ATLAS Production on OSG

Kaushik De

D0SAR Meeting, Iowa State U.

April 5, 2007

ATLAS MC Production

- Computer System Commissioning (CSC)
  - Software integration and operations exercise
  - Started more than one year ago (since end of 2005)
  - Distributed (MC) production on massive scale
    - Validation of Athena software using ‘standard’ samples
    - Continuous production of high statistics physics samples
    - To exercise widely distributed computing infrastructure
- Three different grids used
  - OSG (U.S.), EGEE (the rest) and NDGF (Nordic countries)

CSC Production Statistics

- Many hundreds of physics processes have been simulated
- Release 11 (CSC11) and Release 12 (CSC12) production
- Dozens of sub-releases have been tested and validated
- Thousands of 'bug reports' fed back to software and physics
- 40M+ events done from CSC12
- ~200 TB of CSC12 data on disk
- 12.0.6 recon started (all 50M events for CSC notes)
- Huge team effort!

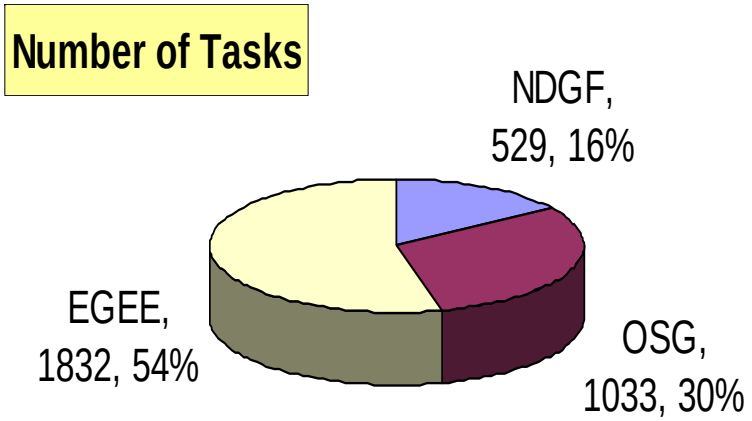

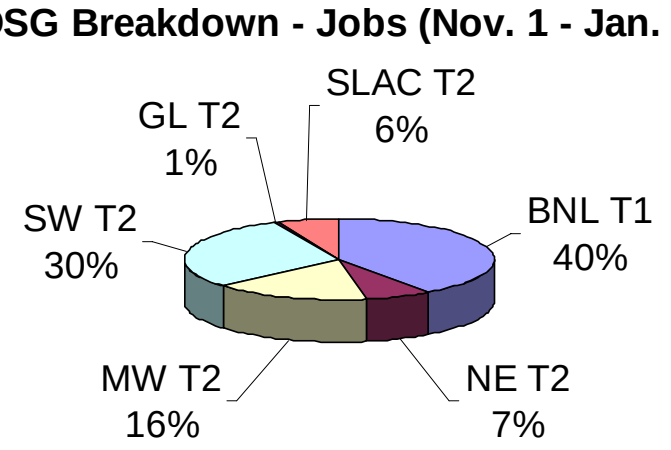

MC Sample Terminology

- Sample A (validation sample)
  - ~10 physics channels, 100-200k events
  - Run for every sub-release ~ every 2 weeks
- Sample B
  - ~25 physics channels, 1M events
  - Run for every stable release, validate for production
- Sample C
  - >100 physics channels, 10-20M events
  - Production for physics analysis
- We have run this sequence successfully many times on the grid for Releases 10, 11 and 12, over the past 15 months

Panda Overview

- PanDA (Production and Distributed Analysis) project was launched in the U.S. in August 2005
- Panda fully functional in very short time
- Panda is fully integrated with DQ2, ATLAS Distributed Data Management system
- All CSC production in the U.S. is done with Panda
- Growing user base for Distributed Analysis
- Panda serves all U.S. ATLAS Tier 1 and Tier 2 facilities

Key Panda Features

- **Service model** – Panda runs as an integrated service for all ATLAS sites (currently U.S.) handling all grid jobs (production and analysis)
- **Task Queue** – provides batch-like queue for distributed grid resources (unified monitoring interface for production managers and all grid users)
- **Strong data management** (lesson from DC2) – pre-stage, track and manage every file on grid asynchronously, consistent with DQ2 design
- **Block data movement** – pre-staging of output files is done by optimized DQ2 service based on datasets, reducing latency for distributed analysis (jobs follow the data)
- **Pilot jobs** – are prescheduled to batch systems and grid sites; actual ATLAS job (payload) is scheduled when CPU becomes available, leading to low latency for analysis tasks
- **Support all job sources** – managed or regional production (ATLAS ProdSys), user production (tasks, DIAL, Root, pAthena, scripts or transformations, GANGA…)
- **Support any site** – minimal site requirement: pilot jobs (locally or through grid), outbound http, and integration with DQ2 services

OSG Production Statistics

- OSG has successfully completed ~25M fully simulated physics events (simul+digit step) – 29% of total central production for CSC
- Also successfully completed >14M single particle events
- Since November, all available CPU's occupied (ran out of jobs only for few days, plus few days of service outages)
- About 380 TB of original data stored at BNL T1 (includes data generated on other grids)
- Additional ~100 TB of replicas kept at ATLAS Tier 2 sites on OSG

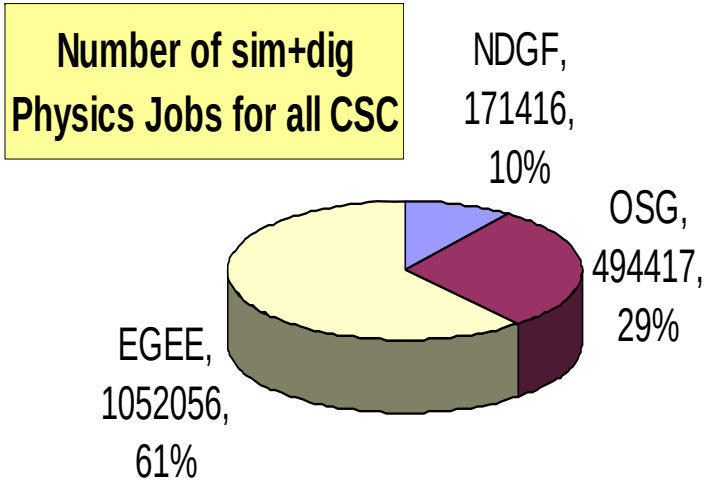

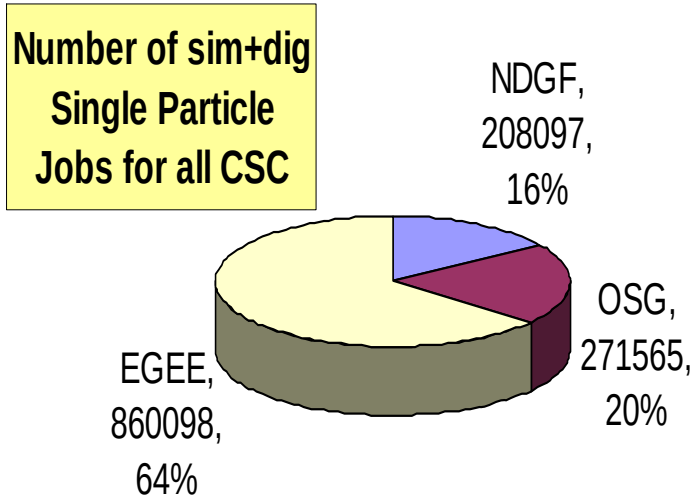

Walltime Usage - Successful Jobs

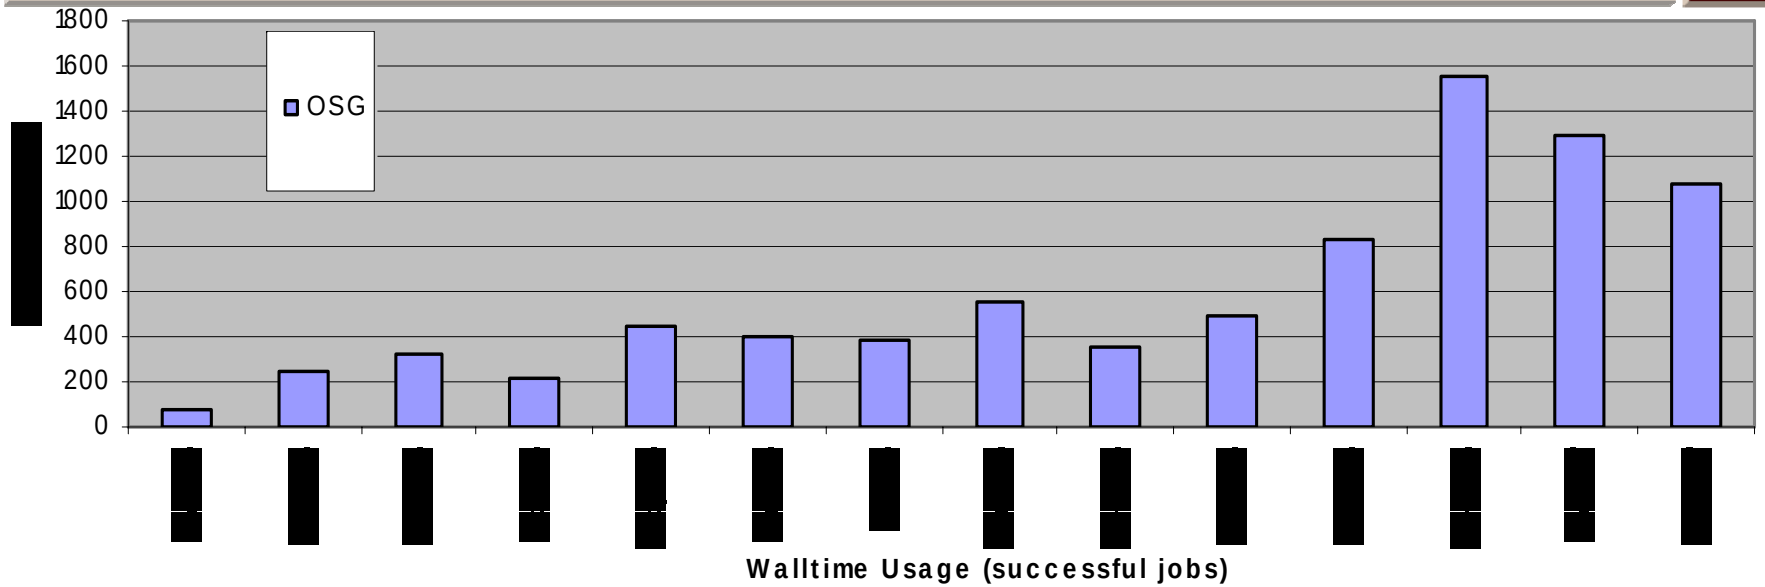

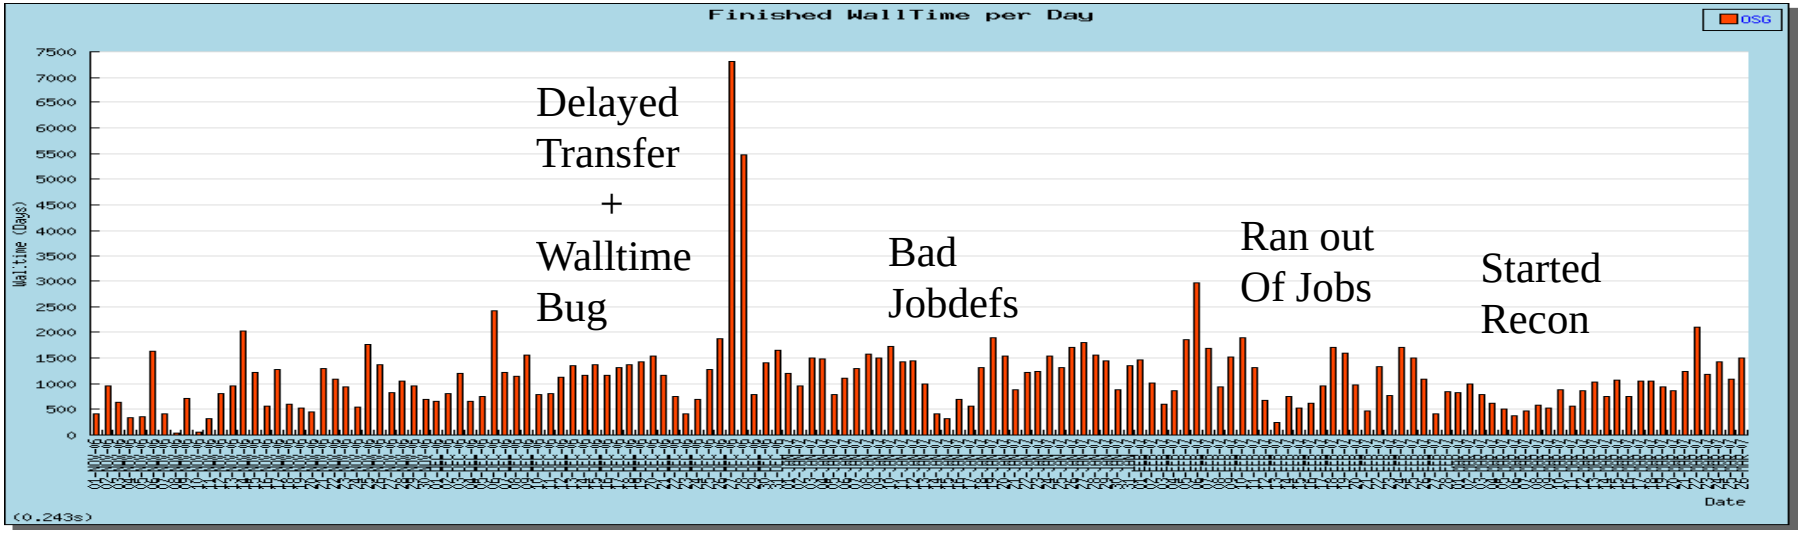

Job Error Rates – Past 2 Months

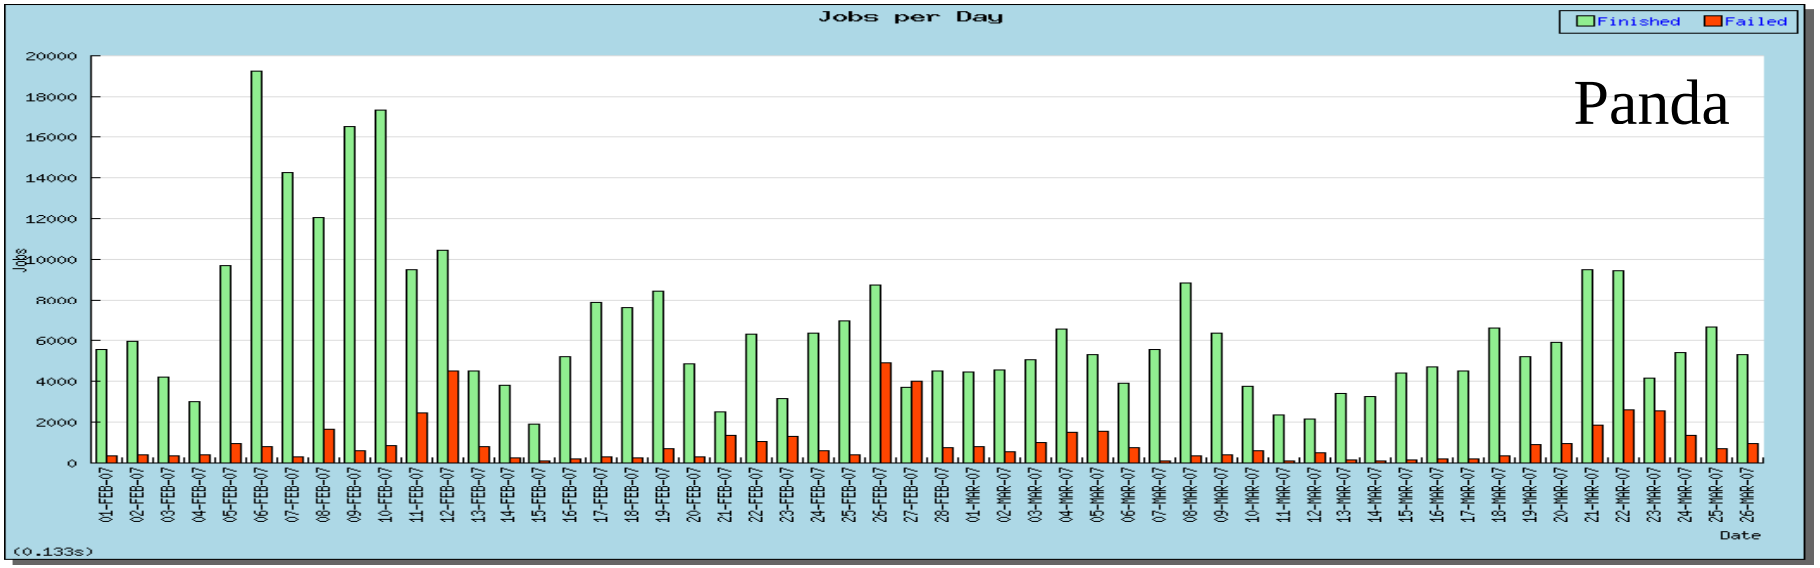

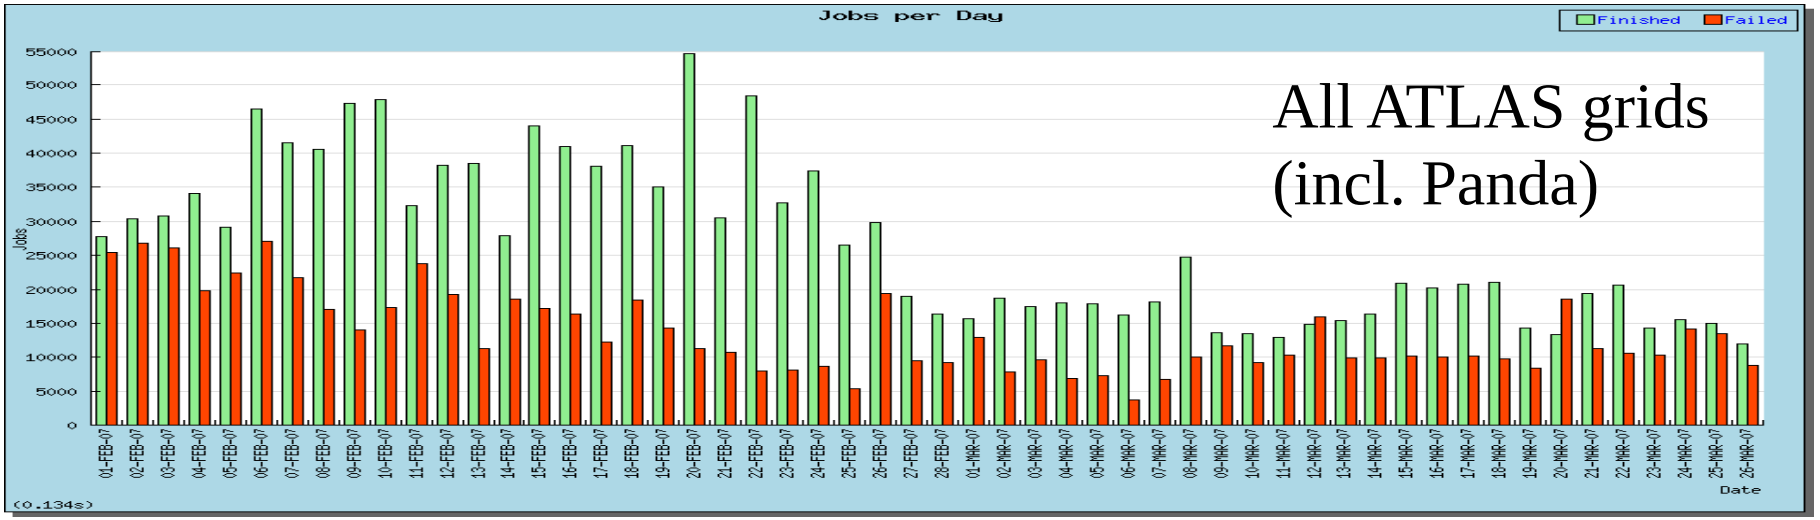

Production job summary, last 12 hours (Details: errors, nodes)

Site	Nodes	Jobs	Latest	defined	assigned	waiting	activated	running	holding	transferring	finished	failed	tot	trf	other
All	849	13337	03-27 06:47	130	311	345	2834	1607	2498	2258	2398	956	29%	0%	28%
BNL_ATLAS_1	182	1989	03-27 06:47	0	0	0	810	475	335	0	350	19	5%	0%	5%
BNL_ATLAS_2	0	0		0	0	0	0	0	0	0	0	0			
BU_ATLAS_Tier2	22	354	03-27 06:47	0	21	0	88	87	29	97	31	1	3%	0%	3%
BU_ATLAS_Tier2o	13	591	03-27 06:45	0	67	0	0	43	30	314	111	26	19%	7%	12%
IU_ATLAS_Tier2	38	381	03-27 06:47	0	1	0	183	64	2	101	30	0	0%	0%	0%
MWT2_IU	41	899	03-27 06:47	0	0	0	195	119	76	165	313	31	9%	0%	9%
MWT2_UC	34	754	03-27 06:47	0	14	0	117	135	58	176	251	3	1%	0%	1%
Unassigned	1	476	03-27 06:43	130	0	345	0	0	0	0	0	1	100%	0%	100%
OU_OCHEP_SWT2	41	607	03-27 06:47	0	0	0	232	82	1	112	176	4	2%	0%	2%
OU_PAUL_TEST	3	4	03-26 17:38	0	0	0	2	0	2	0	0	0			
PROD_SLAC	130	1119	03-27 06:47	0	0	0	217	221	58	368	248	7	3%	0%	3%
SLAC_PAUL_TEST	18	33	03-27 05:16	0	0	0	0	0	0	25	0	8	100%	0%	100%
UBC	0	0		0	0	0	0	0	0	0	0	0			
UC_ATLAS_MWT2	45	893	03-27 06:21	0	162	0	0	1	486	103	133	8	6%	0%	6%
UC_Teraport	4	52	03-26 20:39	0	46	0	0	0	0	6	0	0			
UMATLAS	59	2331	03-27 06:19	0	0	0	0	0	1380	1	162	788	83%	0%	83%
UTA-DPCC	37	529	03-27 06:47	0	0	0	168	70	11	180	92	8	8%	0%	8%
UTA_SWT2	157	2325	03-27 06:47	0	0	0	822	310	30	610	501	52	9%	0%	9%

Some Thoughts For D0SAR

- So far, CSC production (managed) only at T1 & T2 sites
- MC production scaling well – getting ready for data
- Now CSC entering new phase – distributed analysis
- Opportunity for all D0SAR sites/members to participate
- Distributed physics analysis model needs to be exercised
- Need many active users to test model at scale
- Also, develop model for Tier 3 role
- Please let us know what comes out of discussions in this meeting