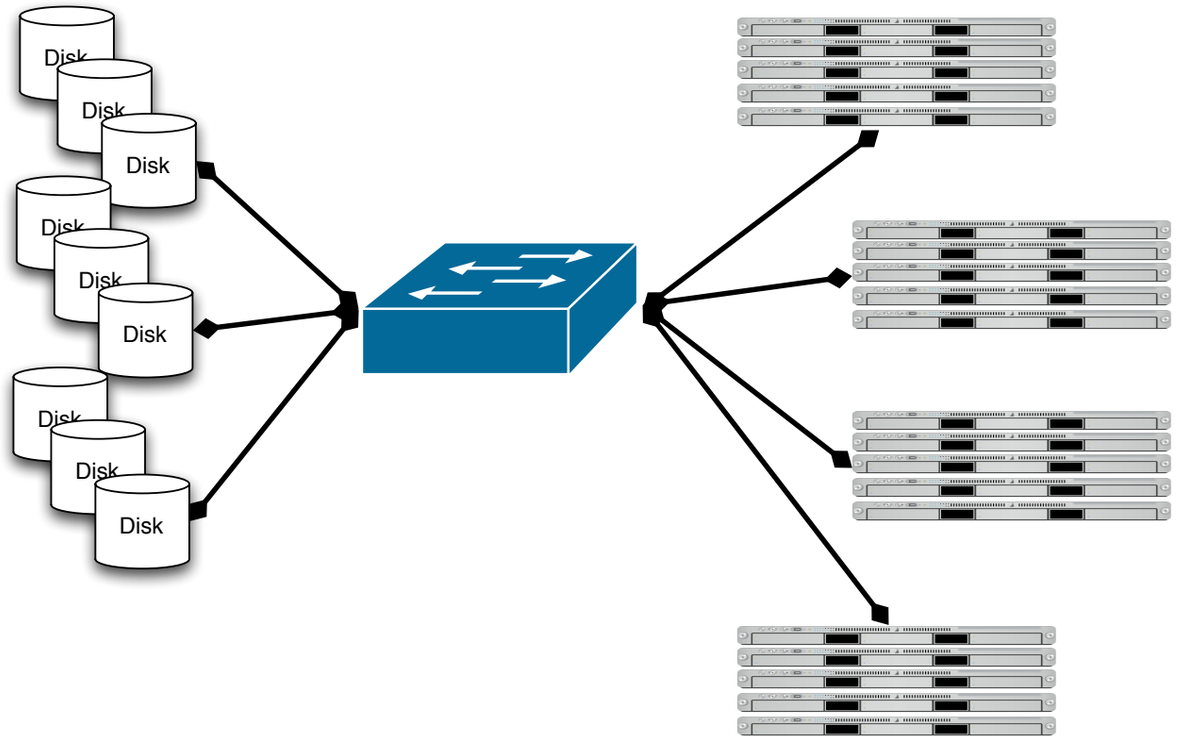

Improvements

	STEP09	Post STEP
DE	499M	800M
UK	400M	1000M
NL	135M	343M

- Events in each cloud in 24 hours

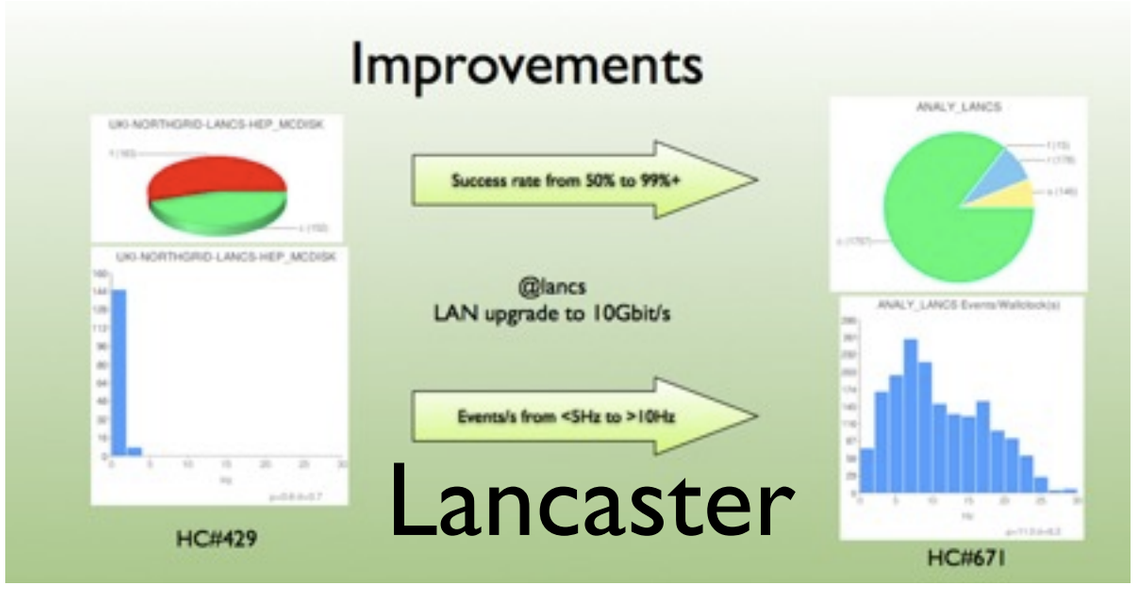

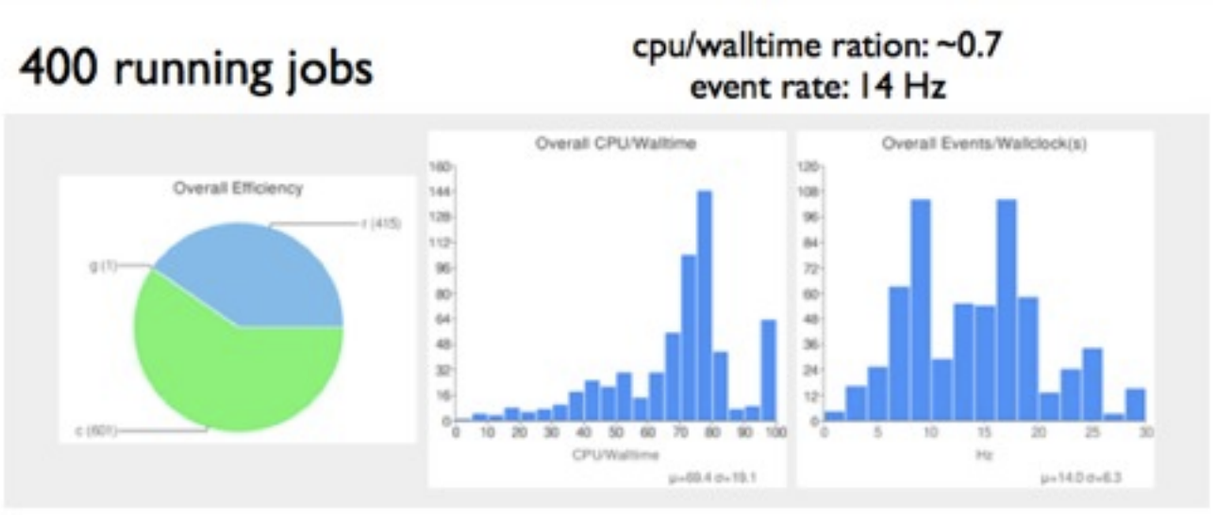

Journées 'Grilles France', Octobre 2009

The good, the bad, the ugly...

Jobs Success	Jobs Failure	Jobs Total	Job Success Rate	Job Event Rate	Total Events	Event Fraction	Site Rank (Max 6)	ATLAS Disk
8941	797	9738	91.8%	8.1	411414711	0.224	5	319
1272	992	2264	56.2%	5.4	34013217	0.019	1	270
1030	90	1120	92.0%	7.0	20781526	0.011	3	90
3522	530	4052	86.9%	10.4	185342695	0.101	5	170
699	0	699	100.0%	11.5	57261465	0.031	5	70
5150	548	5698	90.4%	10.0	214814145	0.117	5	250
13926	204	14130	98.6%	12.9	466536165	0.255	6	475
1804	293	2097	86.0%	12.8	59412253	0.032	4	210
7764	124	7888	98.4%	36.1	377656740	0.206	6	230
350	837	1187	29.5%	1.8	3180094	0.002	0	160
1874	196	2070	90.5%	5.9	2255136	0.001	3	30
46332	4611	50943	90.9%	14.6	1832668147			

- But still not all sites ready
  - It is the responsibility of sites to be ready
  - ATLAS computing operations will not let users suffer from bad sites
- Stephane will present more details on support for physics and performance groups later

Journées 'Grilles France', Octobre 2009

SL5 Migration

- ATLAS want to see site resources on SL5 as soon as possible
  - Many large sites have already migrated
- We hope to move to SL5 native built kits very early in 2010
- We reiterate the undesirability of killing jobs based on VMEM or RSS, especially on 64 bit base OS
  - It seems that sites, by and large, do not do this any longer
    - This is good for the moment, but we need development to get better solutions (TMB?)

Conditions DB Access

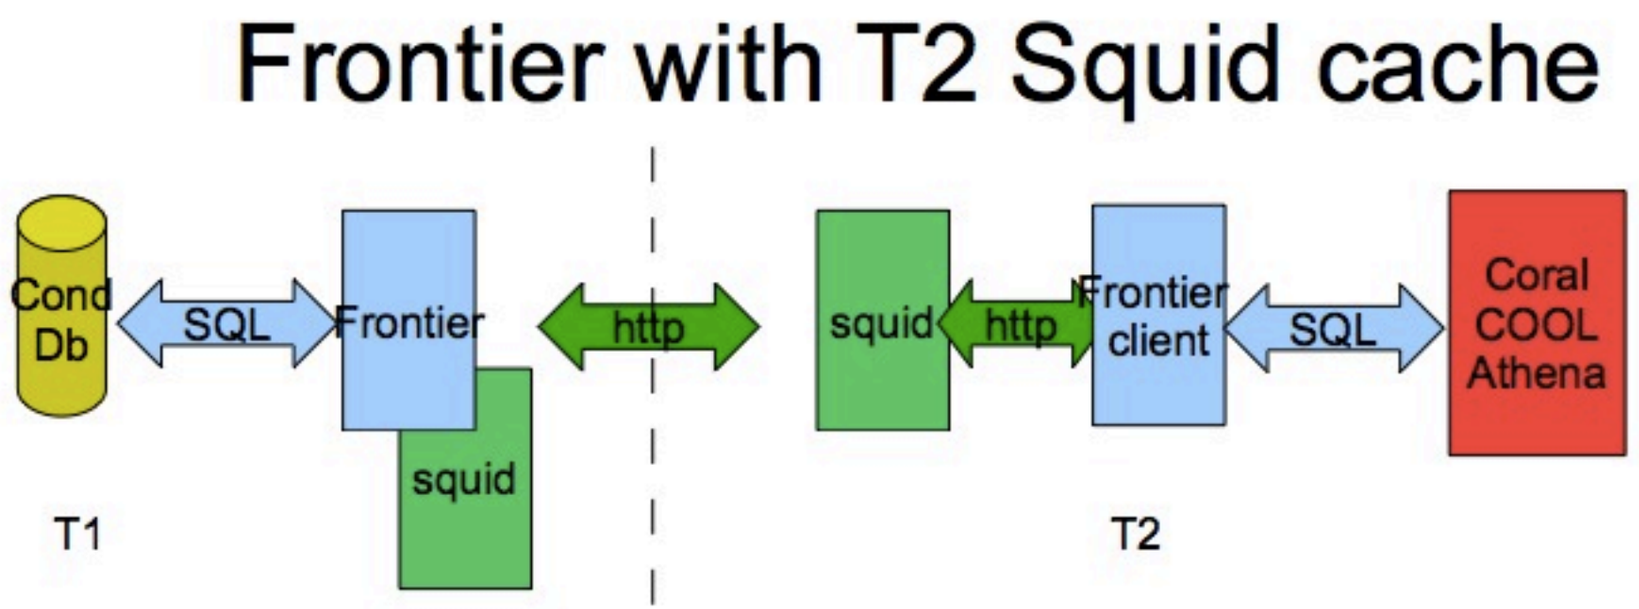

- Fronter server with Squid(accelerator) at FZK using local Oracle
- Squid caching web proxy at the T2

Journées 'Grilles France', Octobre 2009

Frontier

- Need local RAC but not all T1s must have Frontier
  - CMS says "significant" maintenance effort
- Current plan is US,DE,UK,ES,CA +FR?
  - heavy access will scale with number of RAC servers – will not deter other T1s
- T1 RAC direct access used for heavy access patterns
  - each event need new CD, e.g. ID alignment

Squids

- Industry standard web caching service
- Non-standard config to allow large objects to be cached
- cache size depends on TTL and access pattern
  - expect most benefit already with 10GB
  - probably 100GB is sufficient
- Installation instructions linked from
  - https://twiki.cern.ch/twiki/bin/view/Atlas/T2SquidDeployment
  - machine specs low – deploy on what you have and check load
  - 64bit rpms for sl4 and sl5. Tarball make and install.
  - rpms recommended for ease of update

DB Hot Topics

- Squid deployment well advanced in DE and US; started in UK and FR
- It doesn't take long, but needs pushed
  - https://twiki.cern.ch/twiki/bin/view/Atlas/T2SquidDeployment
- We also need access to conditions flat files
  - We have a plan to manage this, through ToA and a local setup.sh
  - There are some issues we need to chase:
    - SL5 DPM access via rfio seems not to work (still need libshift patch for DPM!) - *Work around found*
    - Pre 1.9.3 libdcap.so broken
- SAM test for conditions access
- Ganga Robot test: http://gangarobot.cern.ch/20091006_06/index.html

Batch System Setup

- We want an lcg-CE infrastructure maintained for 2009-10
  - CREAM is too new and condor-g submission not in good shape
- We will have a new activity, group analysis, which we will try to make appear as much as possible like user analysis
  - Picked up by /atlas/Role=pilot
- Results of ADC review conclude panda is 'primary' backend for analysis
  - T2 sites should substantially increase relative share for panda analysis

Storage System Setup

	/atlas/Role=production	/atlas, /atlas/*	Others
ATLASDATADISK	Read/Write	Read	No Access
ATLASMCDISK	Read/Write	Read	No Access
ATLASPRODDISK	Read/Write	No Access	No Access
ATLASDATATAPE	Read/Write	No Access	No Access
ATLASMCTAPE	Read/Write	No Access	No Access
ATLASGROUPDISK	Read/Write	Read	No Access
ATLASSCRATCHDISK	Read/Write	Read/Write	No Access
ATLASLOCALGROUPDISK	Read/Write	Read (Recommended)	No Access

- The single most important thing is that users cannot recall from tape - must happen at T1s
- Space token allocations to follow...

With other Issues...

	Castor	dCache	DPM	StoRM	BeStMan
recommend version	2.1.8	>= 1.9.4 (1.9.5-2)	>= 1.7.2	?	?
tape protection	by DN	by FQAN			
hot files	✔	various ways	mgmt. tool from GridPP	✔	✔
checksum	to be tested	✔	✔	to be fixed	to be tested

Journées 'Grilles France', Octobre 2009

ATLAS UAT: Goals

- ATLAS have planned a large scale User Analysis Test from 21-23 October (delay until 28-30 Oct?)
- The aim is to get a measurement of the "efficiency", time to ntuple, event processing rate, etc.
  - Essentially the same as hammercloud, but with users actively involved
  - Potentially includes a large amount of file movement with dq2-get.
    - We will have to see if this causes interference with other operations.

UAT Timetable

- Now:
  - Sample datasets are being moved to sites now
  - Pre-test with hammercloud
  - Test metric gathering
  - Get users ready
- 21-22 October: Run jobs
- 23 October: Collect outputs
- 26 October - 6 November: Analyse results
- https://twiki.cern.ch/twiki/bin/view/Atlas/UserAnalysisTest

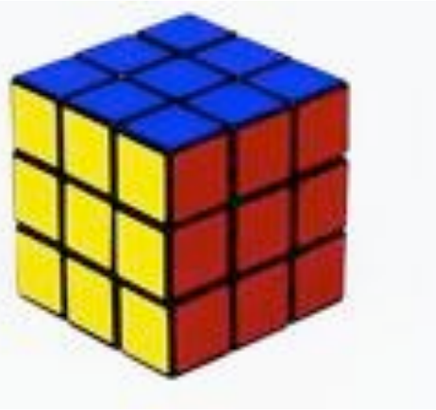

Communications

- Computing strategy is defined and discussed during SW&C weeks
- ADC Operations weekly meetings is the place to discuss issues and to make operational decisions
- Daily Operations hot issues
  - Daily Operations phonecons during major and peak activities
  - Shift(s) : ADCoS and ADC@P1, Virtual Control Room (Skype chat)
  - The first line of expertise is provided by ADC expert on-call
- NO decision making in e-mail discussion or/and coffee, corridor, etc "meeting(s)"
- Activity Twiki page(s)
  - Activity scope, description, general info, and URLs
- Announcements via Operations e-mail list with CC:eLog
  - e.g., activity start/end time, changes, etc

Summary

- A great deal of progress has been made since STEP09
- However, there are still a lot of things to deploy, test and continue to hone
  - SL5 at sites not yet upgraded
  - T1 Frontier & T2 Squids
  - Tape protection at T1s
  - Disk pledges deployed and in correct tokens
  - Physics groups at T1s and T2s
  - Analysis activity
    - Hammercloud and UAT