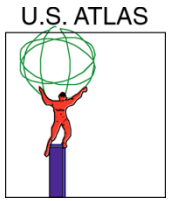

Review of Facility Requirements and Overview

Michael Ernst, BNL

US ATLAS Facilities Workshop at SMU

October 11, 2011

U.S. ATLAS

Outline

- LHC Schedule and Computing Challenges
- Status of computing in ATLAS
- Installed resources and projections for 2012 and 2013
- Overview and developments since last Facility Meeting
- Analysis Performance & Efficiency
- Summary

U.S. ATLAS

From last week's ATLAS weekly

4 fb^{-1}

4.316 fb^{-1} delivered
4.054 fb^{-1} recorded

Maximum Luminosity Delivered in one day	130.71 pb^{-1}	Sunday 02 October, 2011
Maximum Luminosity Delivered in 7 days	532.04 pb^{-1}	Tuesday 20 September, 2011 - Monday 26 September, 2011

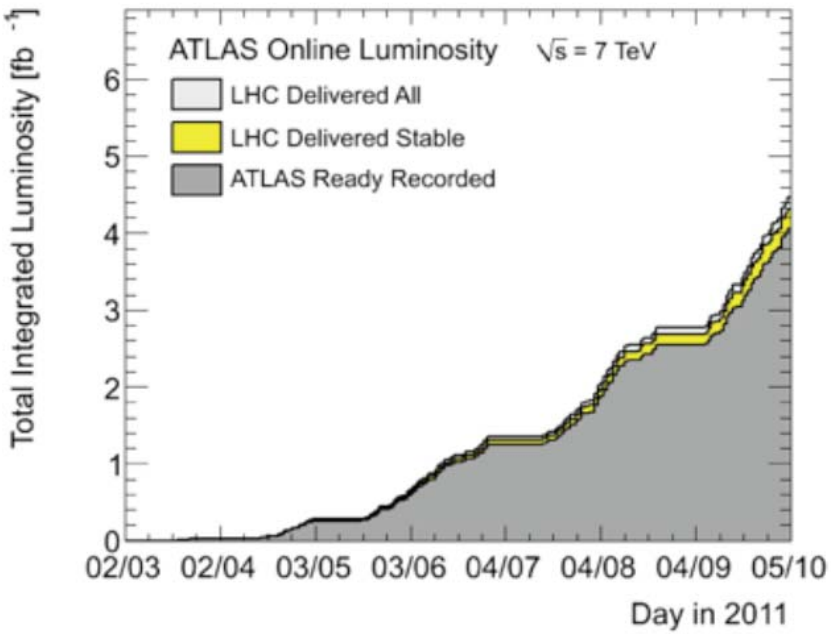

In 2010 the goal was to collect 1fb^-1 in 2011

U.S. ATLAS

The 2011 Schedule

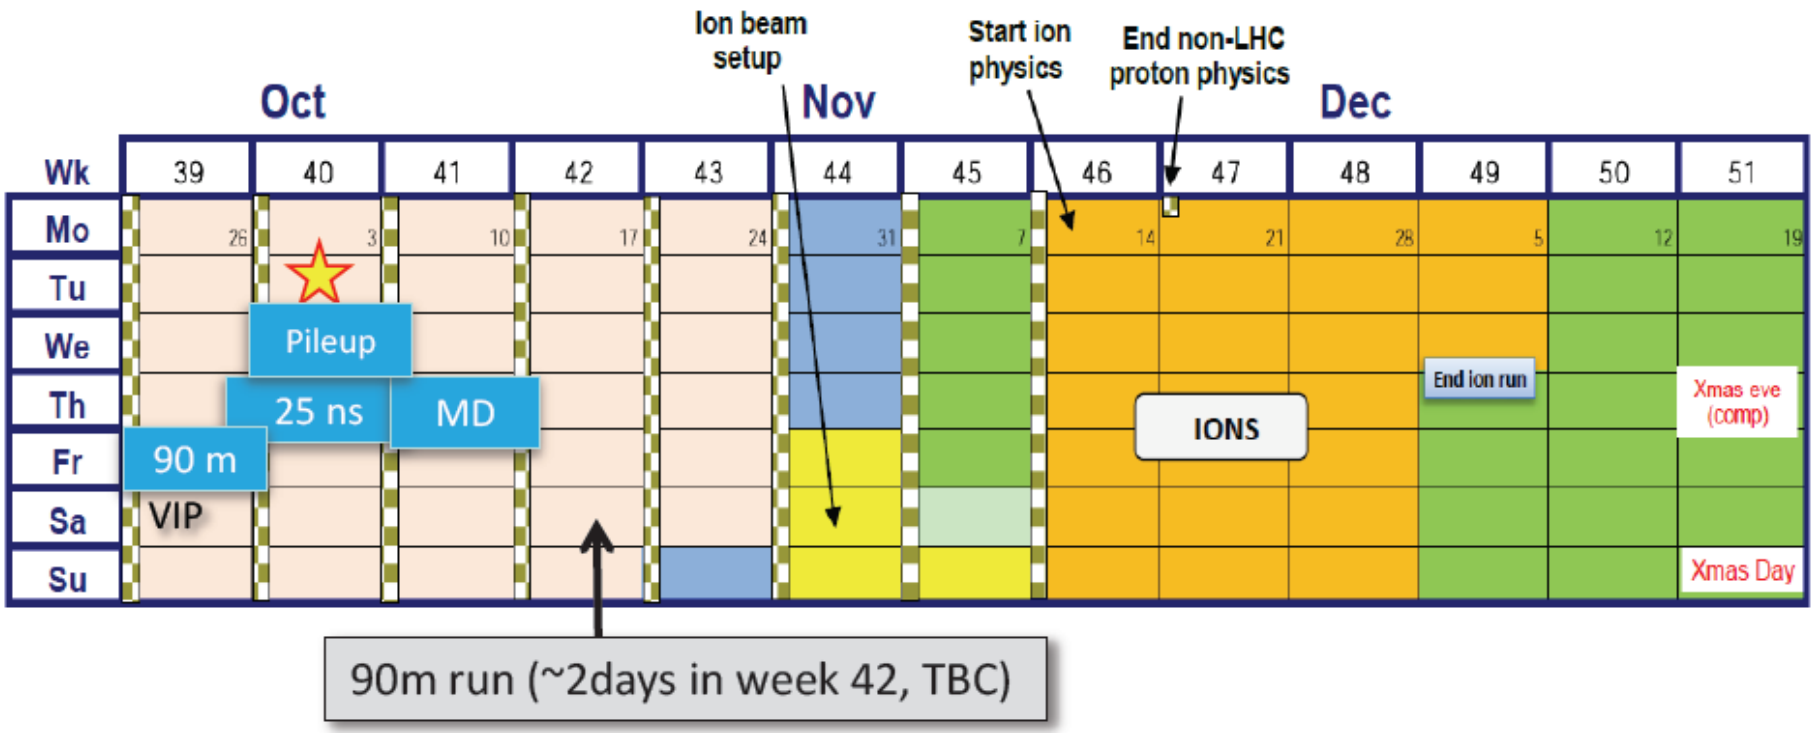

Trade off: Integrated Lumi ↔ pile up

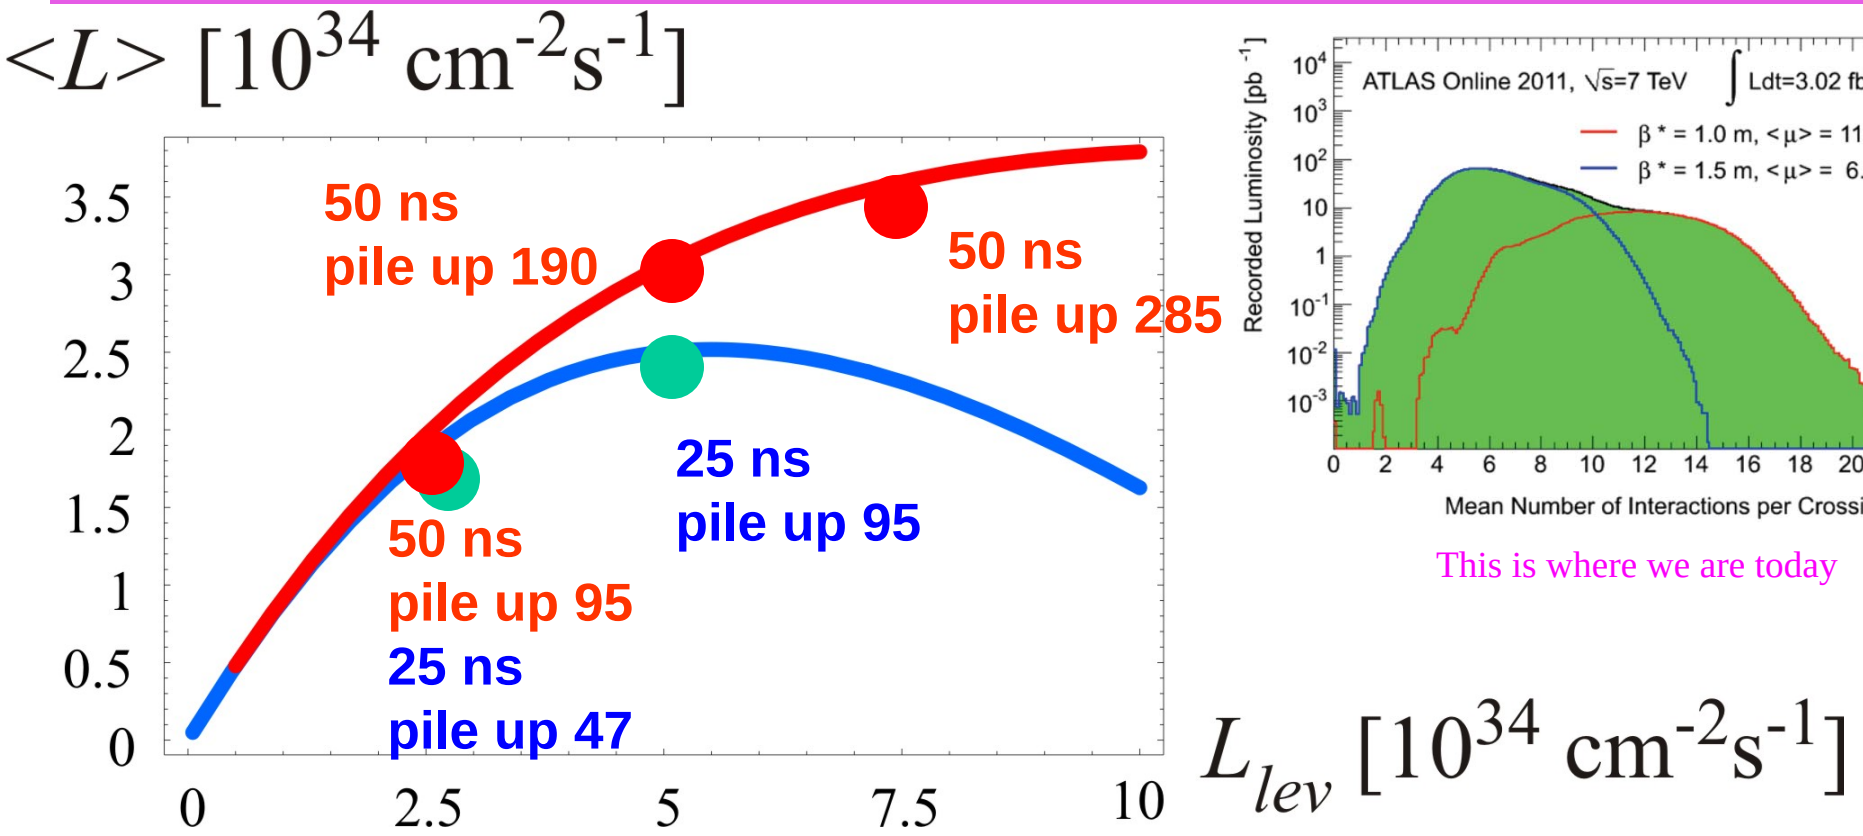

This is where we are today

Roughly for 2 times more integrated luminosity 4 times the pile up

U.S. ATLAS

Improvements in CPU Time and Data Sizes

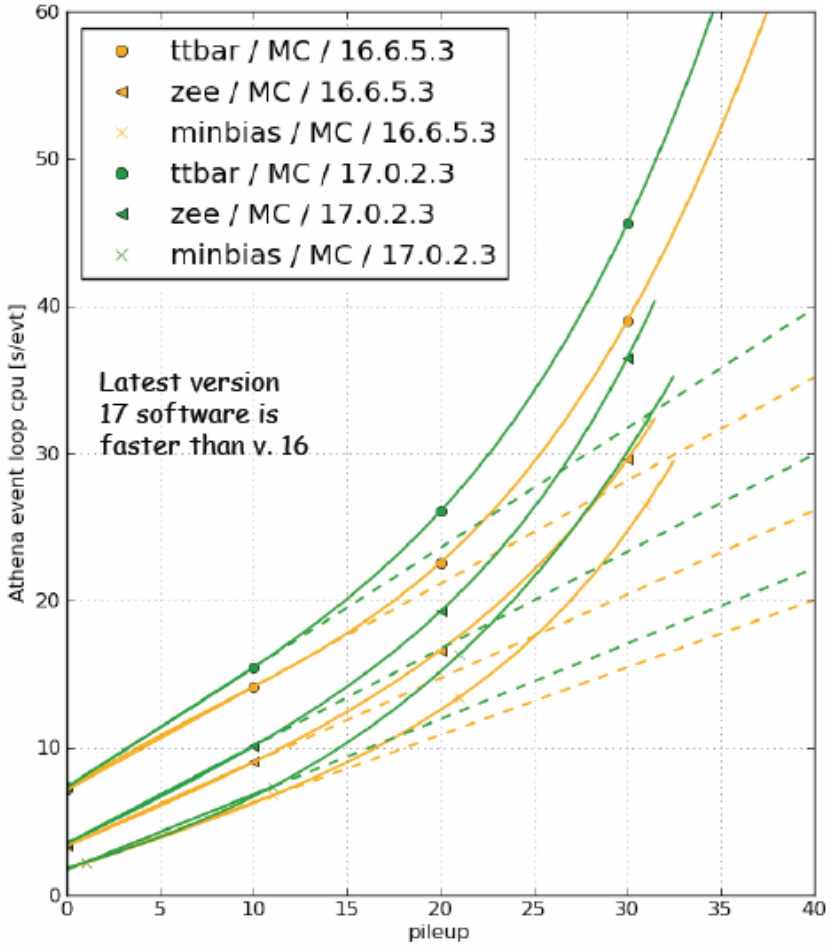

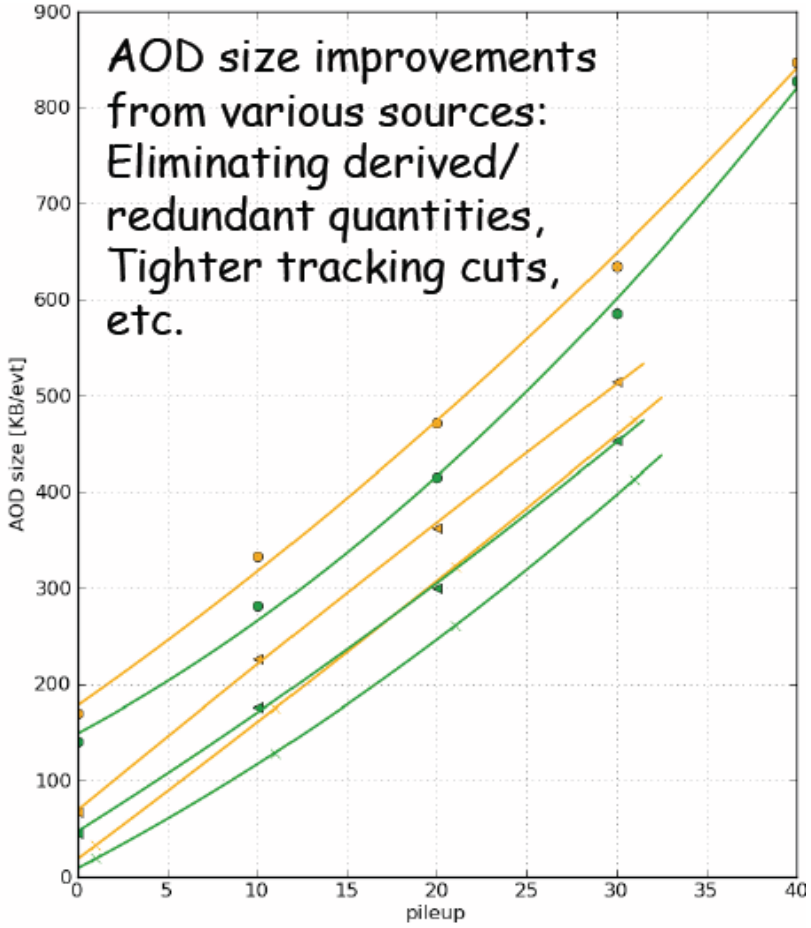

Effect of pileup on reco. time and AOD size for various ATLAS software releases

Physics drives Computing
(J. Shank at September LHCC)

- Our Computing Model and resource usage is driven by our physics goals
  - We don't want computing to be a bottleneck to physics publications
    - We have achieved this so far in the LHC data taking era
      - For example, some recent results presented at the EPS conference were using data that had been taken three weeks before the conference.
- Changes to our data distribution plan made better use of our facilities and allowed us to take data at a higher trigger rate.
- One ingredient allowing us to achieve this is a lot of hard work that went into improving our reconstruction CPU time/event and event sizes in the face of increasing pileup
- Work continues on these improvements and we also are improving our simulation time and our fast Monte Carlo, which now reproduces the data well enough that physics groups are using it for publications in preparation now. (up to 10 times faster than our Geant4 simulation)

U.S. ATLAS

New Data Distribution Policy

- We have eliminated bulk ESD as an analysis format.
  - Just maintain a rolling buffer at T1's of a small fraction of the ESD for special studies
  - Helped successfully guide users to more space-efficient derived physics data.
- RAW (compression)
  - We place one full copy of the RAW data on disk distributed over the ensemble of Tier 1s, but we now compress the RAW, making the data 60% of the uncompressed size.
- Physics analysis is done from AOD, dESD and Ntuples produced by physics groups.
  - Distributed to T1s, and dynamically to T2
- Overall, our data distribution has been flexible and we adapt it in order to optimize our physics output and make maximal use of available CPU and disk resources.

CREM requested T2s to hold 2 primary copies of mc10_7TeV*merge datasets

U.S. ATLAS

Tier-2 CPU Usage

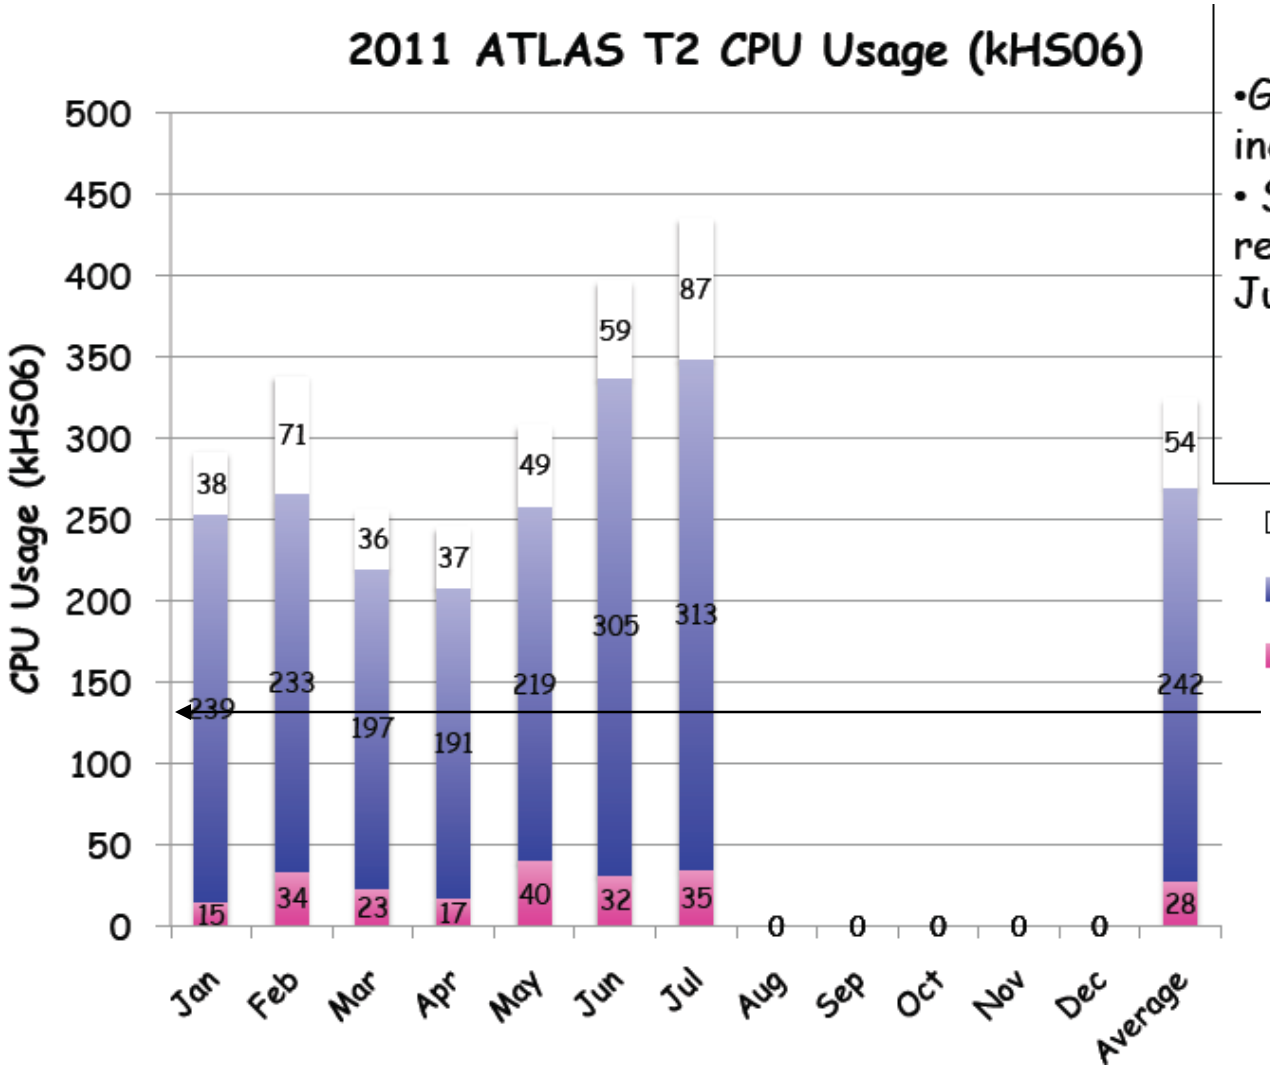

Notes

- Group/User activity increasing
- Still had not completed reprocessing campaign as of 31 July
  - Expect further increased activity after this reprocessing

Installed T2 capacity in the U.S.

U.S. ATLAS

Dynamic Data Placement

- In use in ATLAS for over 1 year now.
- Derived physics datasets are copied to T2 disk only when there is a demand
  - The first time a user submits a job using a dataset as input, that data is copied to a T2 (selected by brokering algorithm)
    - Subsequent jobs will go to that T2 up to a set number
    - Further jobs submitted using that dataset trigger more replicas to other T2's
- Some T2s were being under-utilized (both CPU and Disk)
  - Algorithms (both brokering and data placement) are being tuned to fix this problem
- AOD, Ntuple and DESD are popular ➔ we now are doing some pre-placement of these to level CPU usage across sites
- We are still in the era where physics group's Ntuple samples are small enough (~ few TB) that users can copy them to local resources to do analysis.
  - We expect this to not be possible soon (2012) and users will have to use the T2s for this part of analysis also, so user activity will increase
- Current (20-09-2011) T2 disk usage: 20 PB
  - 35 PB pledged until April, 2012.
    - We expect to fully utilize this space by then

More about PD2P and Caching in Torre's talk

U.S. ATLAS

Updated Resource Request for 2012/2013 (J. Shank @ LHCC)

- Only change from our March, 2011 estimate is in CERN CPU

CPU [kHS06]	2011	2012	2013
CERN	74	73→111	111
Tier-1	202	259	280
Tier-2	275	295	321
Disk [PB]			
CERN	7	9	10
Tier-1	22	27	30
Tier-2	35	49	56
Tape [PB]			
CERN	14	18	18
Tier-1	28	36	40

Baseline for Pledges

US to provide 23% of the total

- Our estimates from last March still stand.
  - Event sizes, reconstruction times, Pileup, trigger rate, etc. have all changed and we expect that to continue.
- We adjust our computing model to fit within the resource constraints we have for 2011 and 2012
- We expect to be able to maintain the current trigger rate with an increase of resources only at the Tier 0

U.S. ATLAS

Usable and Pledged Capacities in the U.S.

2012 Pledged capacities at the US ATLAS Facilities

Site	CPU [HEPSpec 2006] 2012 Pledge	CPU [HEPSpec 2006] Installed Sep, 2011	DISK [TB] 2012 Pledge	DISK [TB] Installed Sep, 2011
Tier-1	60,000	58,000	6,300	6,100
AGLT2	12,500	36,163	2,200	1,910
MWT2	12,500	36,840	2,200	1,302
NET2	12,500	19,035	1,648	1,100
SWT2	12,500	23,220	2,200	1,260
WT2	12,500	15,816	2,200	1,663
Total	**122,500**	**189,074**	**16,748**	**13,335**

2013 Planned to be Pledged capacities at the US ATLAS Facilities

Site	CPU [HEPSpec 2006] 2013 Pledge	CPU [HEPSpec 2006] Installed Sep, 2011	DISK [TB] 2013 Pledge	DISK [TB] Installed Sep, 2011
Tier-1	63,000	58,000	7,000	6,100
AGLT2	13,400	36,163	2,500	1,910
MWT2	13,400	36,840	2,500	1,302
NET2	13,400	19,036	2,500	1,100
SWT2	13,400	23,220	2,500	1,260
WT2	13,400	15,816	2,500	1,663
Total	**130,000**	**189,075**	**19,500**	**13,335**

Site	Tape [TB[2011 Pledge	2012 Pledge	2013 Pledge	Installed 9/11
Tier-1	6,900	8,300	9,200	6,900

Up-to-date summary at
https://docs.google.com/spreadsheet/ccc?key=0AnH5nTdCzYK3dFdrZmhQcnMxelRNSWZGS0FhSlZWbUE&hl=en_US#gid=0

U.S. ATLAS

Service Integration Activities

- **Besides Facility operations at high performance and high reliability we are working on several service integration initiatives**
  - **UIUC joining MWT2**
  - **Moving from "shared area" to CVMFS for S/W distribution and Conditions Data access**
    - **Following difficult start have developed lots of expertise**
    - **Have now several CVMFS-based PanDA sites in production**
  - **Federated Data stores w/ xrootd**
    - **Focused WG has made excellent progress lately**
    - **All tiers participating**
    - **Still a lot of work remaining (authentication, performance, …)**
  - **Cloud Computing**
    - **Lots of ideas and initial steps (e.g. dynamic WN capacity extension, T3, …)**
    - **Need to converge – ATLAS-wide - on a program of work**
    - **We welcome Val from LBNL and Cui et al from Fresno to join our activities**
    - **John Hover will kick off the discussion later today**
  - **AutoPyFactory – A new infrastructure for pilot submission**
    - **Deployment already in progress in the US**
    - **Pilot submission now under full control by Facilities**
    - **Lots of new features – John Hover will present tomorrow**
  - **WAN performance optimization and monitoring with perfSONAR-PS**
    - **Our initiative in collaboration w/ I2 has finally paid off**
    - **LHCOPN monitoring fully implemented at 10 T1s and CERN within only ~3 months**
    - **Shawn and Jason will talk about recent achievements**
- **Rob will provide more information about facility integration in the next talk**

U.S. ATLAS

Summary of Tier-2 "Issues"
Based on Discussion Rob and I had with PIs et al

Very useful thanks to excellent preparation and openness

- **Capacity Provisioning according to defined milestones**
- **CPU/Disk capacity balance (too much CPU in 2011/2012)**
- **Limited WAN b/w to storage affecting job efficiency**
- **Analysis performance: number of concurrent analysis jobs below 50% of total number of job slots**
- **Facility operations critically dependant upon 1 person**
- **Power capacity and stability**
- **Delayed implementation of agreed upon new services and updates**
- **Resilience of core services against component failures**

➢ **Together w/ sites will track progress**

➢ **Will have these discussions every ~6 months**

U.S. ATLAS Institutional Computing Policy

- Applies to “Tier-3” sites with different functionalities, e.g.
  - Follows ATLAS-wide policy, adding specifics for U.S. region
  - ATLAS Analysis Site or “Tier-3g”
  - Site with full grid services or “Tier-3gs” (Tier-2 like) contributing to production and/or analysis
  - DDM site with SE to receive data via subscriptions
- Site requirements and obligations
  - Computing, storage, networking resources, effort, uptime and response time in case of problems, cloud support, official affiliation with ATLAS
  - Site Certification Process
    - Depends on functional category
    - Site validation (coordinated through Facilities Integration Pgm)
- Status change and termination
- Document in preparation
  - Important to have crisp requirements and official process for site inclusion in place to avoid disruption of ATLAS workflow

Network Evolution – "Flattening" today's Hierarchical Model

- This is what we have now
- It is a push model
- And has been very successful
- But has also drawbacks

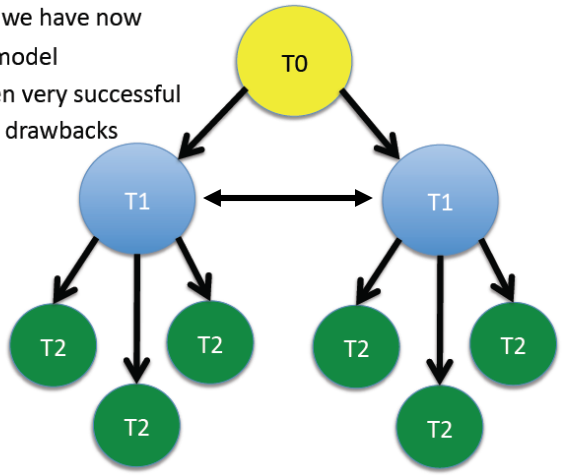

Hierarchical (Monarch) Model (restricted)

- Data can be pulled from anywhere
- Needs another network
- Not more bandwidth
- Different topology

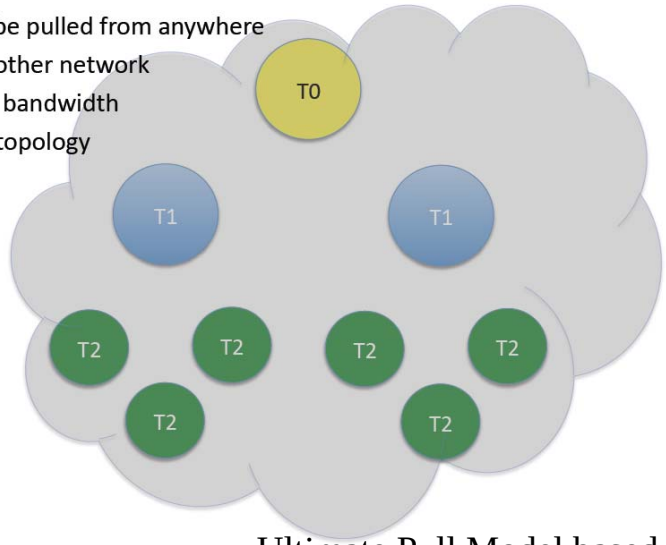

Ultimate Pull Model based on unrestricted access across sites

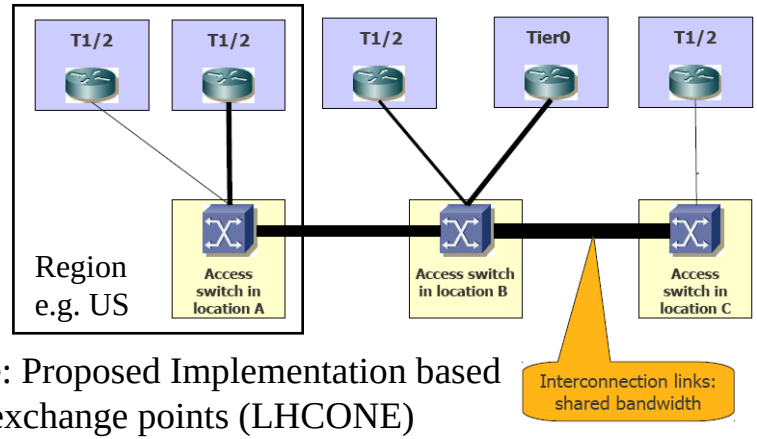

Compromise: Proposed Implementation based on regional exchange points (LHCONE)

U.S. ATLAS

Networking - LHCONE

- "Prototype" is shaping in several regions of the world
  - North America (US & CA, USLHCNet), Europe, Asia
  - Connectivity between R&E Networks in Europe and the US
  - Growing number of sites are connected and start making throughput measurements
    - We need to keep in mind that LHCONE is implemented as an overlay network
    - There is no new bandwidth – traffic is using pretty much the same infrastructure
      - In some cases less throughput than w/ General IP
    - LHCONE infrastructure allows engineering of traffic flows
  - Still much work in progress and far away from production
    - Different providers have different perspectives of LHCONE
    - No imminent problem of saturating existing R&E IP networks
    - Constructing independent IP network for HEP has limited value
    - All R&E providers are investigating in initiatives (DYNES, DICE, GLIF) to better control large science flows -> create global initiative
    - LHCONE should be technically ambitious/far beyond IP networks
      - Should engage with early adopter user communities in a managed way
      - To be coordinated, for the LHC, with the experiments

More in Shawn's talk

Analysis jobs at T1s & T0 (CERN)

Slides presented by Chair of the ATLAS International Computing Board in 07/2011

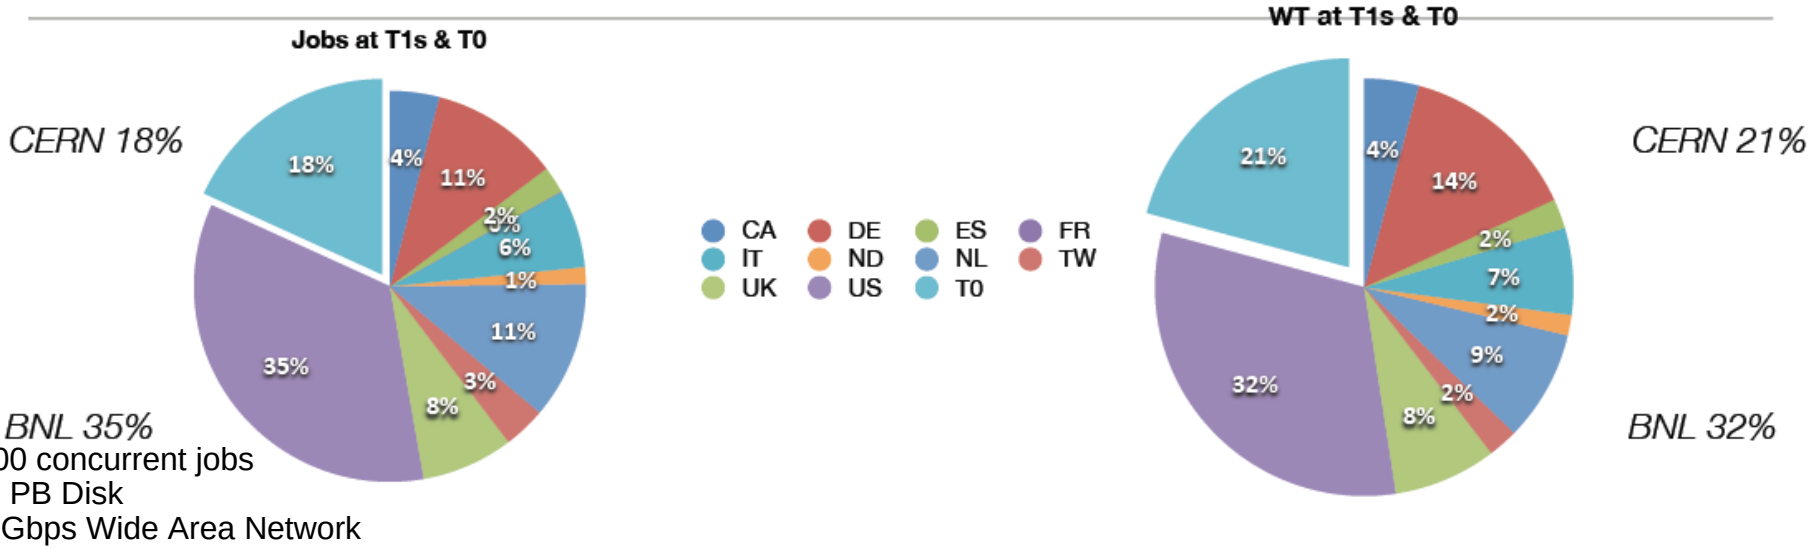

7500 concurrent jobs
9,6 PB Disk
60 Gbps Wide Area Network

Relative distribution of analysis jobs compared to CPU pledges

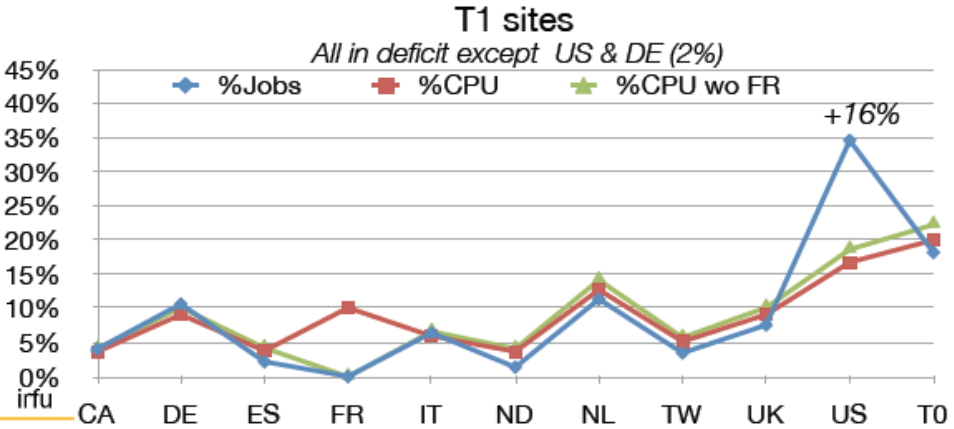

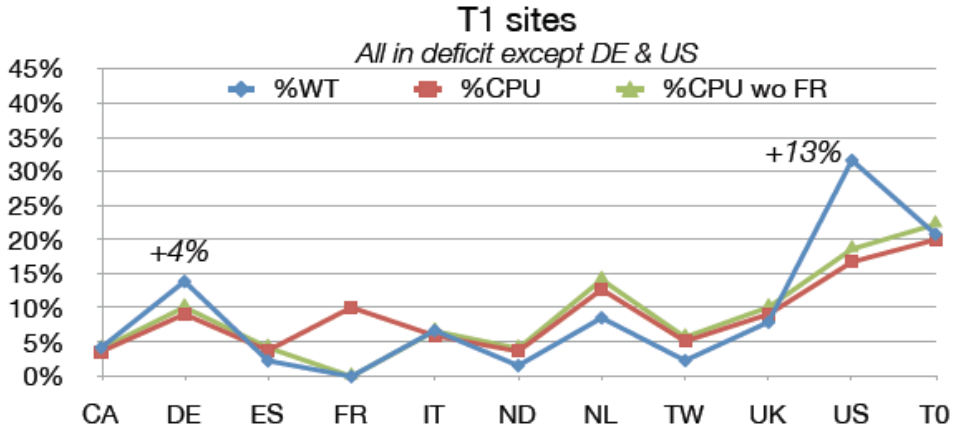

- **The US ATLAS Tier-1 Center is the most attractive ATLAS Analysis Site**

Analysis share by Tier-2 Site (as of July 2011)

NET2 (w/o HU) working on improvements

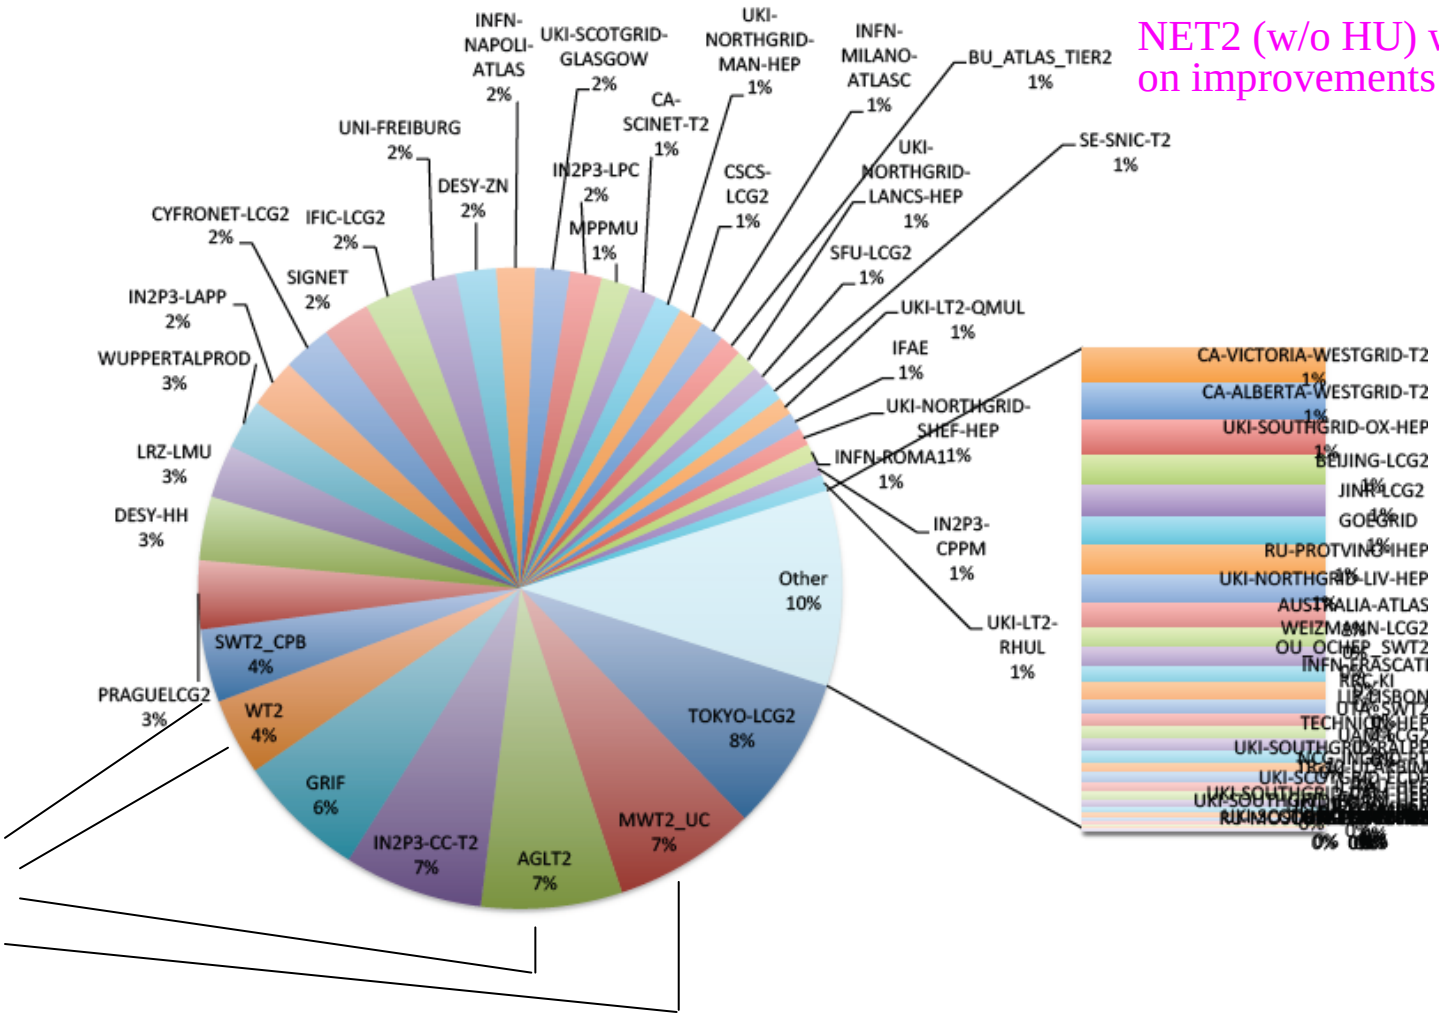

4 (out of 5) US T2 Sites

~50% of analysis jobs in ~10 sites

Relative excess/deficit wrt CPU share (T1s & T2s)

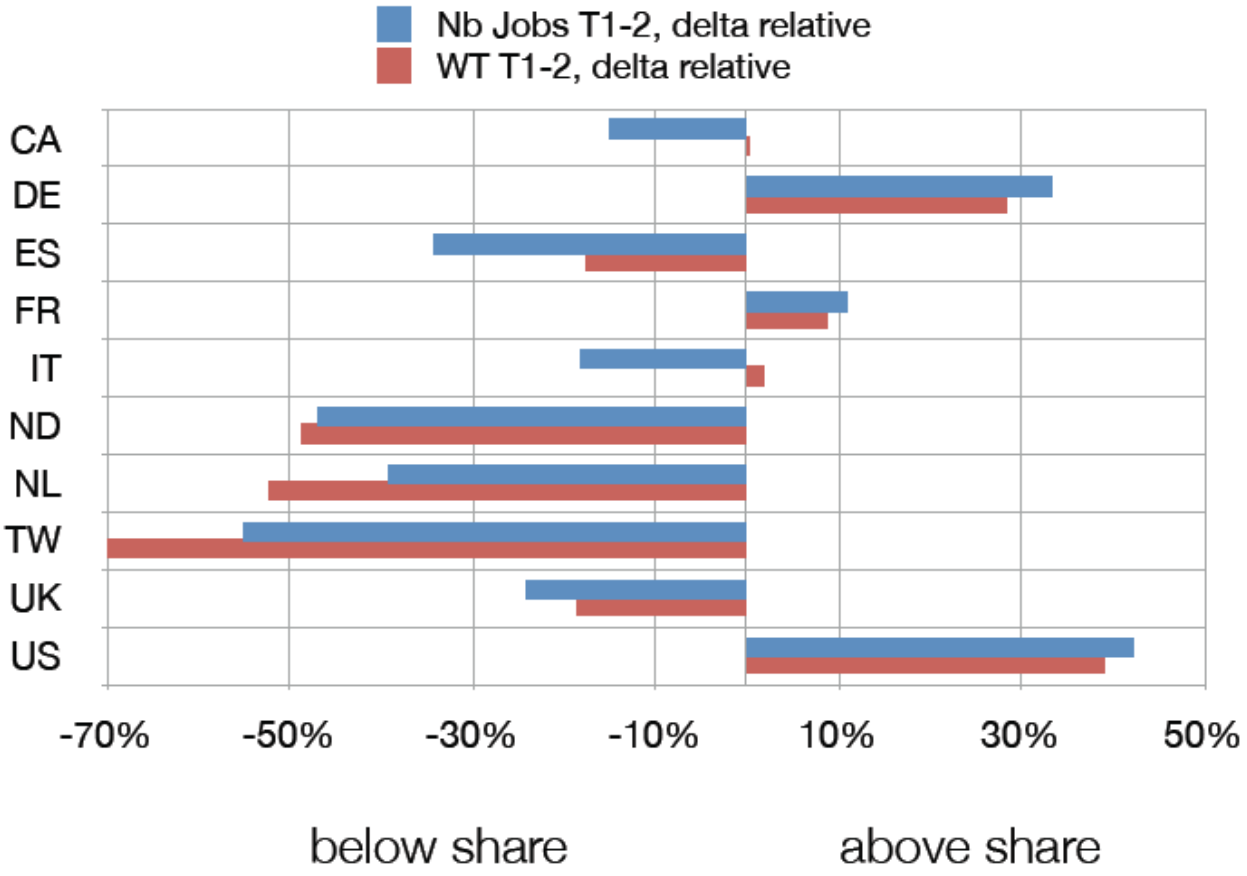

U.S. ATLAS

Analysis Efficiency
(J. Elmsheuser @ CREM Meeting)

Introduction

https://indico.cern.ch/getFile.py/access?contribId=5&resId=0&materialId=slides&confId=156964

Short analysis of the ATLAS distributed analysis job efficiencies:

- All user analysis jobs and GangaRobot functional tests (AFTs)
- per month June, July, Aug, Sep 2011
- GangaRobot functional tests (AFTs) in later plots split by T0/1, T2, T3
- GangaRobot AFTs:
  - 4 tests: 15.6.14 (MC AOD), 16.0.3 (MC AOD), 16.0.3 (Data dESD), 16.6.7 (Data AOD)
  - 24h HammerCloud tests, keeps always 1 job in ANALY_* queues
  - after 24h HammerCloud test is stopped, jobs are cancelled and new HammerCloud test is started

Plot References

References:

- `http://dashb-atlas-job.cern.ch/dashboard/request.py/dailysummary`
- 1st set of plots:
  - *Sites:* All T0-T3, *Activities:* all or gangarobot, *Granularity:* daily, *Time Rage:* Jun, Jul, Aug or Sep 2011 (-26th)
  - *Select:* Success/Failure category
- 2nd set of plots:
  - *Sites:* T0-T1, T2, T3, *Activities:* gangarobot, *Granularity:* daily, *Time Rage:* Aug or Sep 2011 (-26th)
  - *Select:* Success/Failure category

Caveats:

- Historic dashboard classifies *cancelled* jobs as *failed*
- $\rightarrow$ lowers GangaRobot site efficiencies by $\sim$1-3%
- Efficiency could also be higher for ALL analysis jobs
- Dashboard team has been notified and promised to address the issue in the next days
- Downtime handling is not fully included in efficiencies

ALL ANALYSIS SEPTEMBER 2011

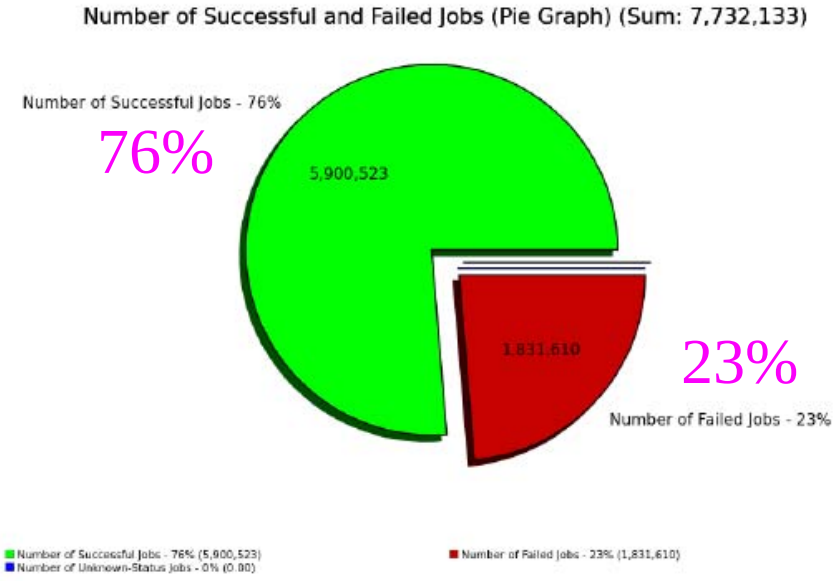

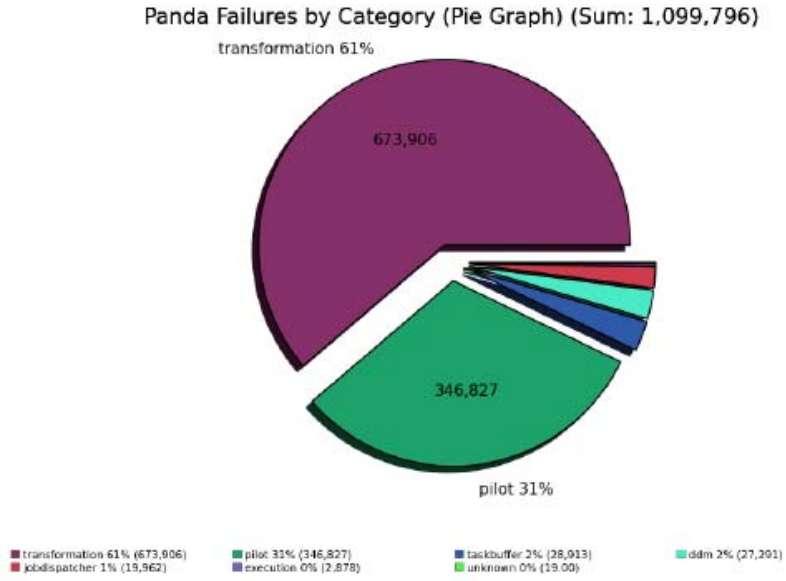

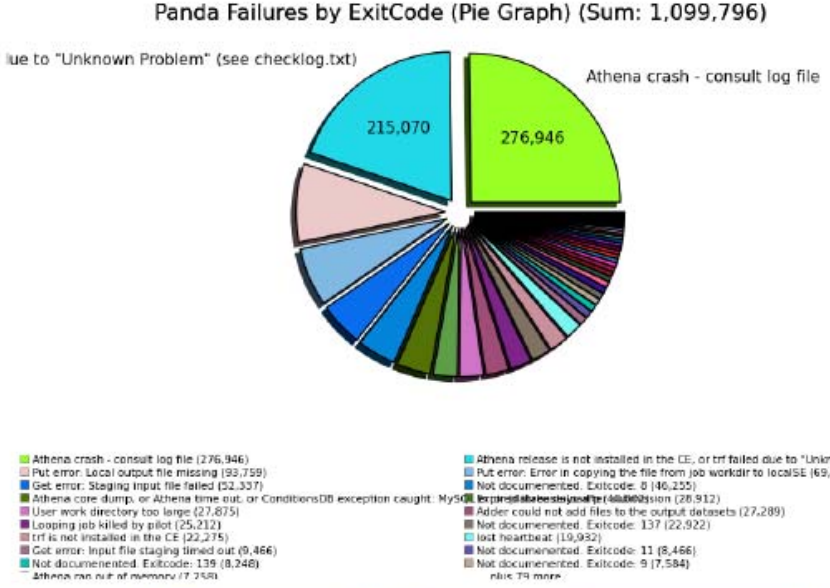

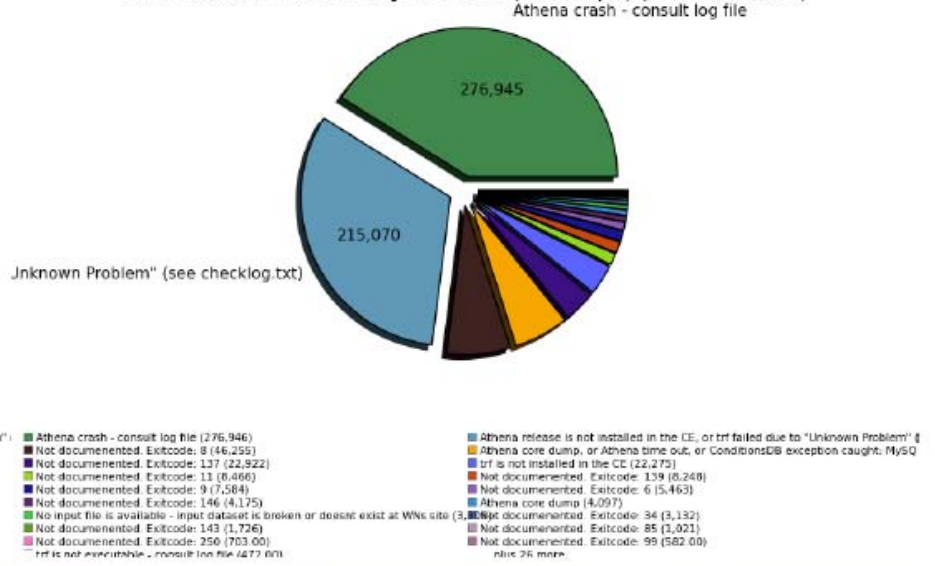

GANGAROBOT ALPHA SEP 2011

http://atladcops.cern.ch:8000/drmon/crmon_TiersInfo.html

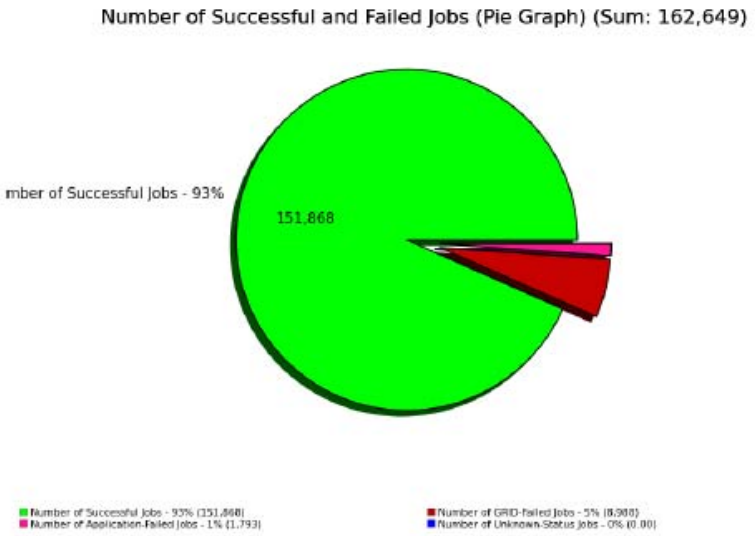

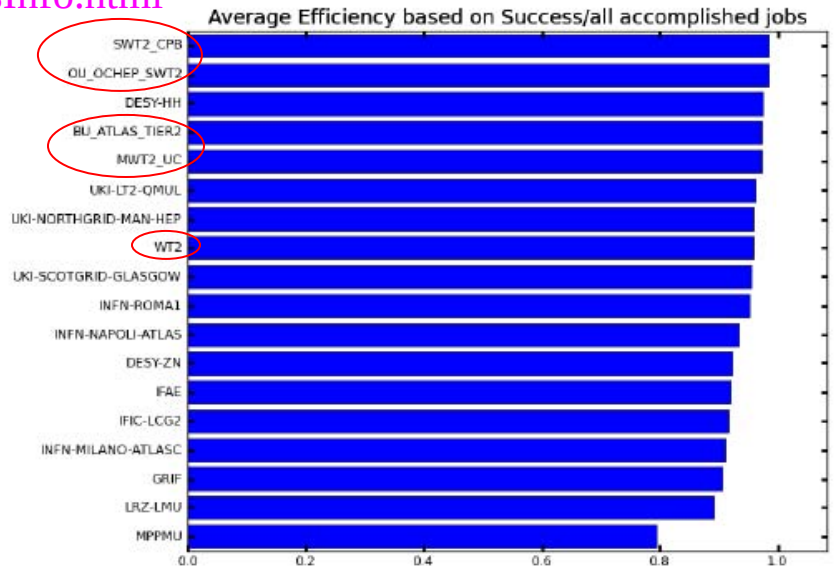

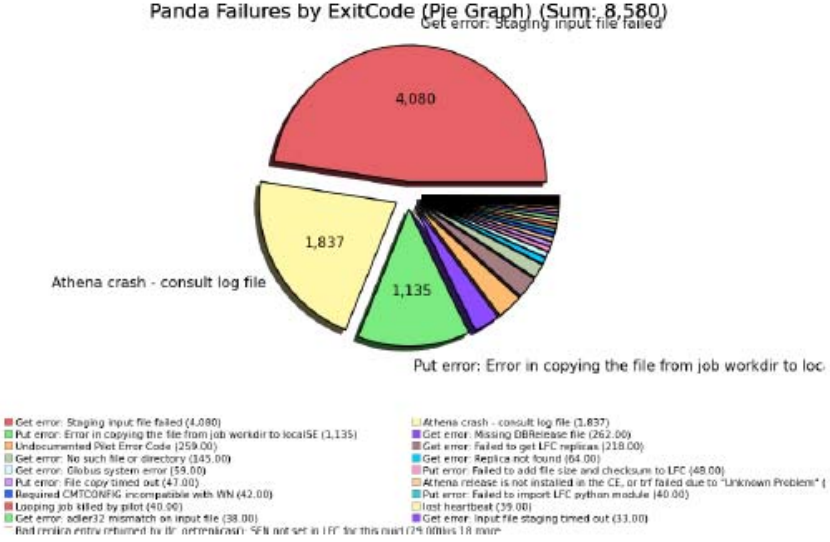

T1	tier	%	T2 Group	tiers status
US	AGLT2_DATADISK	50%	charlie	registered: 6 ready: 6
	ILLINOISHEP_DATADISK	100%	delta	
	MWT2_DATADISK	50%	alpha	
	NET2_DATADISK	50%	alpha	
	SLACXRD_DATADISK	50%	alpha	
	SWT2_CPB_DATADISK	50%	alpha	

Legend:

- ** ready status
- ** waiting status
- ** not validated status
- ** site in downtime

Associated tiers groups are highlighted with different colors.

Total Tier2s in 8 clouds: 64 (ready: 62)

GANGAROBOT CHARLIE SEP 2011

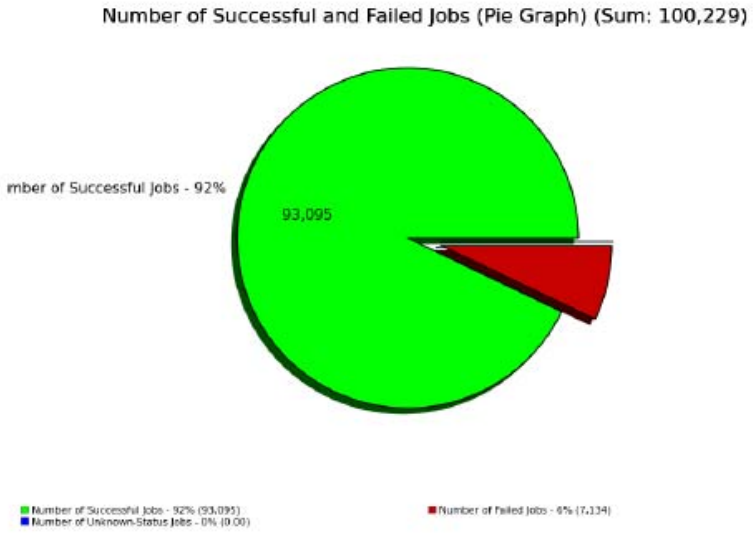

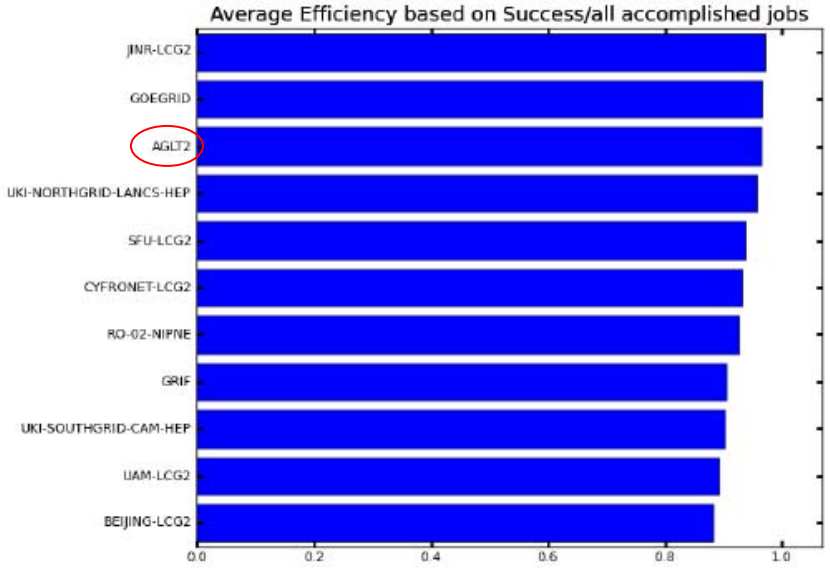

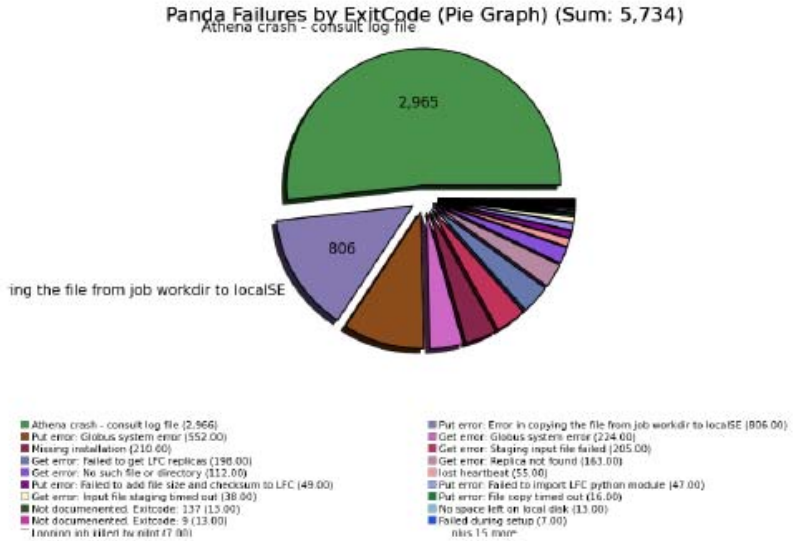

Comments and Conclusions I

Comments:

- AFTs (GangaRobot) job efficiencies about constant with slight improvement over the last 4 months
- Efficiency is ∼1-3% higher due to *cancelled*-job mis-classification
- All analysis jobs error rate also constant (but contains user code problems)
- Direct input file access problems (dcap/root) or conditions data access problems hidden in *Athena crash* category
- Differences in job efficiencies between T0/1, T2 and T3 visible
- Main crash reasons of AFTs:
  - Athena crash (direct input, cond), stage-in, stage-out problems
- Some sites are not running jobs from time to time: queue is brokeroff, no pilots running, site is totally full, ATLAS share at site has been fully used for a certain time window, or the site is blocking jobs with my name ?
- Downtime handling is problematic, not accounted for in statistics
- Setting ANALY_* queues offline still often manual

Comments and Conclusions II

Further comments:

- alpha, bravo, charlie sites are basically equivalent, only delta is showing lower efficiency
- Have improved Grid analysis efficiency compared to beginning of the year by aggressively black-listing sites
- But SE problems can always happen with this large number of sites, files and jobs
- How to improve on systematic intrinsic SE troubles ?

Possible solutions:

- More Pilot retries on clearly classified error categories (stage-in, stage-out) => Local Site Mover (LSM) as "Insulation Layer" between Pilot and SE
- CVMFS for conditions data flat file access
- What about other errors ?

U.S. ATLAS

Summary

- Computing was a big success as enabler for physics, on its own metrics but also on the ultimate metric of timely physics output
- The Facilities, the Tier-1 and the Tier-2' centers, have performed well in initial LHC data taking and analysis
  - Production and Analysis Operations Coordination provides seamless integration with ATLAS world-wide computing operations
  - We have a very effective Integration Program in place to ensure our share to production & analysis operations with real data and simulation and to integrate new components and services
  - Excellent contribution of the Tier-1 and Tier-2 Sites to high volume production (event simulation, reprocessing) and analysis
- The U.S. ATLAS Computing Facilities need sufficient funding to be on track to meet the ATLAS performance and capacity requirements
  - Tier-2 funding uncertainties beyond 2011
    - New Cooperative Agreement w/ NSF for 2012 – 2016 – awaiting the numbers
  - Tier-1 equipment target reduced to minimum
    - Extended equipment lifetime increases risk
- Smooth integration of the Tier 3s in the U.S. – thanks to Doug
  - Tier-3s largely in place & operational
  - Doug has moved on to Analysis support
- U.S. ATLAS is actively pursuing continuation of OSG

➢ Overall, the Facilities in the U.S. have performed very well during the 2011 run, and I have no doubts that this will hold for 2012 and beyond !

In Loving Memory of...

TEVATRON

1983-2011

"...new knowledge has all to do with honor and country but it has nothing to do directly with defending our country except to help make it worth defending."

- Robert R. Wilson