



ATLAS

Grid Computing



ICFA Workshop on HEP
Networking, Grid, and Digital
Divide Issues for Global e-Science

THE CENTER FOR HIGH ENERGY PHYSICS
KYUNGPOOK NATIONAL UNIVERSITY

May 23, 2005


Rob Gardner
University of Chicago

Outline

- ATLAS project
- Computing model
- Data Challenges
- Case in point: focus on production on Grid3
- Outlook



Scale of the detector



<i>Diameter</i>	<i>25 m</i>
<i>Barrel toroid length</i>	<i>26 m</i>
<i>End-cap end-wall chamber span</i>	<i>46 m</i>
<i>Overall weight</i>	<i>7000 Tons</i>



Recent Cavern Photo



5/23/2005

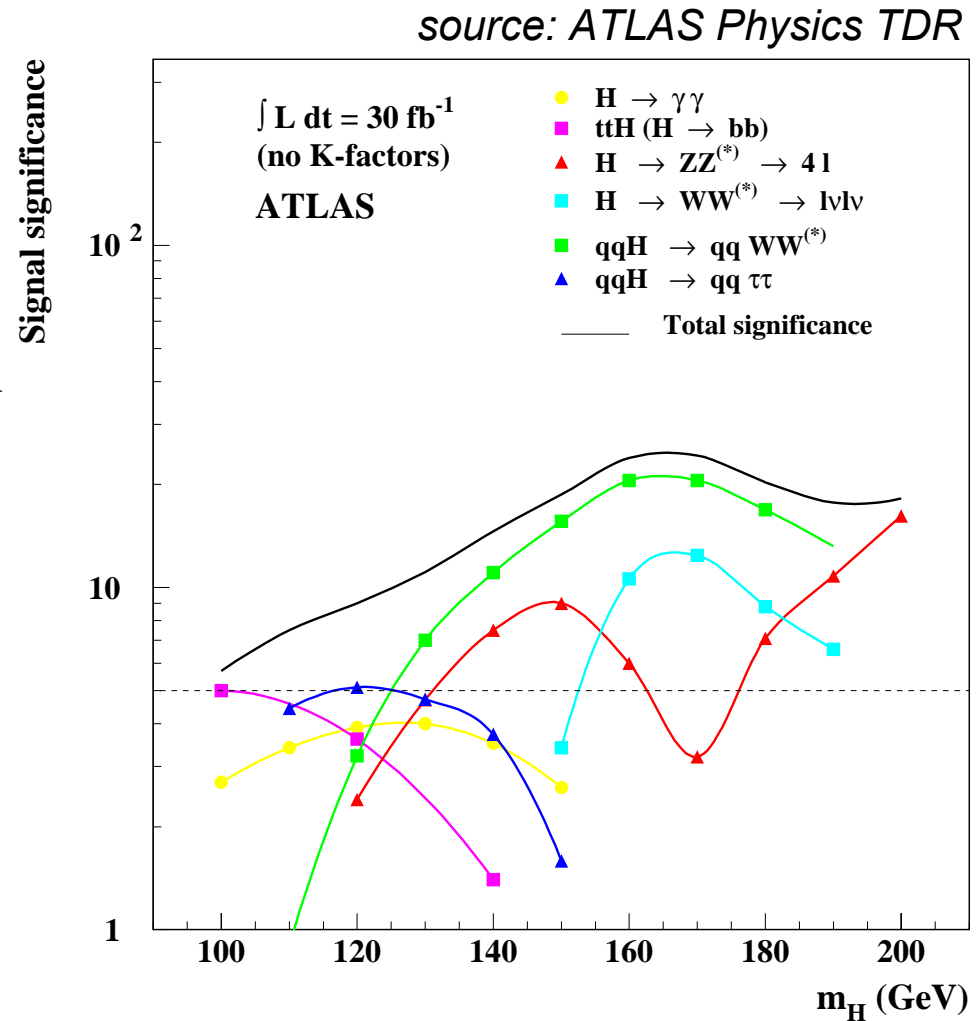
R. Gardner

ATLAS Grid Computing

4

Higgs Discovery Potential

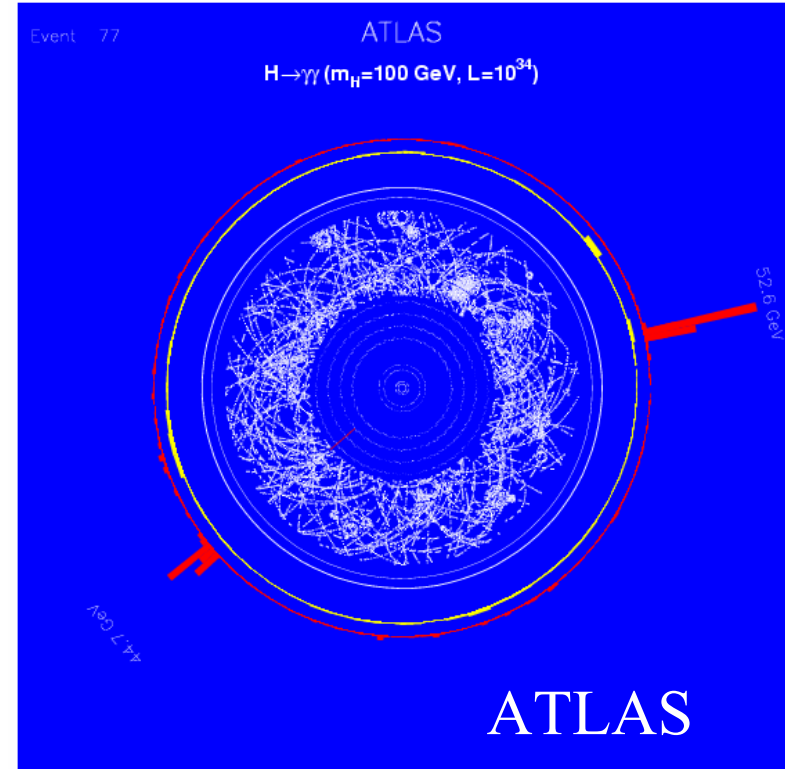

If we can start up at $1/10^{\text{th}}$ design luminosity, we'll discover a Higgs with mass greater than 130 GeV within 1 year. Will cover entire theoretically allowed range with 1 year of design luminosity.



Higgs Signatures in ATLAS

$$H \rightarrow ZZ^* \rightarrow e^+e^-\mu^+\mu^-$$

$$H \rightarrow \gamma\gamma$$

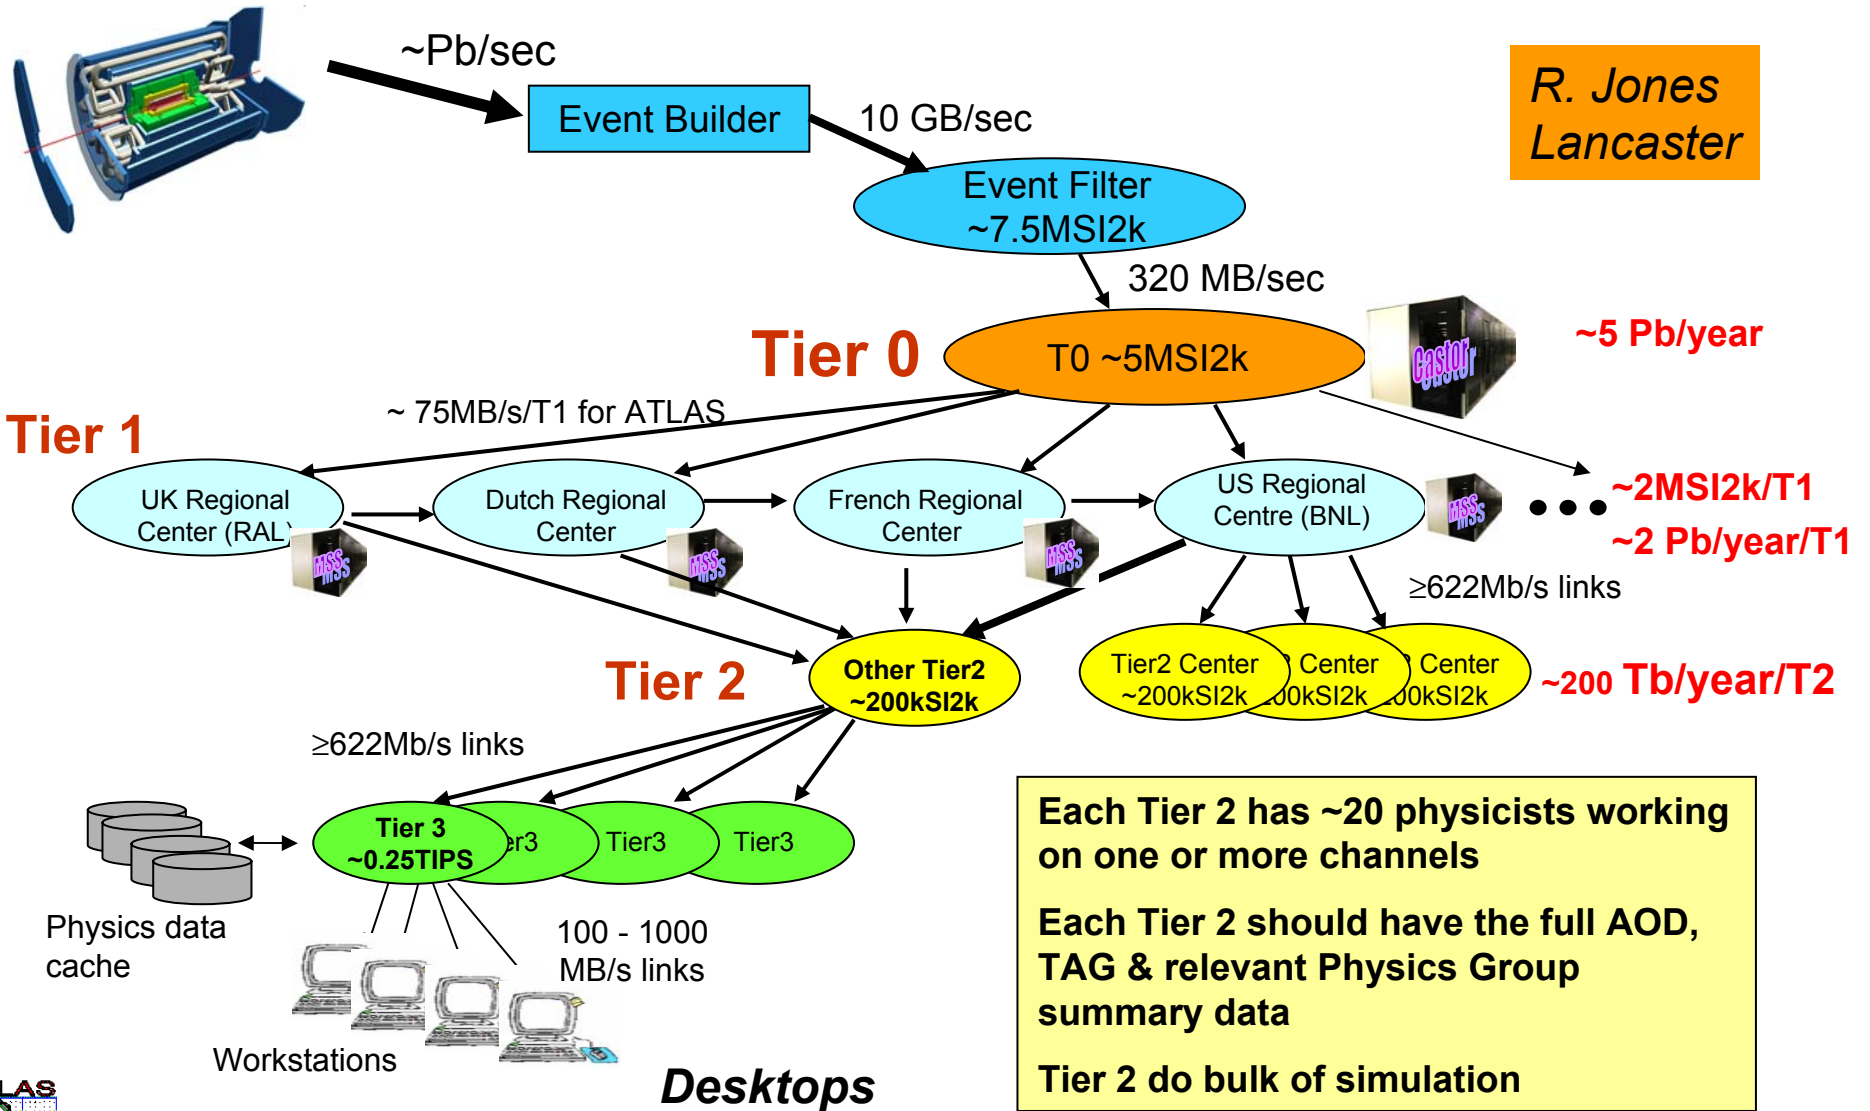


ATLAS Computing Model

- Developed to meet physics processing requirements for all phases – event filtering, reconstruction, re-processing, Monte Carlo production, and analysis of highly summarized data sets.
- Consists of a hierarchy of distributed processing centers – the “Tier0 – Tier1 – Tier2 – Tier3” model
- The model has been refined a number of times, and still is understudy as we get further along.



LHC Tiered Computing Model for ATLAS



Computing Model Assumptions

- The ATLAS Computing model has been designed to meet the following assumptions
 - 100 days running in 2007 (5×10^6 sec live)
 - 200 days running in 2008 and 2009 at 50% efficiency (10^7 sec live)
 - Luminosity:
 - $0.5 \times 10^{33} \text{ cm}^{-2}\text{s}^{-1}$ in 2007
 - $2 \times 10^{33} \text{ cm}^{-2}\text{s}^{-1}$ in 2008 and 2009
 - $10^{34} \text{ cm}^{-2}\text{s}^{-1}$ (design luminosity) from 2010 onwards



Data Sets Among Tiers

- Tier-0 has raw+calibration + first-pass ESD
- CERN Analysis Facility has AOD, ESD and RAW samples
- Tier-1s hold RAW data and derived samples and ‘shadow’ the ESD for another Tier-1
- Tier-1s also house simulated data
- Tier-1s provide reprocessing for their RAW and scheduled access to full ESD samples
- Tier-2s provide access to AOD and group Derived Physics Datasets



Processing Among Tiers

■ Tier0

- ❑ First pass processing on express/calibration lines
- ❑ 24-48 hours later, process full primary stream with reasonable calibrations

■ Tier1

- ❑ Reprocess 1-2 months after arrival with better calibrations
- ❑ Reprocess all resident RAW at year end with improved calibration and software

■ Tier2

- ❑ Full simulation load



ATLAS Data Challenge 2

- Phase I: MC production started July 2004
 - Event generation, GEANT4 simulation, digitization & pile-up
 - ran on three Grids, worldwide, finished in ~ 6 months
- Phase II: test of Tier-0 operation
 - 10 day exercise to test 1 day of real data-taking
 - input is “Raw Data” like
 - output (ESD+AOD) will be distributed to Tier-1’s in real time for analysis
- DC2 Scale
 - Physics channels ~10 million events
 - Calibration samples ~few million events



DC2 Phase I

■ LCG Grid

- Uses RB and LCG Grid-Canada interface

■ NorduGrid

- Scandinavian resources + sites in Australia, Germany, Slovenia, Switzerland

■ Grid3

- Dedicated and opportunistic resources

ATLAS DC2 Production Status

Overview of Grids

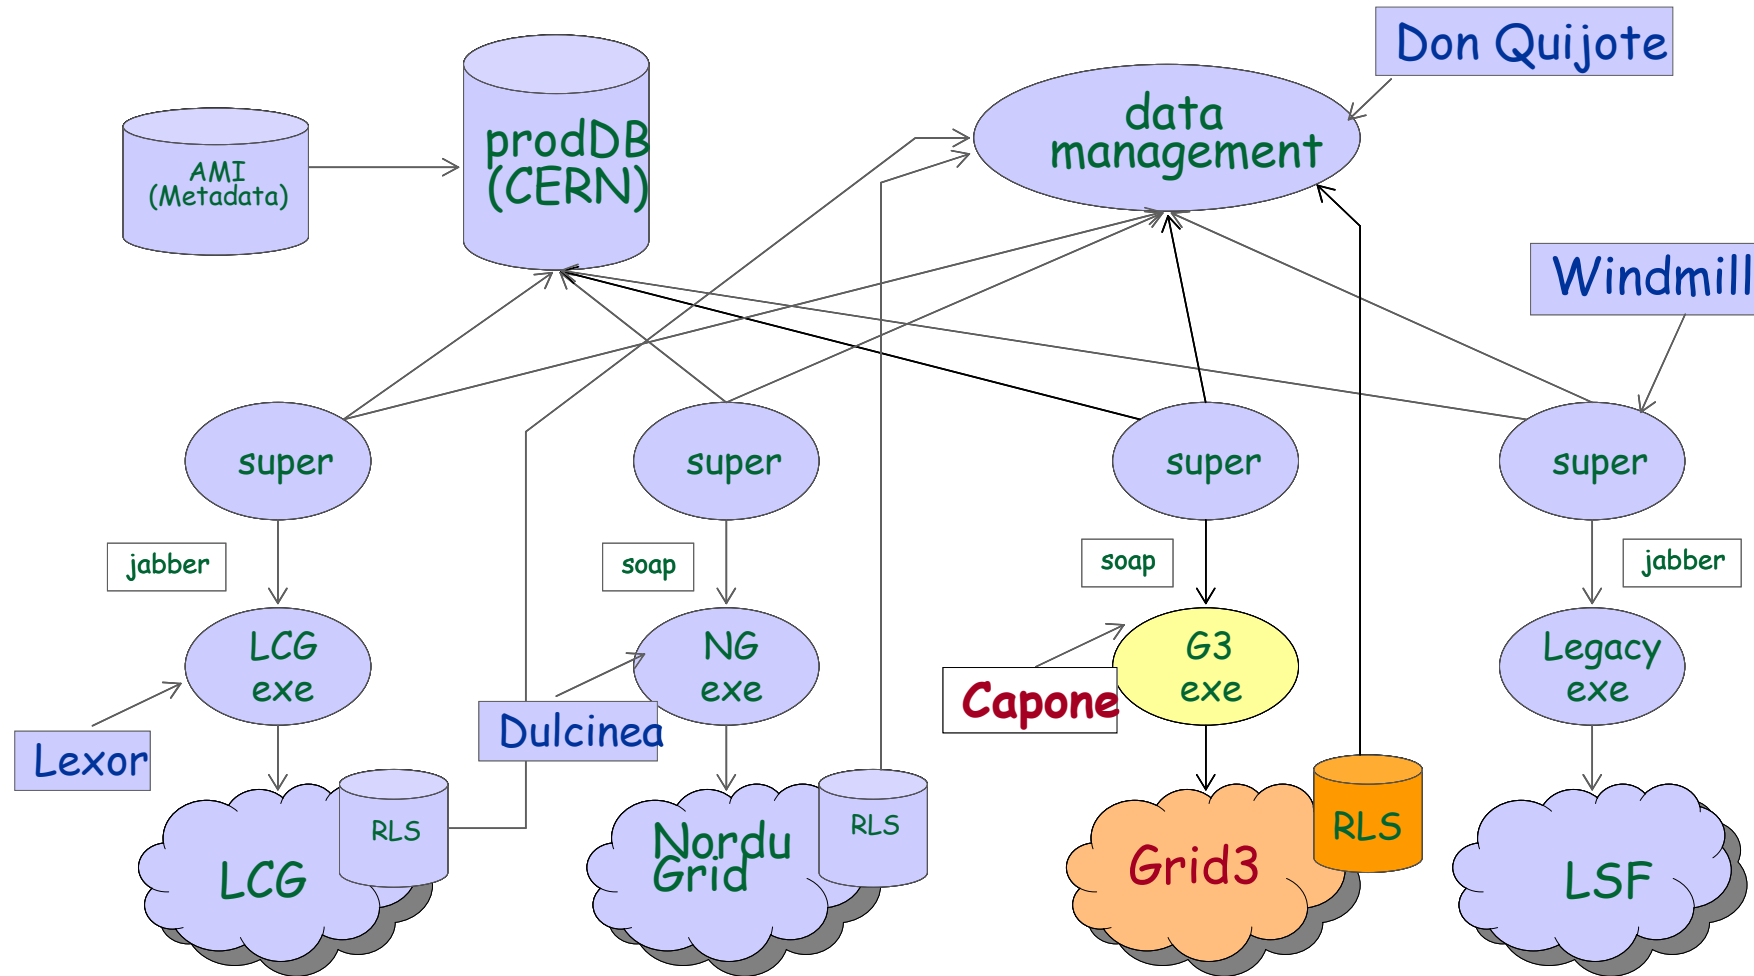
as of 2005-01-28 10:33:00

Grid	submitted	pending	running	finished	failed	efficiency
Grid3	27	10	215	122531	41385	75 %
NorduGrid	16			111108	66311	63 %
LCG	32	380	432	122688	215105	36 %
TOTAL	75	390	647	356327	322801	52 %

[Prodsys homepage](#)

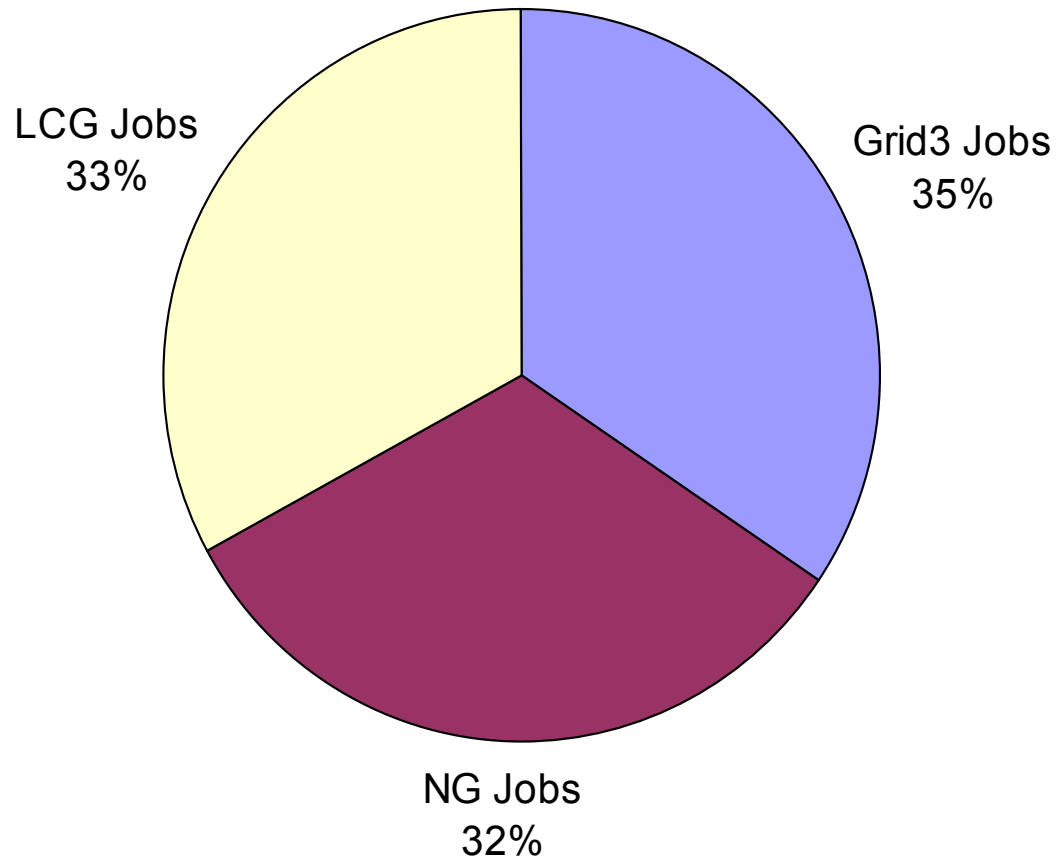


ATLAS Production System





DC2 Jobs Per Grid

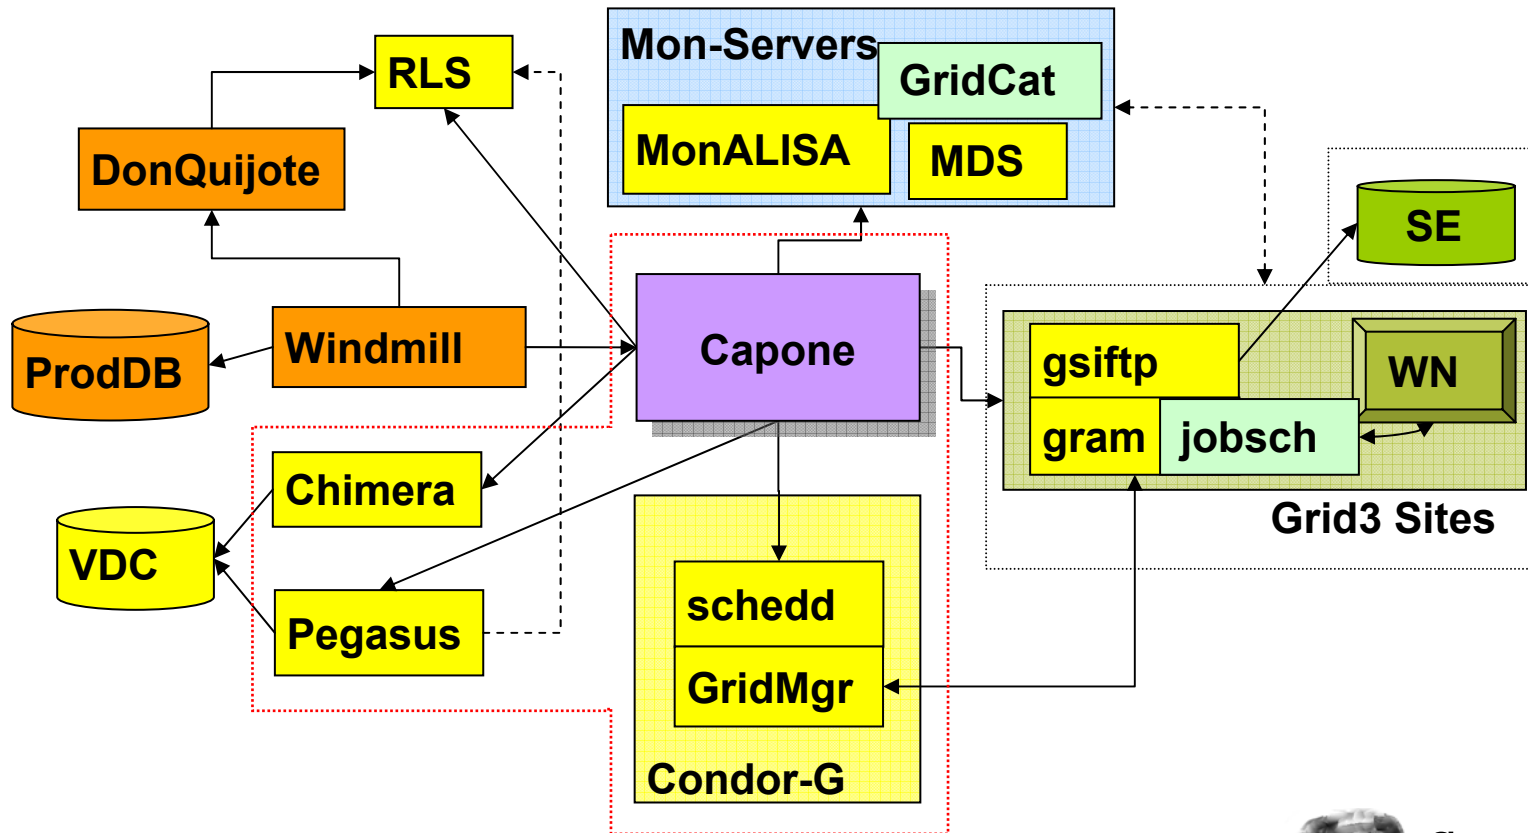
DC2 Jobs (till Jan 22, 2005)



DC2 Events Processed (till Jan 22, 2005)



Focus: US ATLAS Grid3 Production System



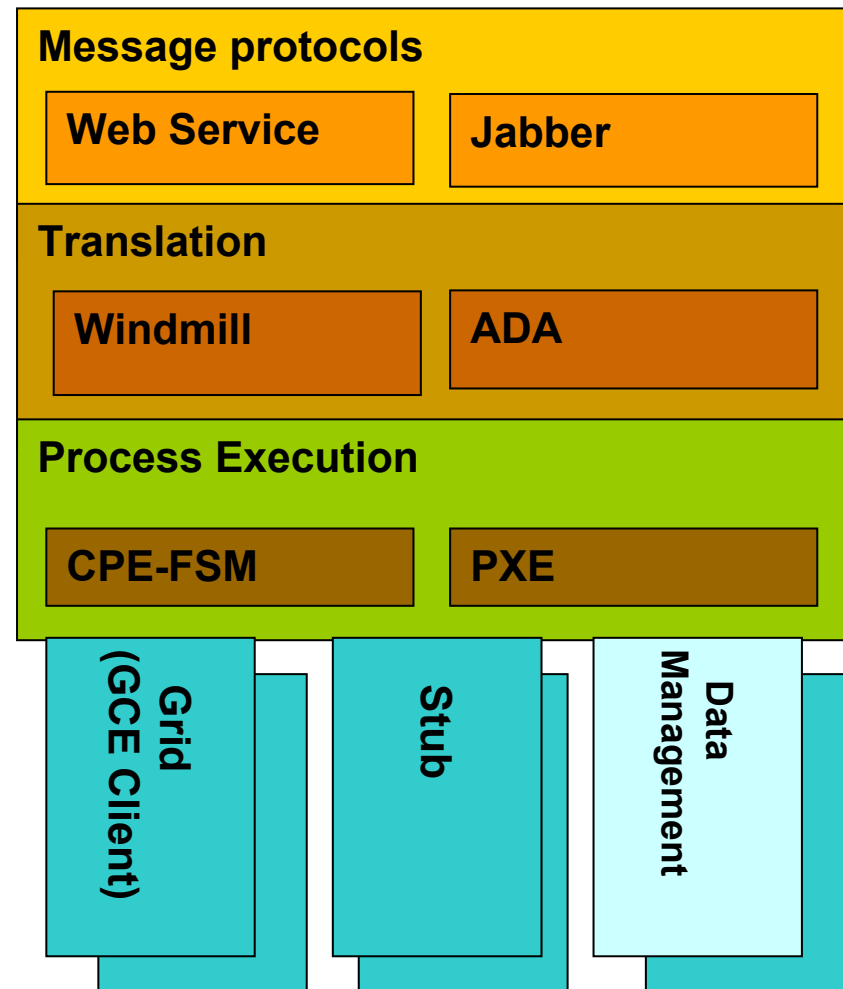
Capone
inside



Capone Workload Manager



- Message interface
 - Web Service & Jabber
- Translation layer
 - Windmill (Prodsys) schema
- Process execution
 - Capone Finite State Machine
- Processes
 - Grid interface
 - Stub: local shell script testing
 - Data management



Capone + Grid3 Performance

ATLAS DC2 Overview of Grids as of 2005-02-24 18:11:30

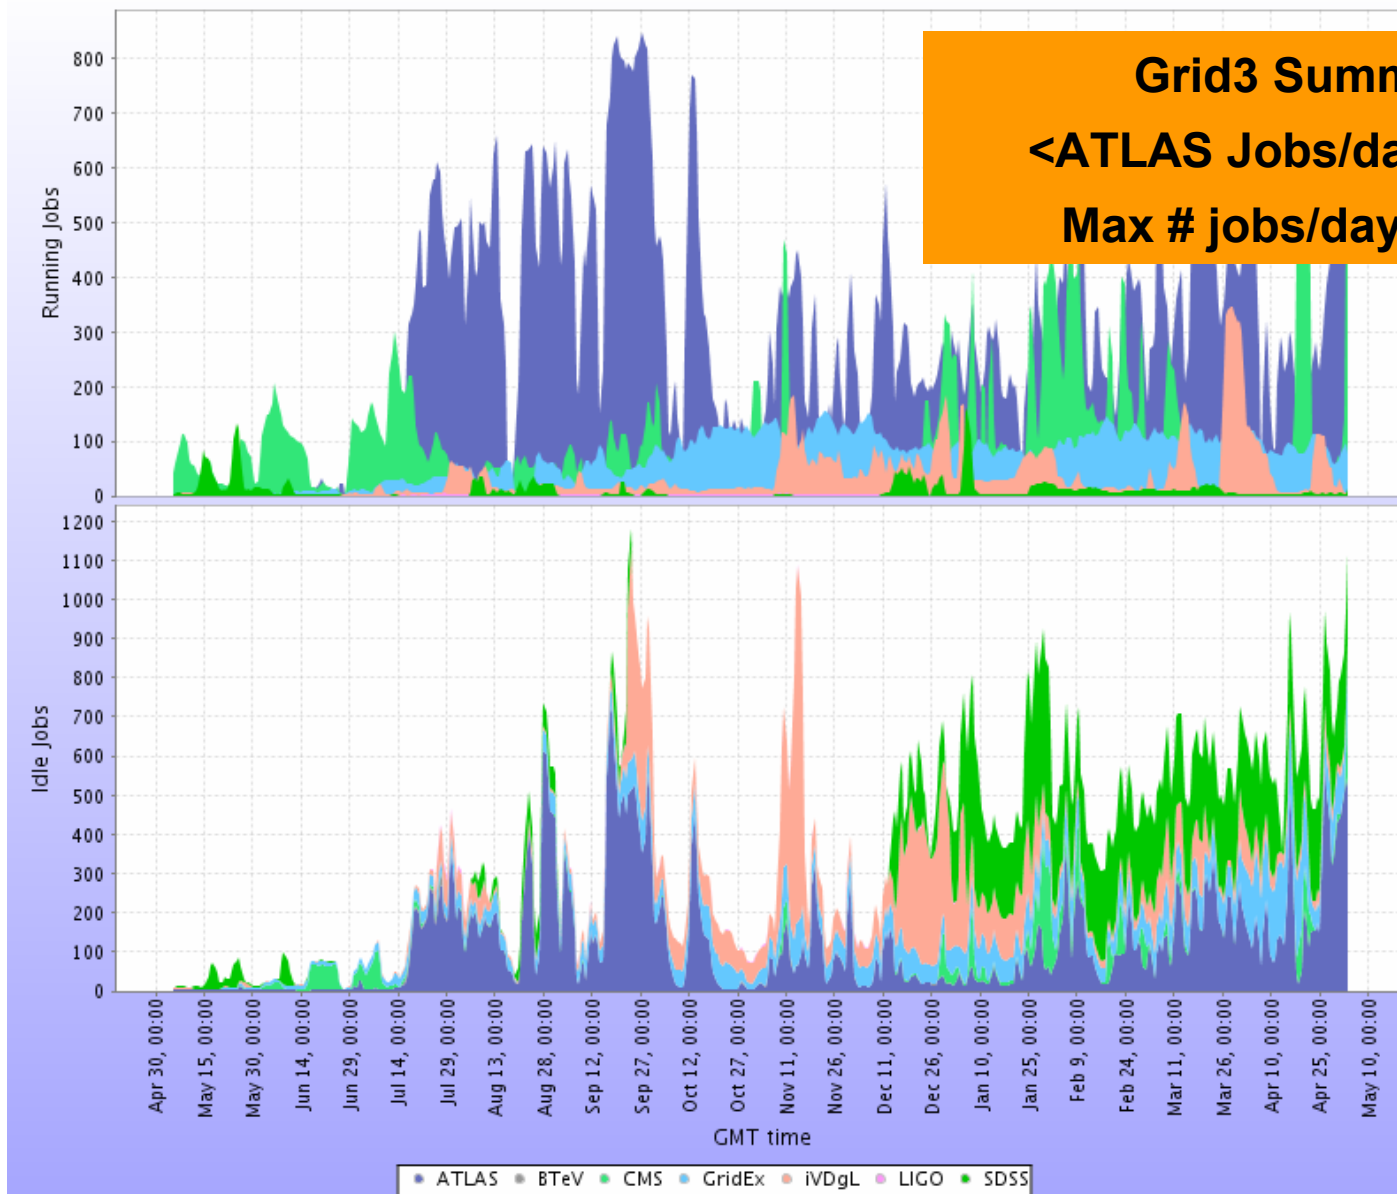


Grid	submitted	pending	running	finished	failed	efficiency
Grid3	36	3	814	153028	46943	77 %
NorduGrid	29	130	1105	114264	70349	62 %
LCG	60	528	610	145692	242247	38 %
TOTAL	125	661	2529	412984	359539	53 %

- Capone submitted & managed ATLAS jobs on Grid3 > **150K**
- In 2004, **1.2M CPU-hours**
- Grid3 sites with more than **1000** successful DC2 jobs: **20**
- Capone instances > 1000 jobs: **13**



Global jobs view



Grid3 Summary
<ATLAS Jobs/day> = 350
Max # jobs/day = 1020



Summary

- ATLAS is ramping up for turn-on of the LHC in 2007
- Major production exercises have provided many lessons and experience for moving forward
 - both at the (ATLAS-specific) application layer and in the reliability of the basic middleware services
- Major challenges to address are improvements in scalability of production services and fault tolerance (reducing failure rates)

