



ATLAS Computing

J. Shank

Future computing in particle physics. Edinburgh 15 June, 2011




Overview

- Current Computing Operations
 - Data Throughput, Production and Analysis jobs
- Resource estimates in the next few years
- Software development issues in the future



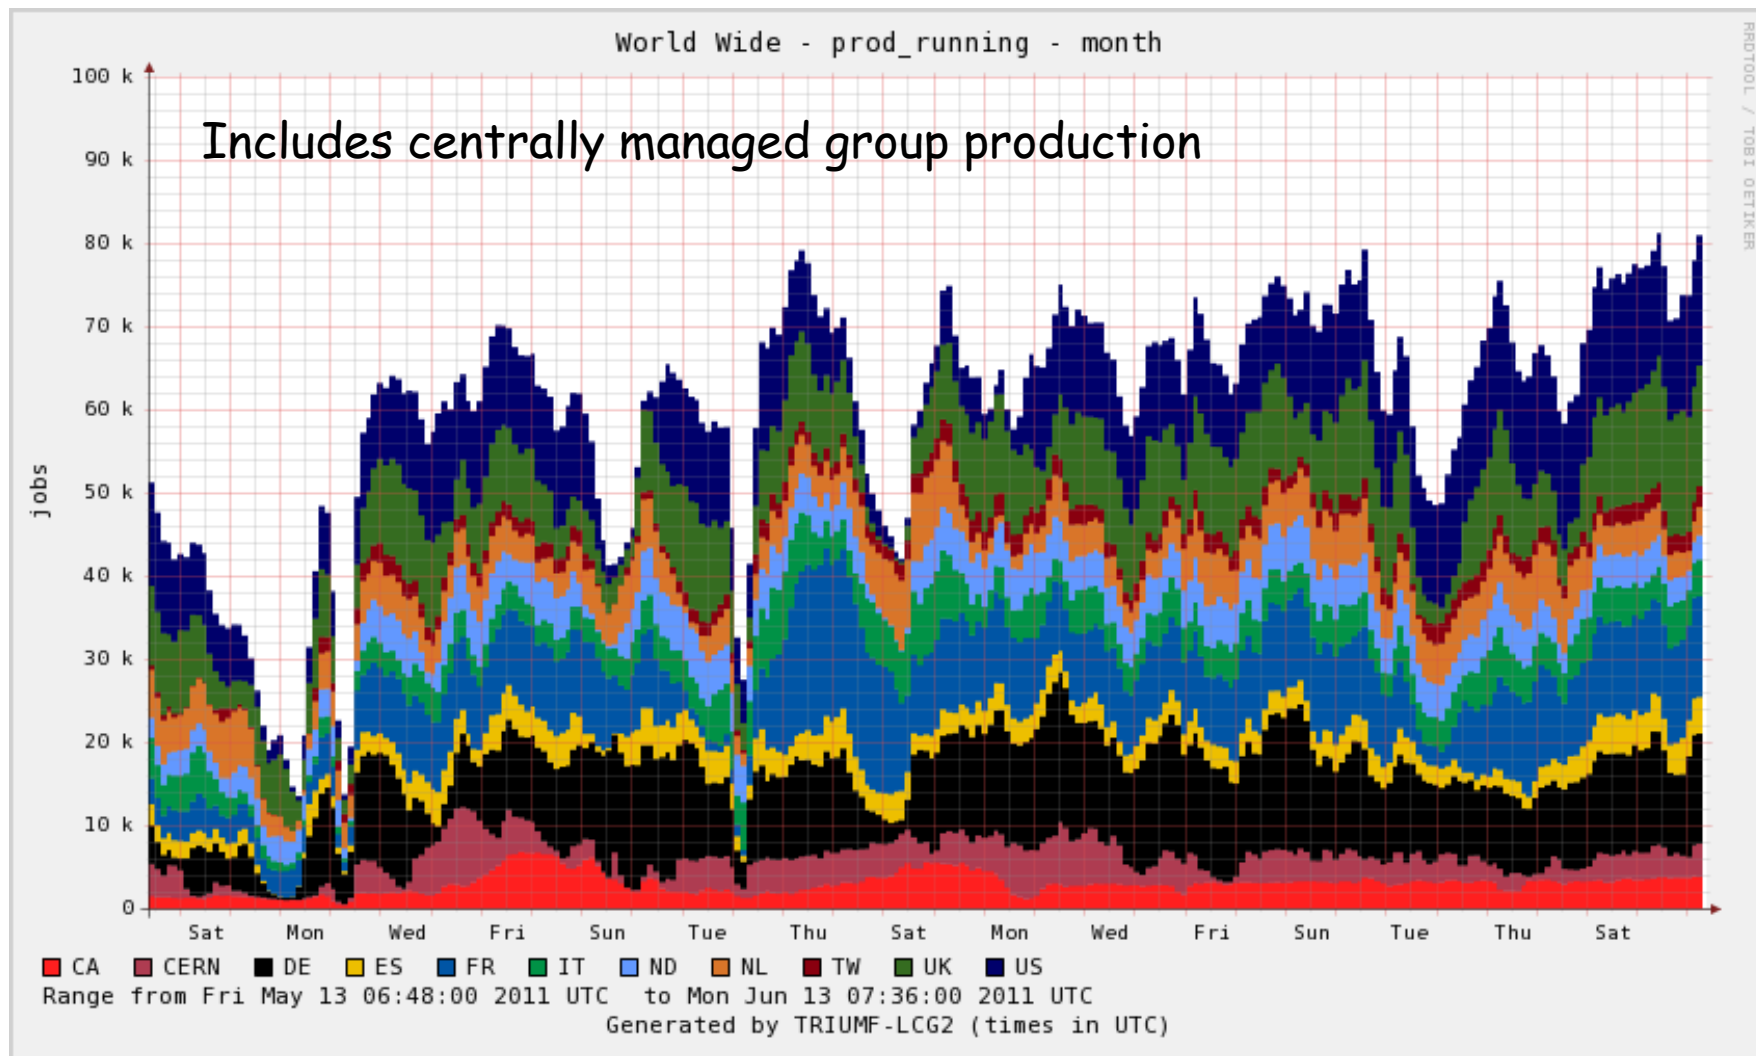
ATLAS Computing: Data Throughput

Daily average for last 30 days.
Throughput (MB/sec) from CERN to T1 centers.






ATLAS Computing: Production (Simulation and data processing)





ATLAS Computing: Analysis (Group and User)





Rolled up Summary of Resource Request

CPU [kHS06]	2010	2011	2012	2013
CERN	74	74	73	60
Tier-1	178	202	259	280
Tier-2	226	275	295	321
Disk [PB]				
CERN	4.7	7	9	10
Tier-1	22	22	27	30
Tier-2	24	35	49	56
Tape [PB]				
CERN	9	14	18	18
Tier-1	18	28	36	40



Some Current Issues

- Reconstruction Software
 - Dealing with Pileup
 - Memory footprint
 - Dealing with Multicores
 - See talk of D. Rousseau
- Data Placement and Management
 - PanDA Distributed Data Placement (PD2P)
 - Distributed Data Management (DDM)
 - Future of CASTOR?, xrootd?
- Physics Analysis
 - See talk of A. Farbin

- DDM architecture needs to go through a major revision
 - Robust and scalable, but more and more difficult to introduce new functionalities
 - Re-engineering would benefit from new technologies and solutions
- DDM team is working on a new system (Rucio)
 - Not necessarily backward compatible with the existing one
 - In design phase, architecture and timescale by July SW and Computing week
- At the same time the current system should be supported
 - Maintained but also improved if needed

- The current machinery of job definition needs to be re-factored
 - More and more difficult to introduce new functionalities
 - Need tighter integration with DDM
 - Manpower has been allocated, work will start now
 - Timescale October SW week for beta version
 - As for DDM, the current system needs to be supported and in this case the change should be adiabatic
- Many improvements foreseen in Panda
 - Within the year
 - Integration with DDM file level callbacks
 - Possible integration of GlideinWMS
 - Optimize usage of AthenaMP/allnode
 - Explore leverage and commonality with Rucio

- Streamlining
 - Backend has been consolidated (Panda)
 - Two frontends (pAthena/Ganga) currently supported maximizing shared code
- Improve the “user experience”
 - Automatic resubmission of failed jobs
 - More tools to favor usage of the Grid (pROOT?)
 - Timescale of some weeks
- Introduce “analysis tasks” and task management
 - Production tasks obviously already exist
 - Timescale by Sept SW week

- Identified as one of the critical areas last year
 - Therefore need more resources and attention
- Converged in an “unofficial” cocktail of components
 - Backend data exposed via HTTP in JSON format
 - JQuery+plugins for visualization
 - Not mandatory, but a standard de facto
- Converged in a common set of requirements
 - Filtering, sorting, grouping
- Main goal is to provide the needed information in a correct way (no inconsistencies)
 - But quality of visualization is now also pursued

- Medium to Long term perspective
 - Explore new solutions and products
- Need to keep the effort focused
 - Fixed targets and objectives
 - Regular updates
- from the real beginning, involve and collaborate with other parties than ADC
 - ATLAS software community
 - Other experiments
 - WLCG, OSG, CERN IT and PH-SFT

- Cloud Computing and Virtualization
 - Successful workshop in April
 - Test plan defined
 - Hammercloud using CERNVM on LXCLOUD at CERN (results in 6 weeks)
 - Other cloud installations will follow
 - Need to think at the data layer (will start with simple local storage approach)
- File & event level caching
 - Test Plan defined in Dubna
 - remote IO between CERN and BNL
 - Using again Hammercloud
 - Thinking about Brokering for reuse - strategy and tools

- NoSQL databases
 - Two testbeds in place at BNL and CERN
 - First use cases identified:
 - DDM trace monitoring (in place)
 - DDM accounting, Panda Monitoring archive
- Network Monitoring
 - Build a network monitoring infrastructure
 - perfSONAR infrastructure in place in the US.
 - ping and throughput test
 - Extend it to other clouds
 - Need development of components for data aggregation and monitoring

- Multicores
 - ATLAS Software community working on AthenaMP
 - ADC working on the Gridification on multicore jobs
 - How to initiate them at the site
 - Scheduling in Panda
 - ...