The Grid for Particle Physics

Glenn Patrick

IOP HEPP Annual Conference

Lancaster, April 2008

Science & Technology Facilities Council
Rutherford Appleton Laboratory

GridPP
UK Computing for Particle Physics

In the beginning…

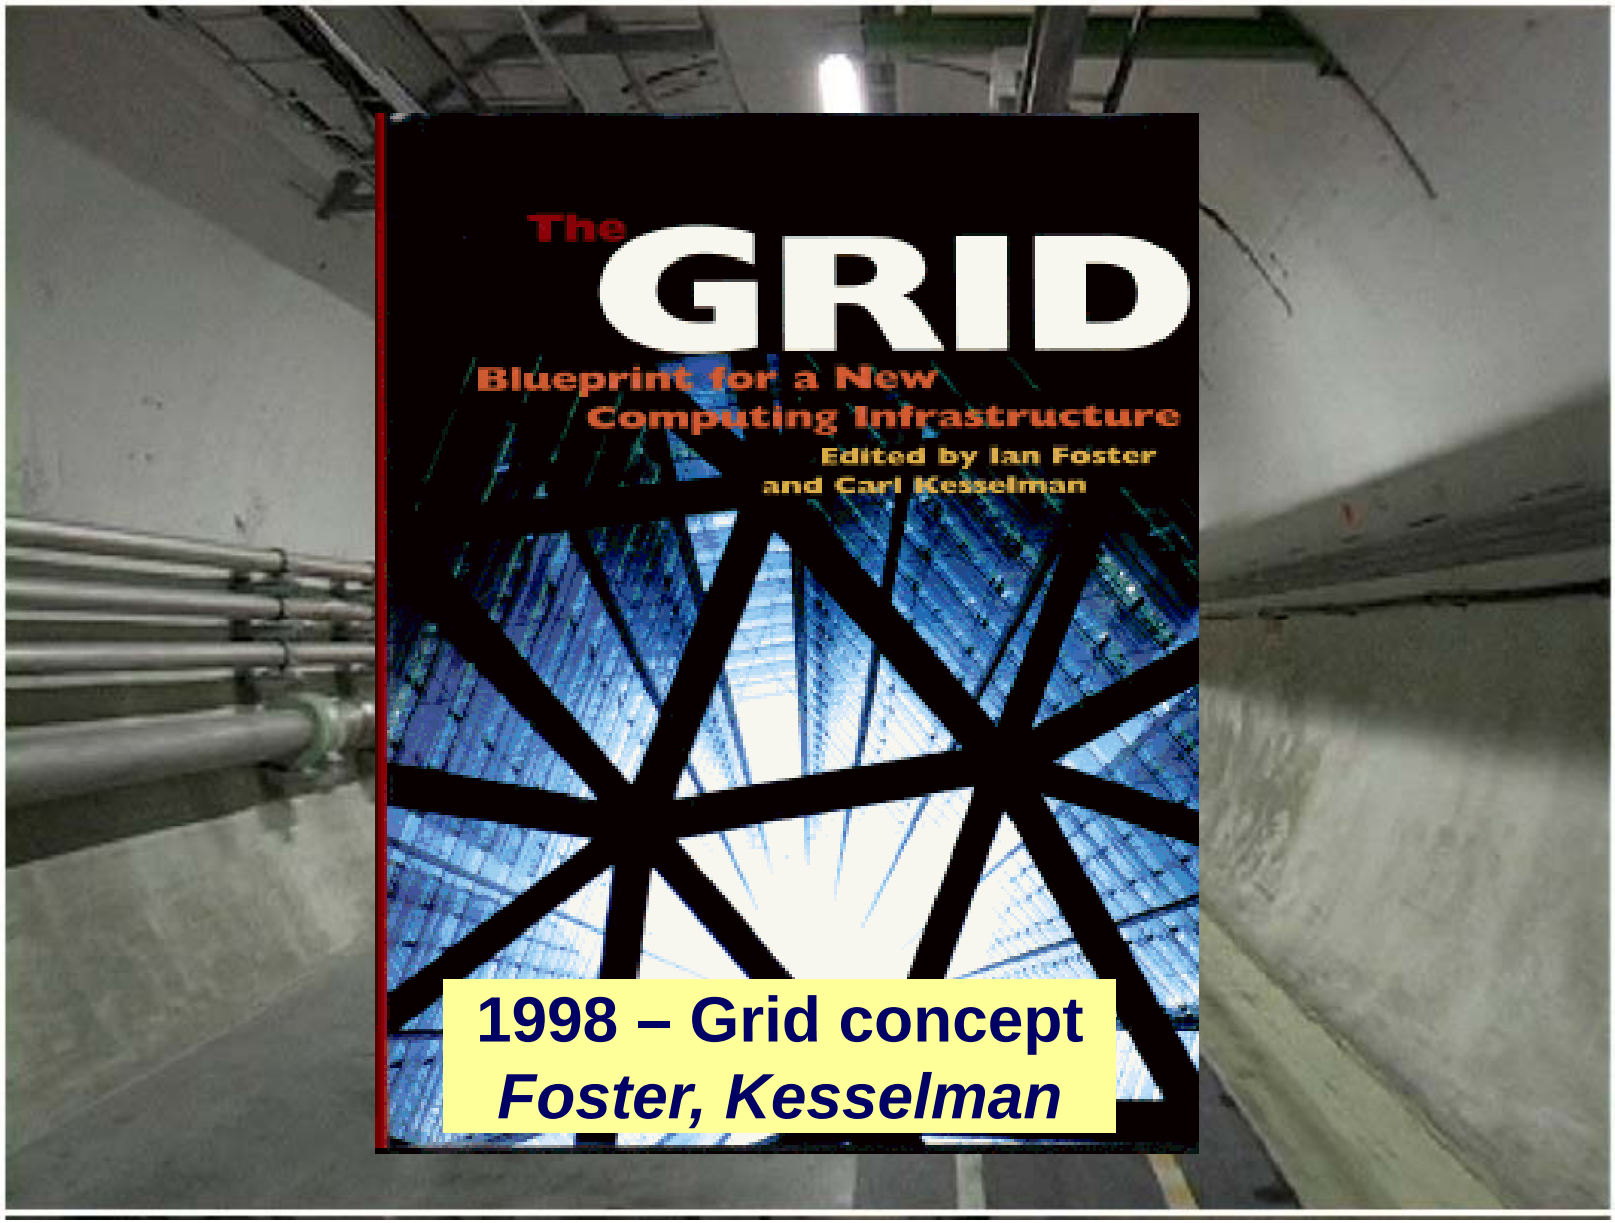

1998 – Grid concept
Foster, Kesselman

What is a Grid?

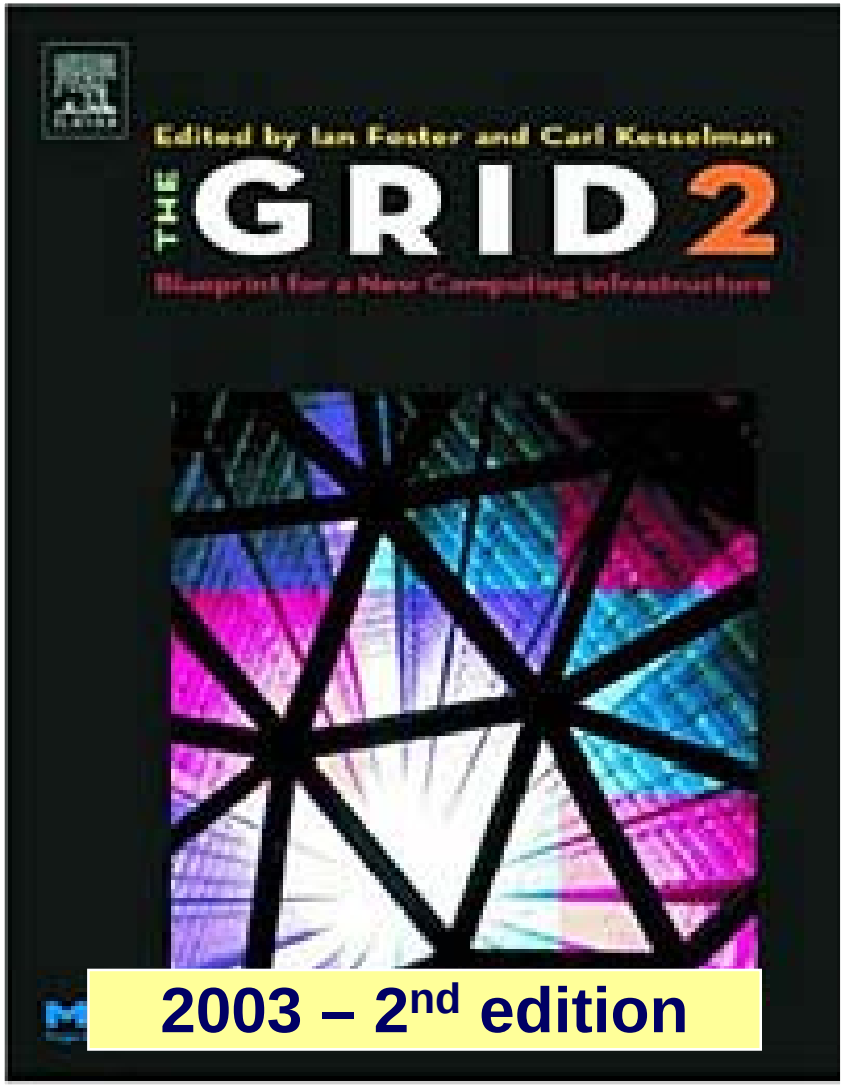

"A computational grid is a hardware and software infrastructure that provides dependable, consistent, pervasive, and inexpensive access to high-end computational capabilities".
Ian Foster & Carl Kesselman, 1998

Grid Checklist (2002):

1. Co-ordinates resources that are not subject to centralized control.
2. … using standard, open, general-purpose protocols and interfaces.
3. … to deliver non-trivial qualities of service.

Science & Technology Facilities Council

Something like this…

GridPP UK Computing for Particle Physics

… or this

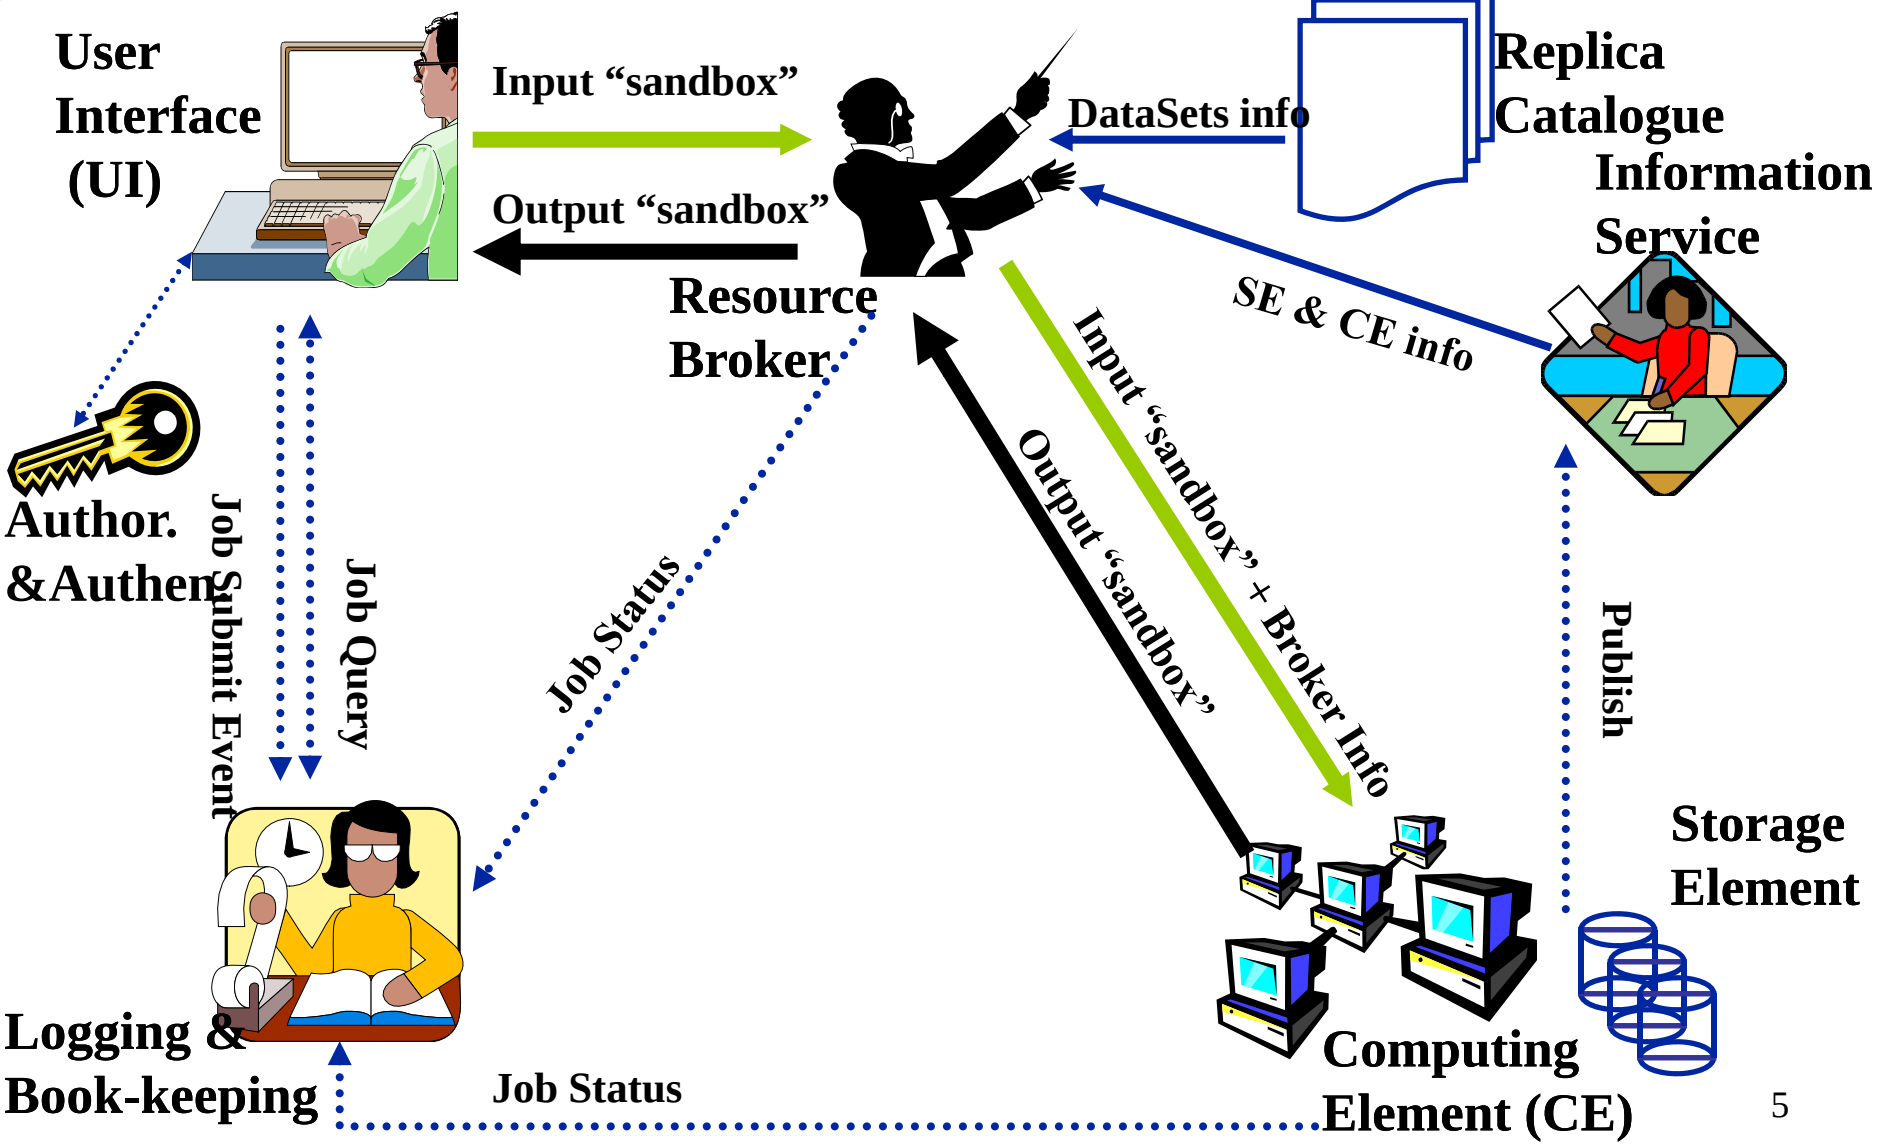
User Interface (UI)
Input "sandbox"
Output "sandbox"
Resource Broker
DataSets info
Replica Catalogue
Information Service
SE & CE info
Author. &Authen
Job Submit Event
Job Query
Job Status
Input "sandbox" + Broker Info
Output "sandbox"
Publish
Storage Element
Logging & Book-keeping
Job Status
Computing Element (CE)

Physical Tier Structure

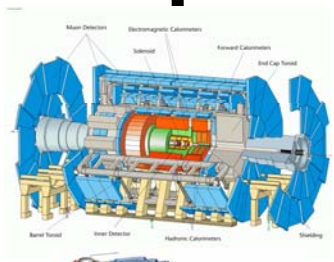

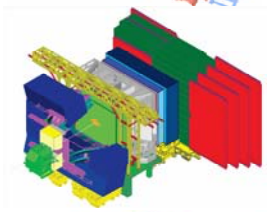

Offline farm

Online system

Tier 0

CERN computer centre

11 T1 centres

Tier 1

National centres

RAL,UK

Spain

ermany

Italy

rance

Tier 2

Regional groups

ScotGrid

NorthGrid

SouthGrid

London

Institutes

Glasgow

Edinburgh

Durham

Workstations

Structure chosen for particle physics.
Different for others.

Rise of the UK Grid

Particle Physics e-Science Programme Proposal

The UK Grid for Particle Physics Collaboration

GridPP

April 2001

University of Birmingham,
University of Bristol,
Brunel University,
CERN, European Organization for Nuclear Research,
University of Cambridge,
University of Durham,
University of Edinburgh,
University of Glasgow,
Imperial College of Science, Technology and Medicine,
Lancaster University,
University of Liverpool,
University of Manchester,
Oxford University,
Queen Mary, University of London,
Royal Holloway, University of London,
Rutherford Appleton Laboratory,
University of Sheffield,
University of Sussex,
University of Wales Swansea,
University College London.

Contacts
Dr. Tony Doyle – A.Doyle@physics.gla.ac.uk
Dr. Steve Lloyd – S.L.Lloyd@qmw.ac.uk

GridPP – UK Grid

UK's contribution to LHC computing:
- 19 UK Universities, STFC and CERN

GridPP1 (2001- 2004) £17M
"From Web to Grid"

GridPP2 (2004 - 2008) £16M
"From Prototype to Production"

GridPP3 (2008 – 2011) ~£30M
"From Production to Exploitation"

CERN
UNIVERSITY OF CAMBRIDGE
THE UNIVERSITY OF BIRMINGHAM
University of Bristol
BRUNEL UNIVERSITY WEST LONDON
University of Durham
THE UNIVERSITY OF EDINBURGH
Imperial College London
LANCASTER UNIVERSITY
UNIVERSITY of GLASGOW
THE UNIVERSITY of MANCHESTER
THE UNIVERSITY of LIVERPOOL
Queen Mary University of London
Royal Holloway University of London
Science & Technology Facilities Council
THE UNIVERSITY OF SHEFFIELD
THE UNIVERSITY OF SUSSEX
PRIFYSGOL CYMRU ABERTAWE
UNIVERSITY OF WALES SWANSEA
UCL
WARWICK

Science & Technology Facilities Council

…and the experiments

GridPP UK Computing for Particle Physics

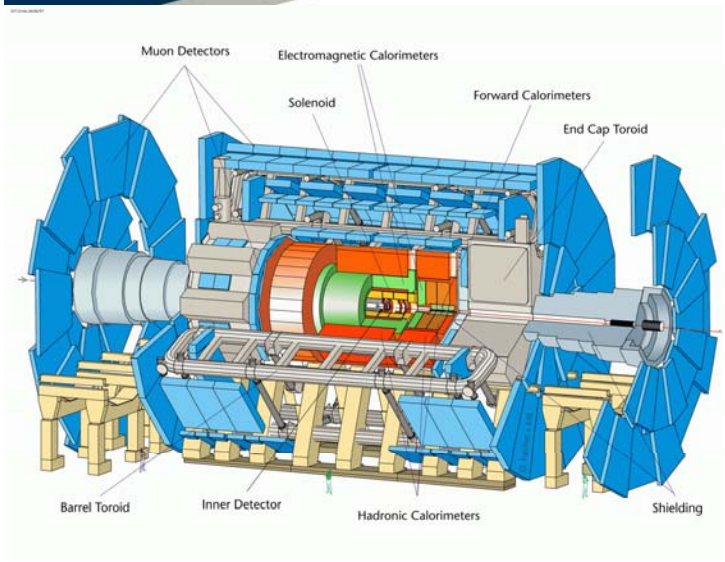

ILC

SNO

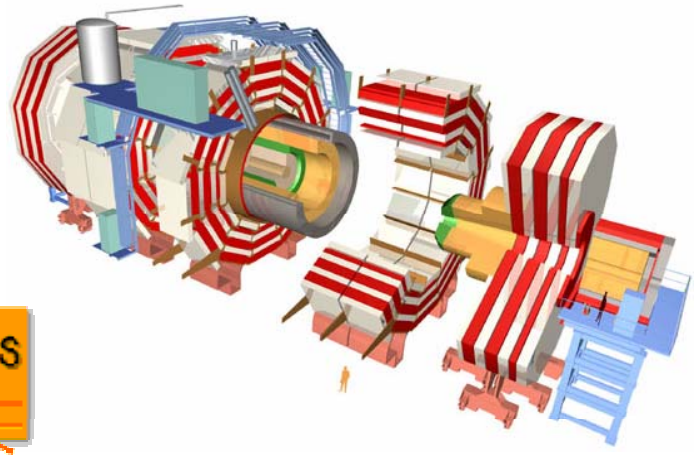

ZEUS

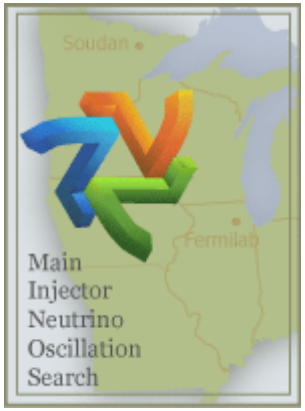

MICE

T2K

BABAR

H1

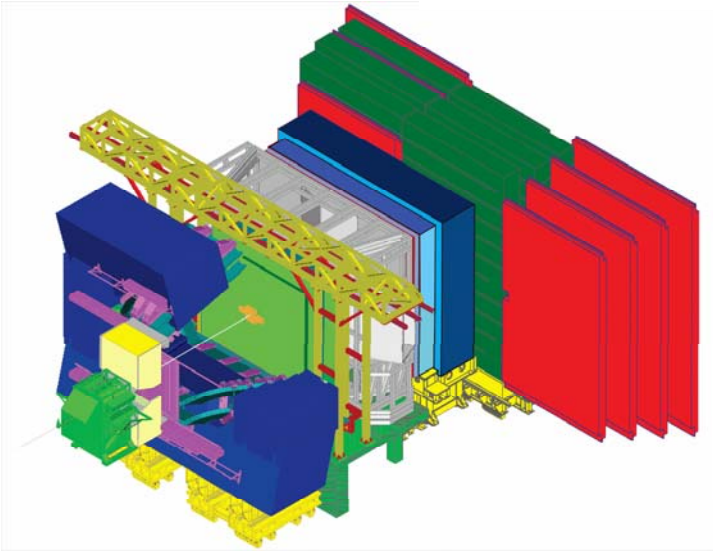

UKQCD collaboration

DØ

CDF II

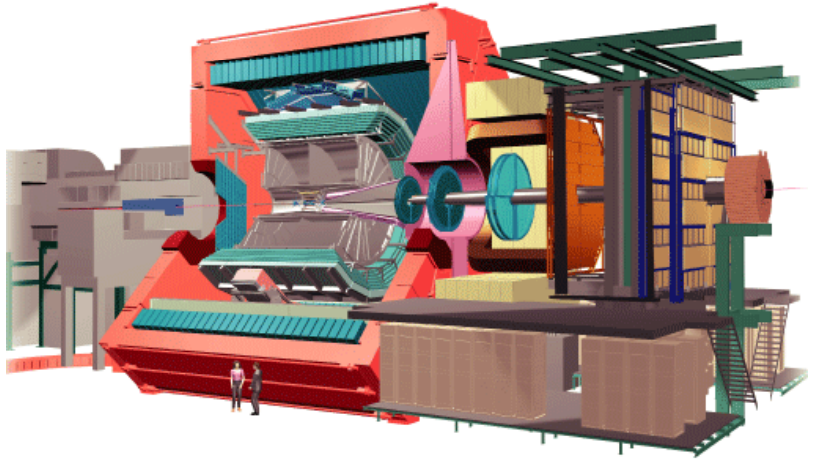

Worldwide LHC Computing Grid

Computing Technical Design Reports
WLCG + 4 LHC experiments – June 2005

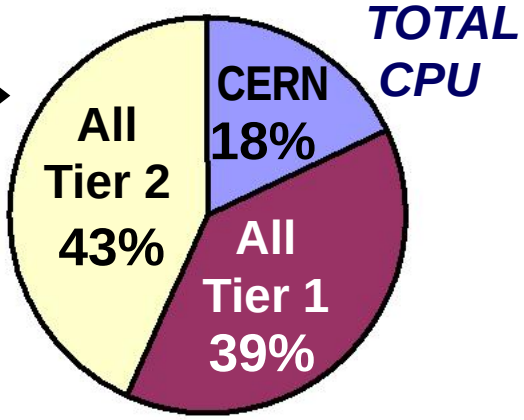

Summary of Computing Resource Requirements

All experiments - 2008

From LCG TDR - June 2005

New data will grow at ~15 PetaBytes/year

~100,000 fastest CPU cores

	CERN	*All Tier-1s*	*All Tier-2s*	*Total*
CPU (MSPECint2000s)	**25**	**56**	**61**	**142**
Disk (PetaBytes)	**7**	**31**	**19**	**57**
Tape (PetaBytes)	**18**	**35**		**53**

Need more than twice as much by 2010!

Grid Infrastructure

WLCG based on EGEE and OSG Grid infrastructure.

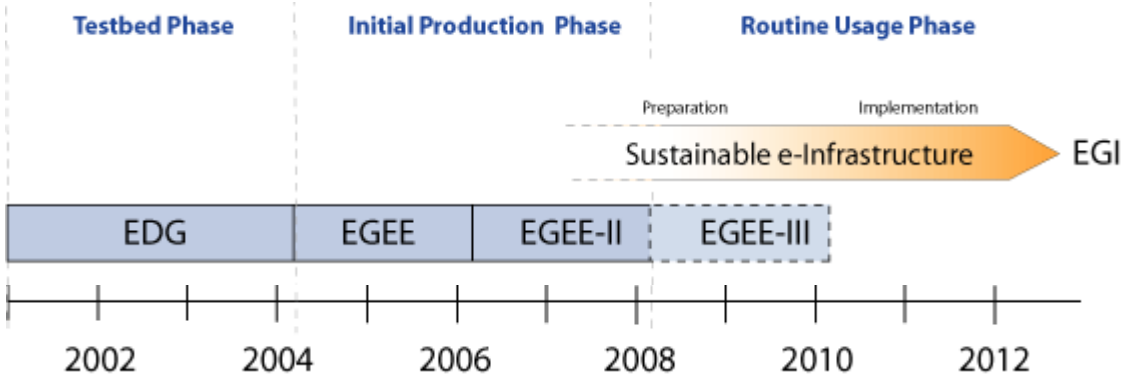

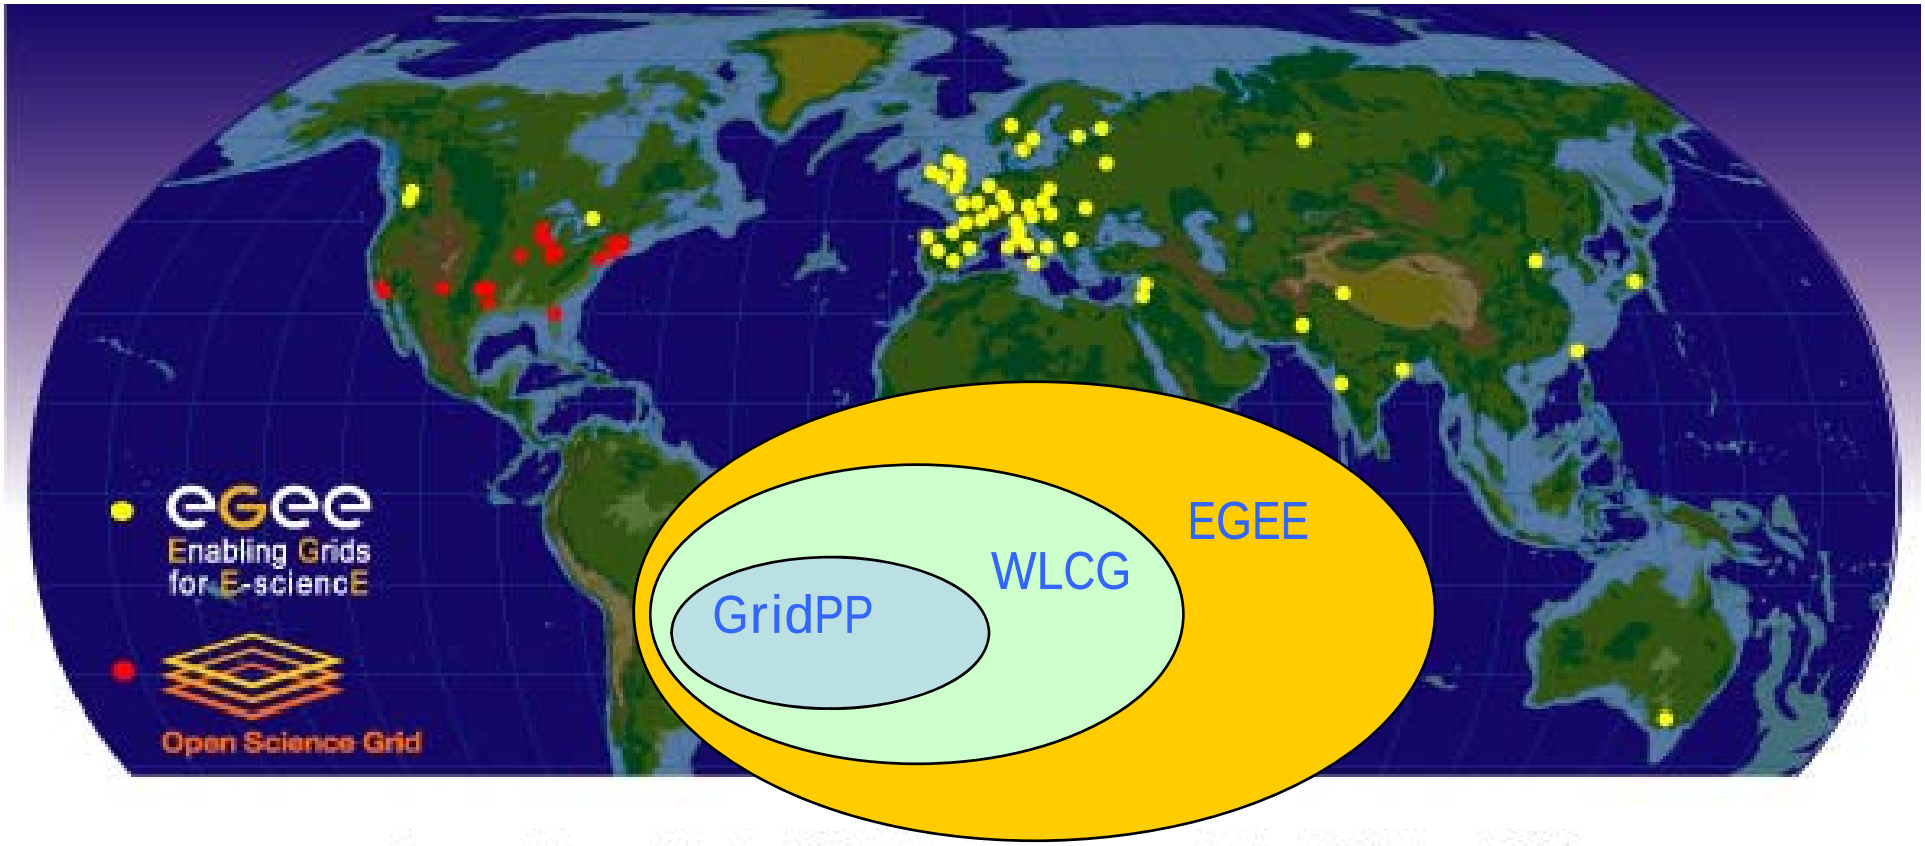

A map of the worldwide LCG infrastructure operated by EGEE and OSG.

WLCG Status

March 2007

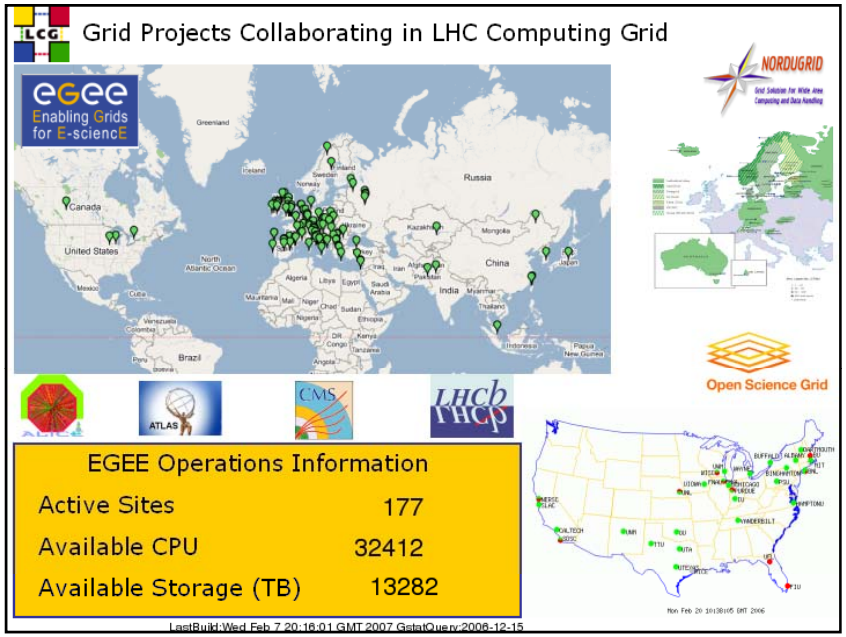

Status in **2007**:
177 sites, 32,412 CPUs,
13,282 TB storage

(Dave Britton, IOP Conference Guildford, April 2007)

March 2008

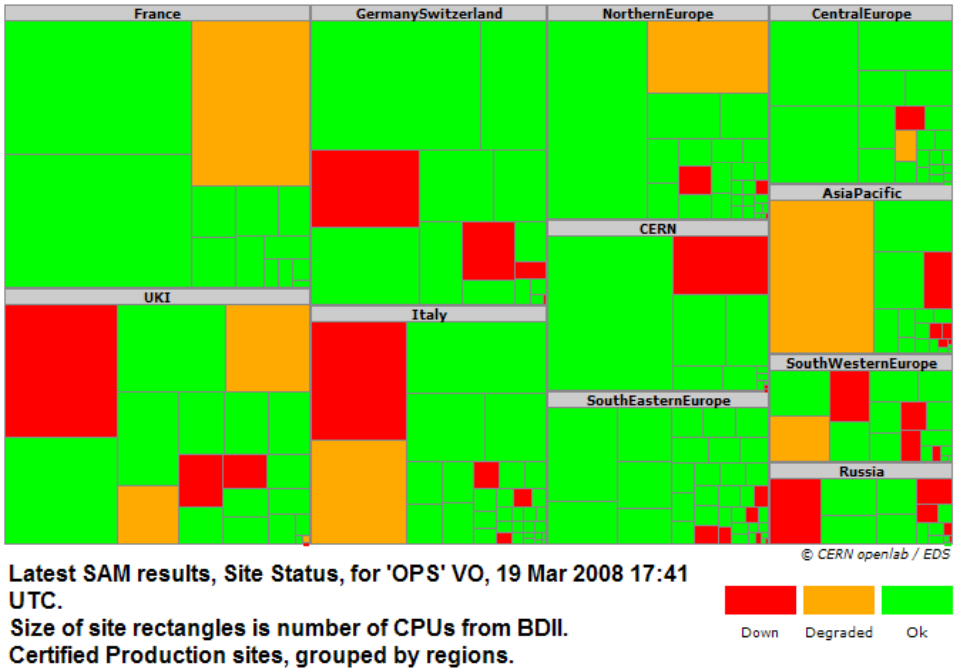

Latest SAM results, Site Status, for 'OPS' VO, 19 Mar 2008 17:41 UTC.
Size of site rectangles is number of CPUs from BDII.
Certified Production sites, grouped by regions.

Down Degraded Ok

Status on 29 March 2008:

250 sites, 50 countries

55,094 CPUs (MaxCPU=70,973)

21,079 TB storage (?)

UK - CPU contribution

In 2006, UKI provided **29%** of the EGEE CPU

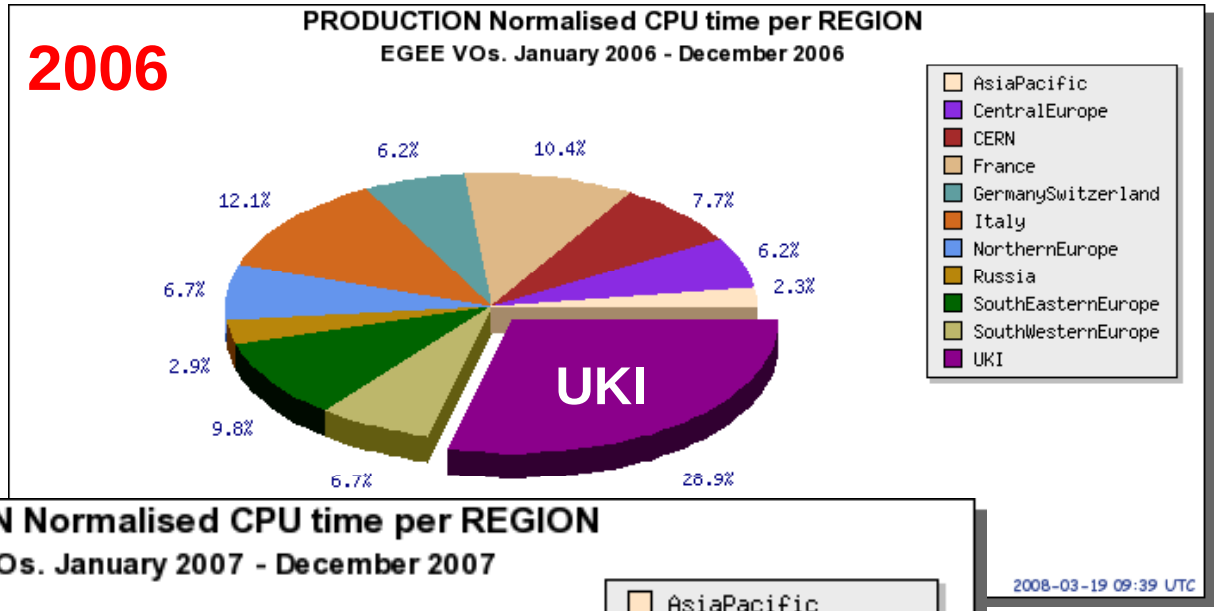

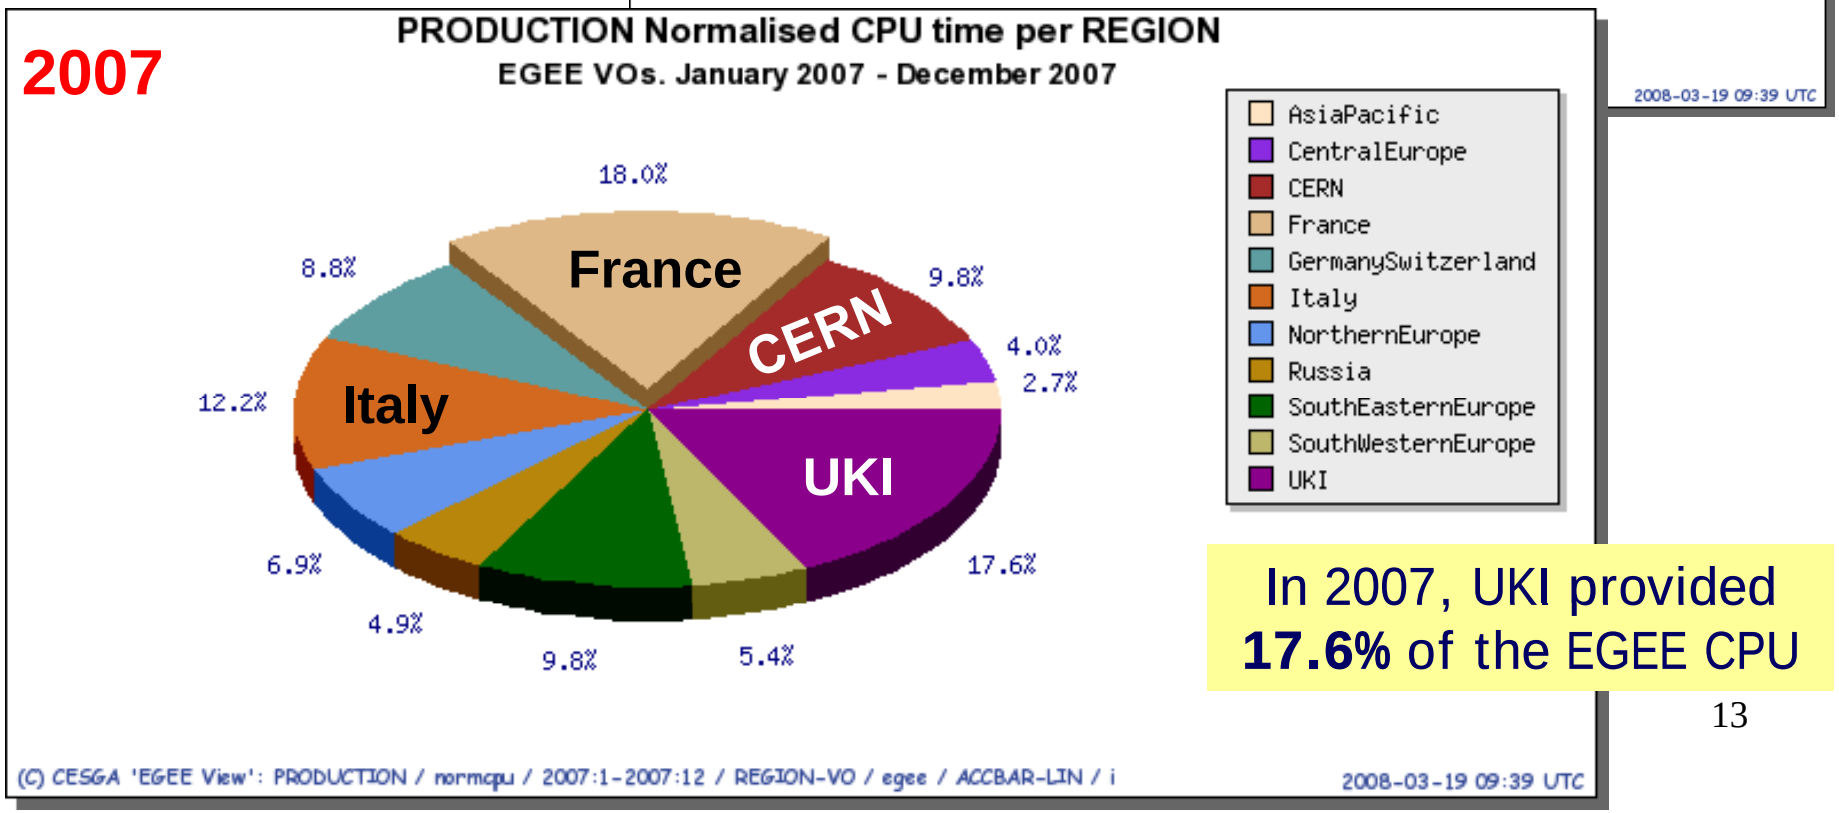

In 2007, UKI provided **17.6%** of the EGEE CPU

Experiments – CPU Use

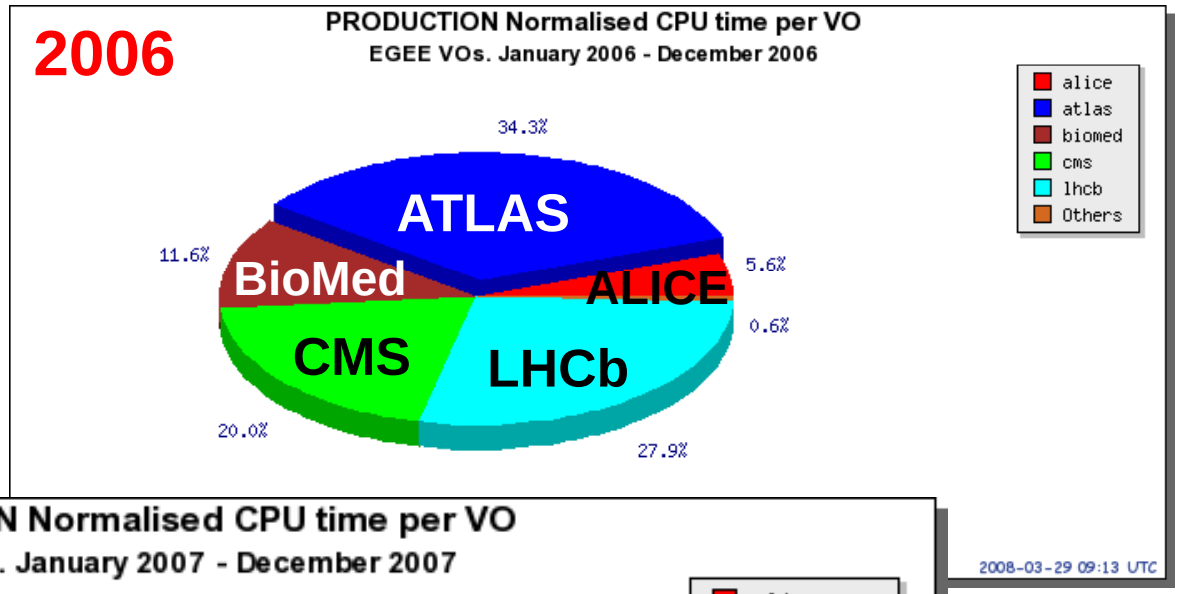

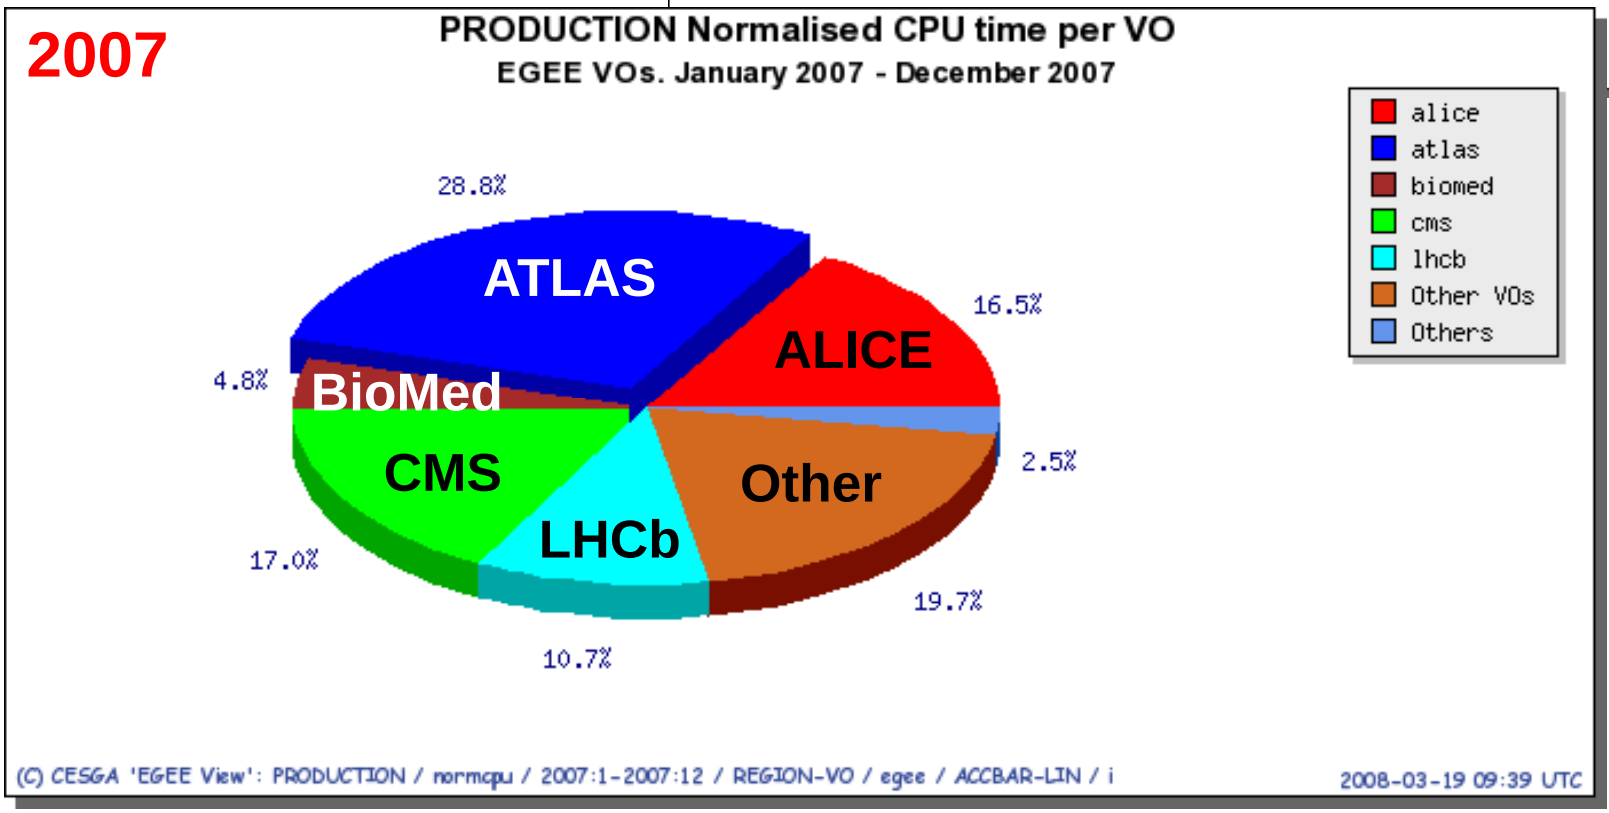

Building the UK Tier 1

Science & Technology Facilities Council

GridPP UK Computing for Particle Physics

27 March 2007

History repeats …

Replaces ATLAS Centre.

Current home to UK Tier 1.

UK Tier 1 Hardware

From ~April 2008

CPU Cluster with over **3,200 cores** delivering **4,500 KSI2K**.

Disk **340 servers** spinning over **5,400 drives** and providing over **2.3 PetaBytes** of disk storage.

Tape SL8500 tape robot with **2 PetaBytes** of media, 18*T10K drives and 8*9940B drives. 10,000 slots → **5 PetaBytes**


CPU Cluster 2004


Disk Cluster 2003


Disk Cluster 2008


Tape Robot 2005

LHCOPN – current status

LHC Optical Private Network

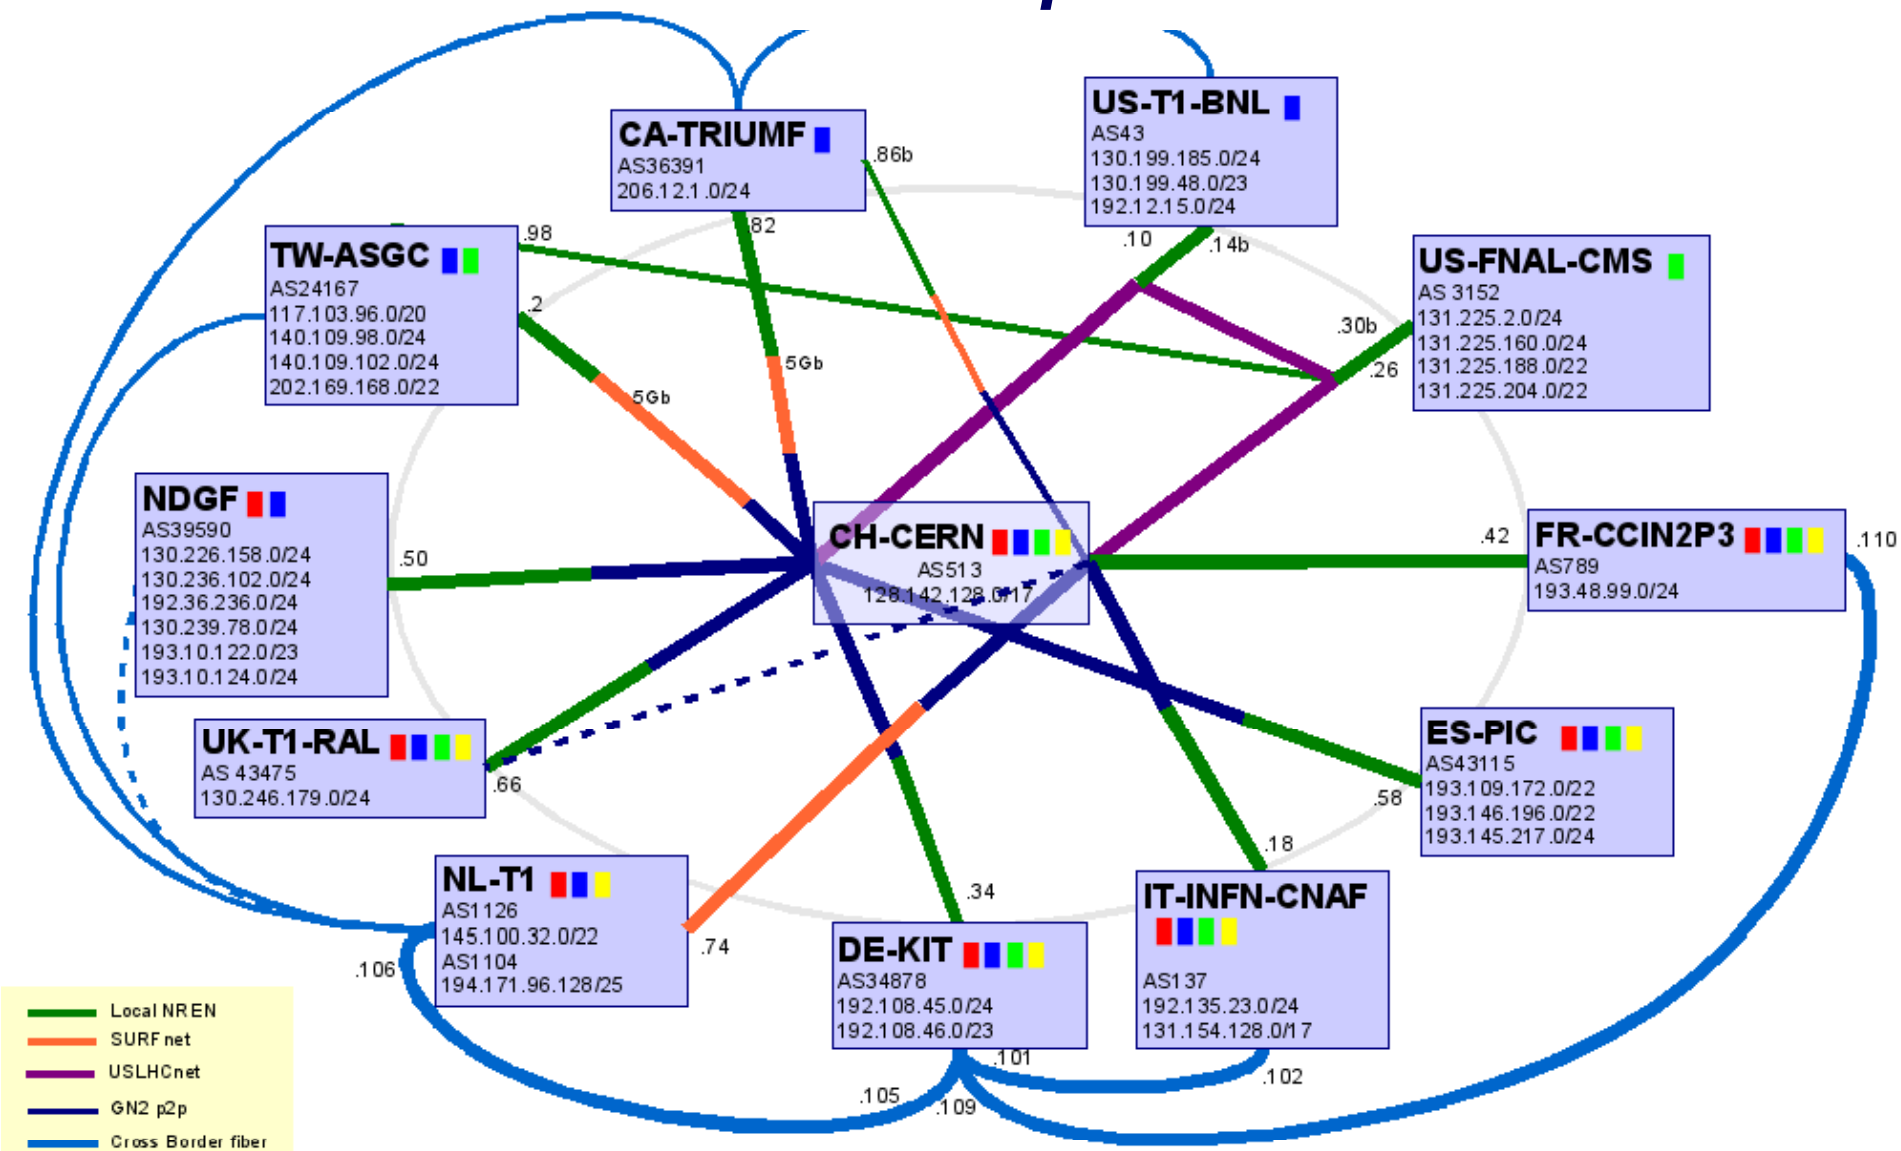

10 Gb/s dedicated light path RAL - CERN

edoardo.martelli@cern.ch 20080312

Optical Private Network Traffic

UK Tier 1 will receive data from CERN at 350 MB/s (average)
Deliver to Tier 2 centres at 500 MB/s (peak)

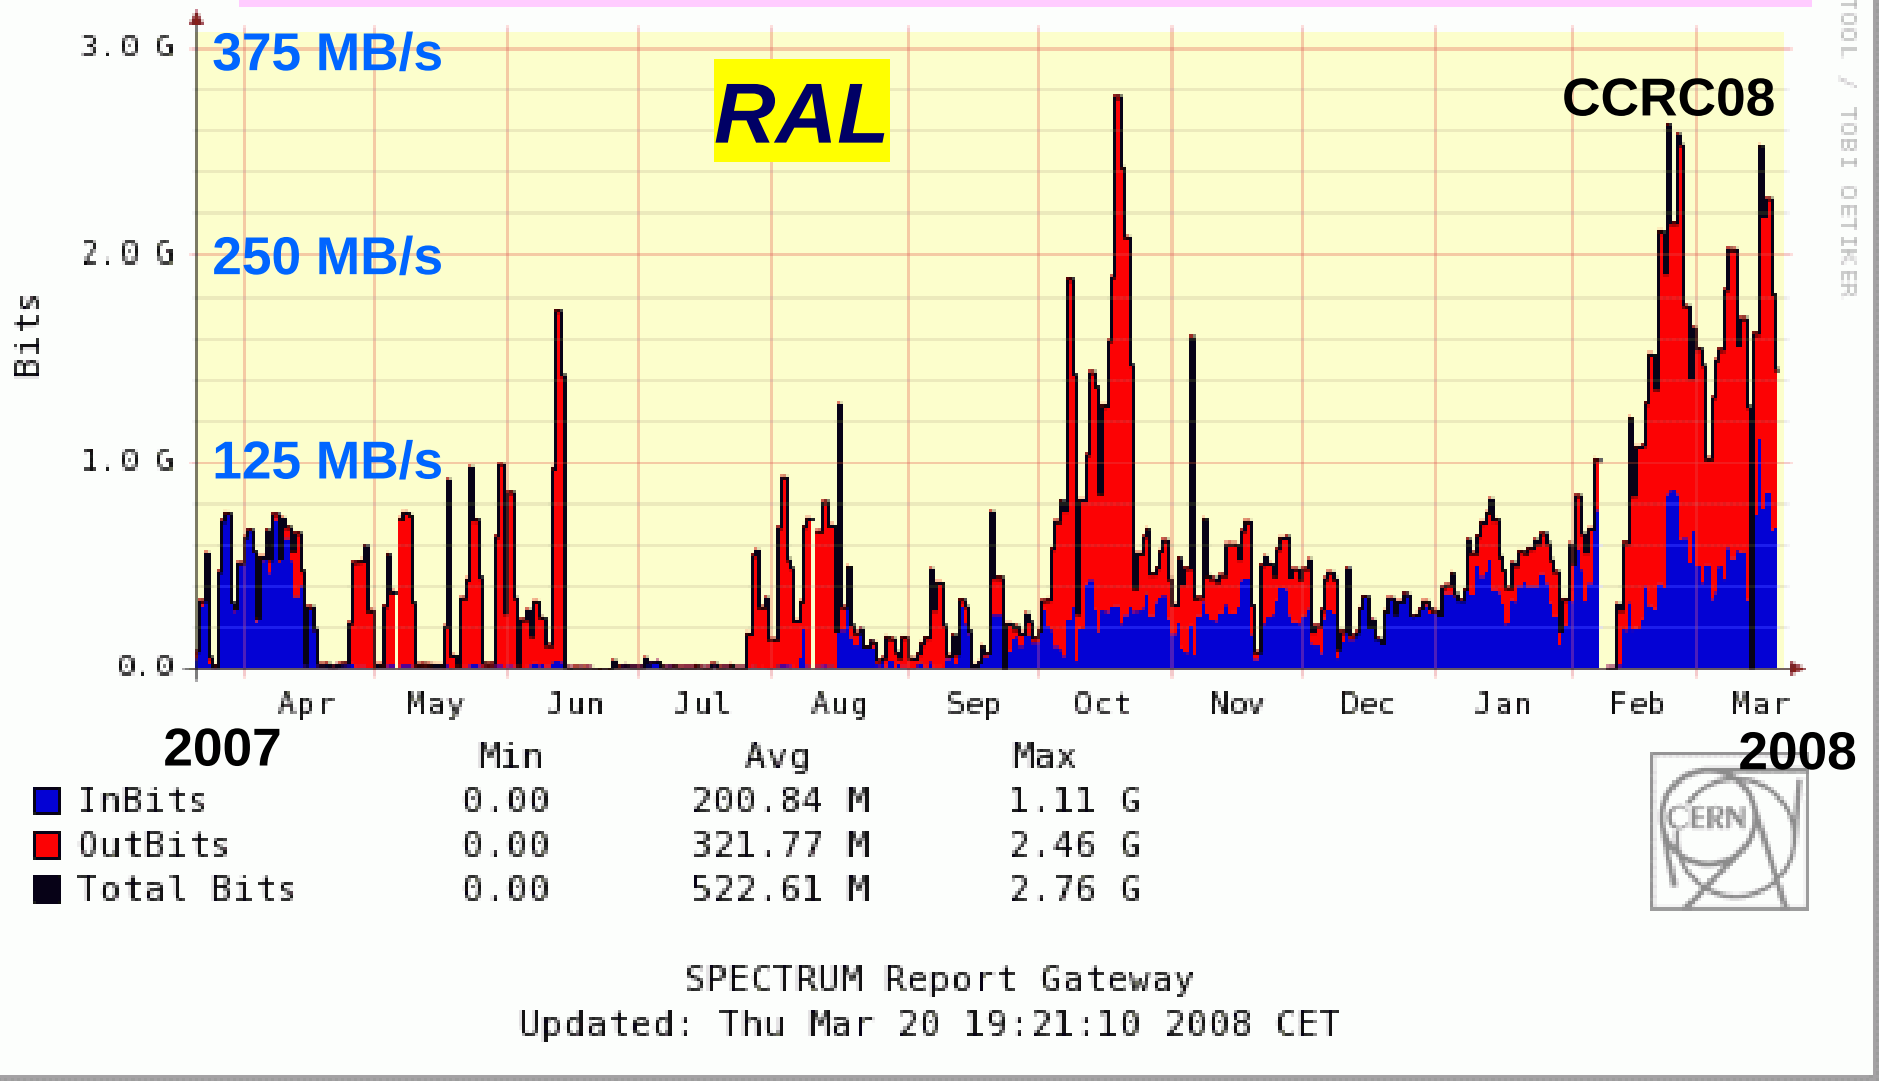

Experiment Data Transfers

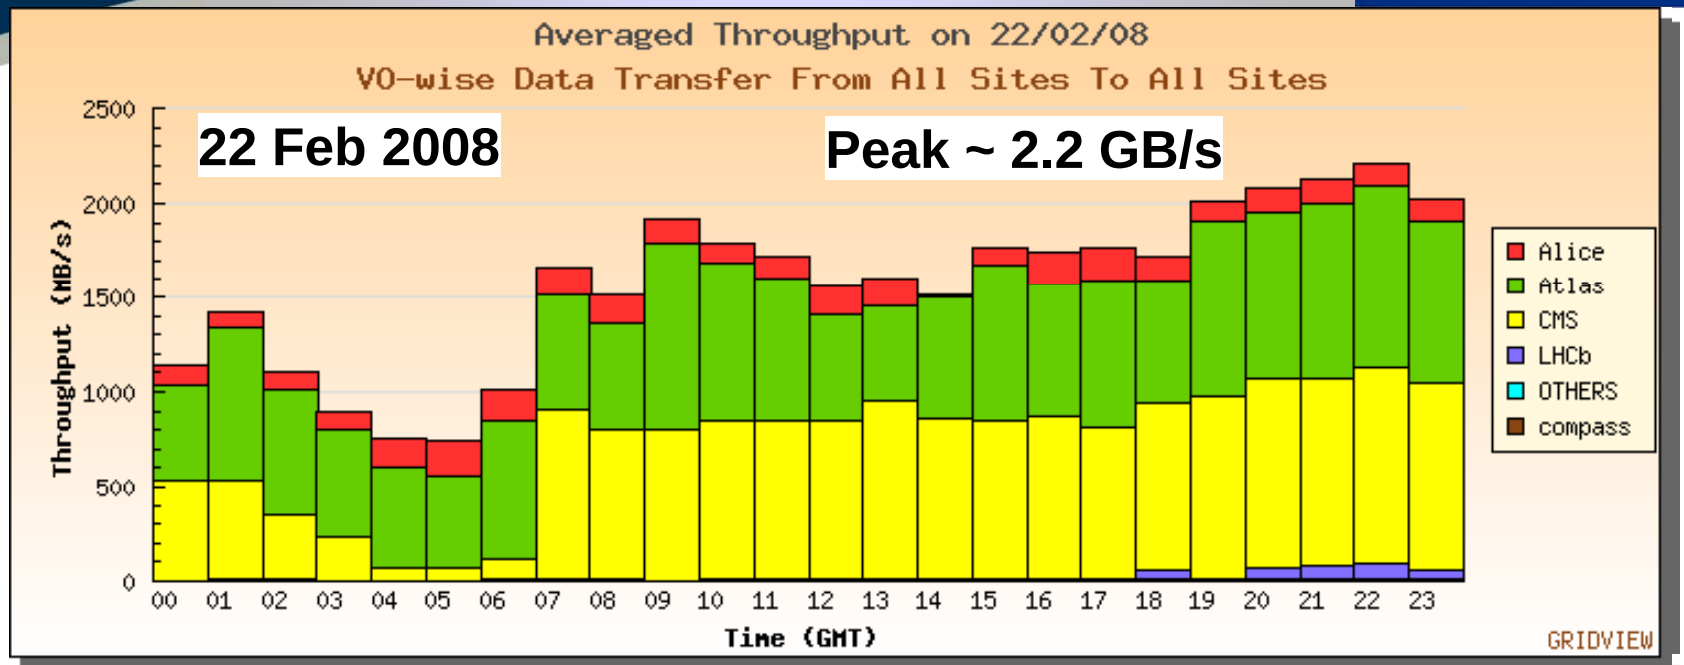

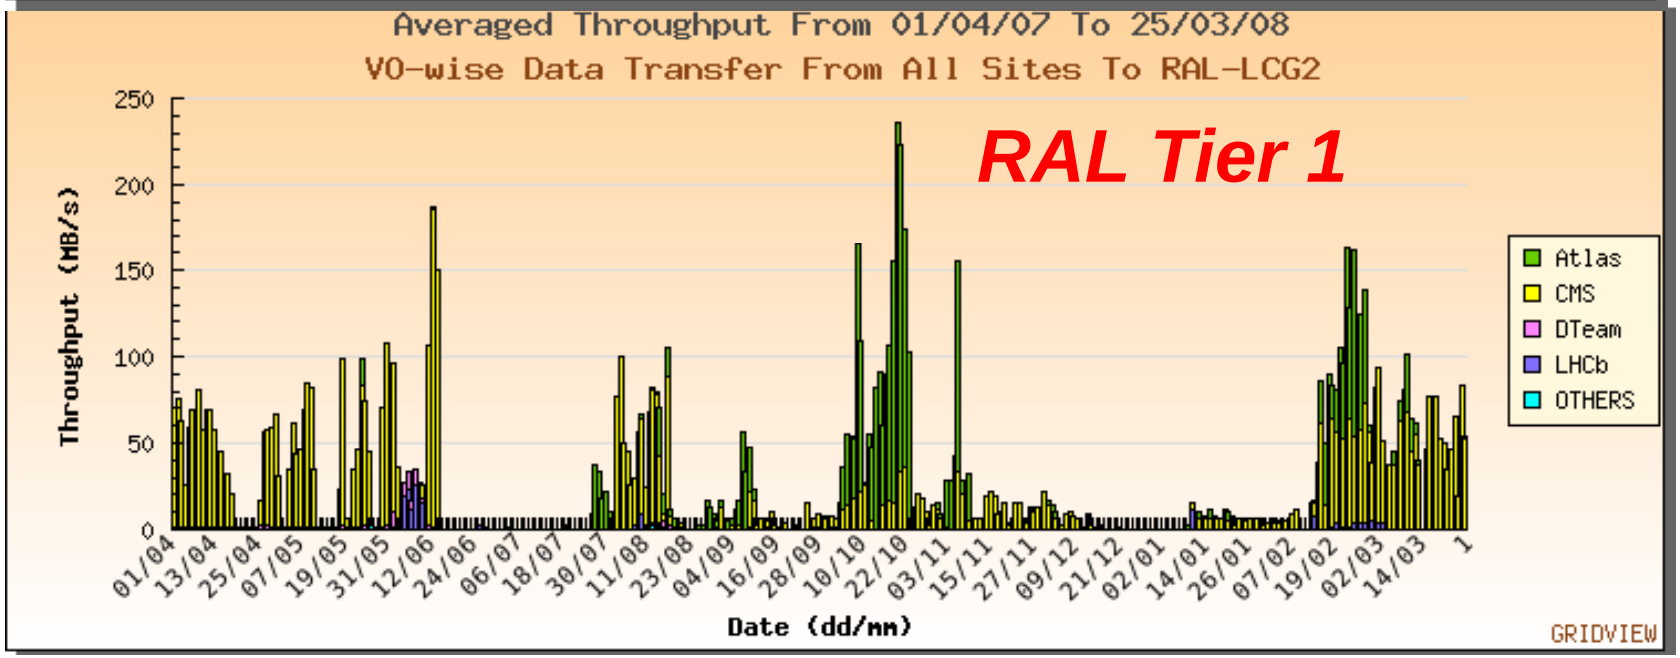

Storage Middleware

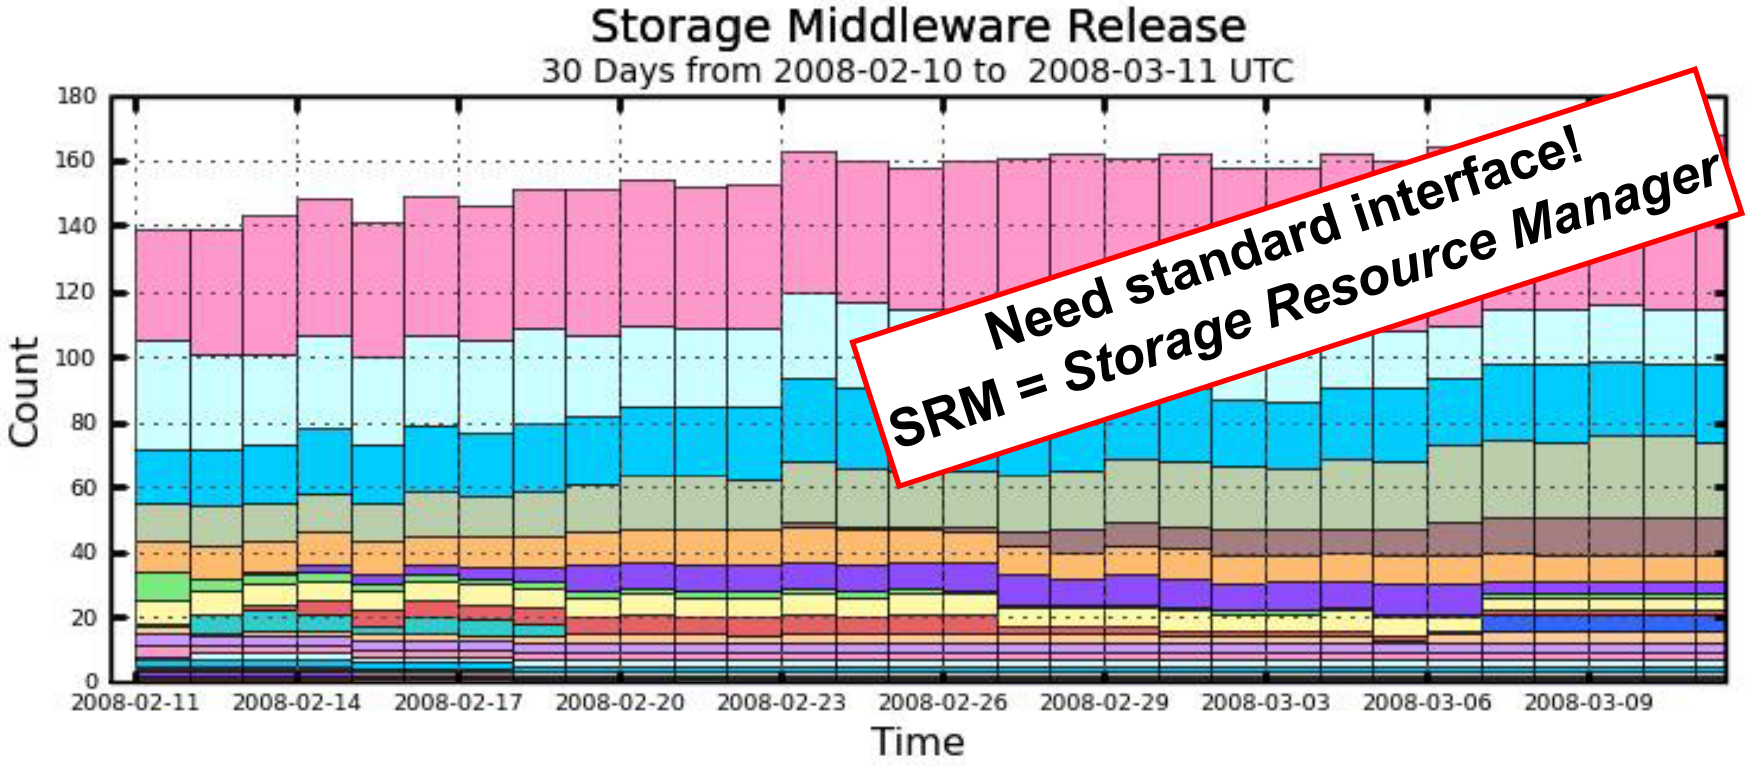

c/o Jens Jensen

RAL Tier 1 - migrating from dCache to Castor 2
A journey which started in 2006

UK T1 – Castor2 Architecture

Separate Castor instances for each experiment essential!

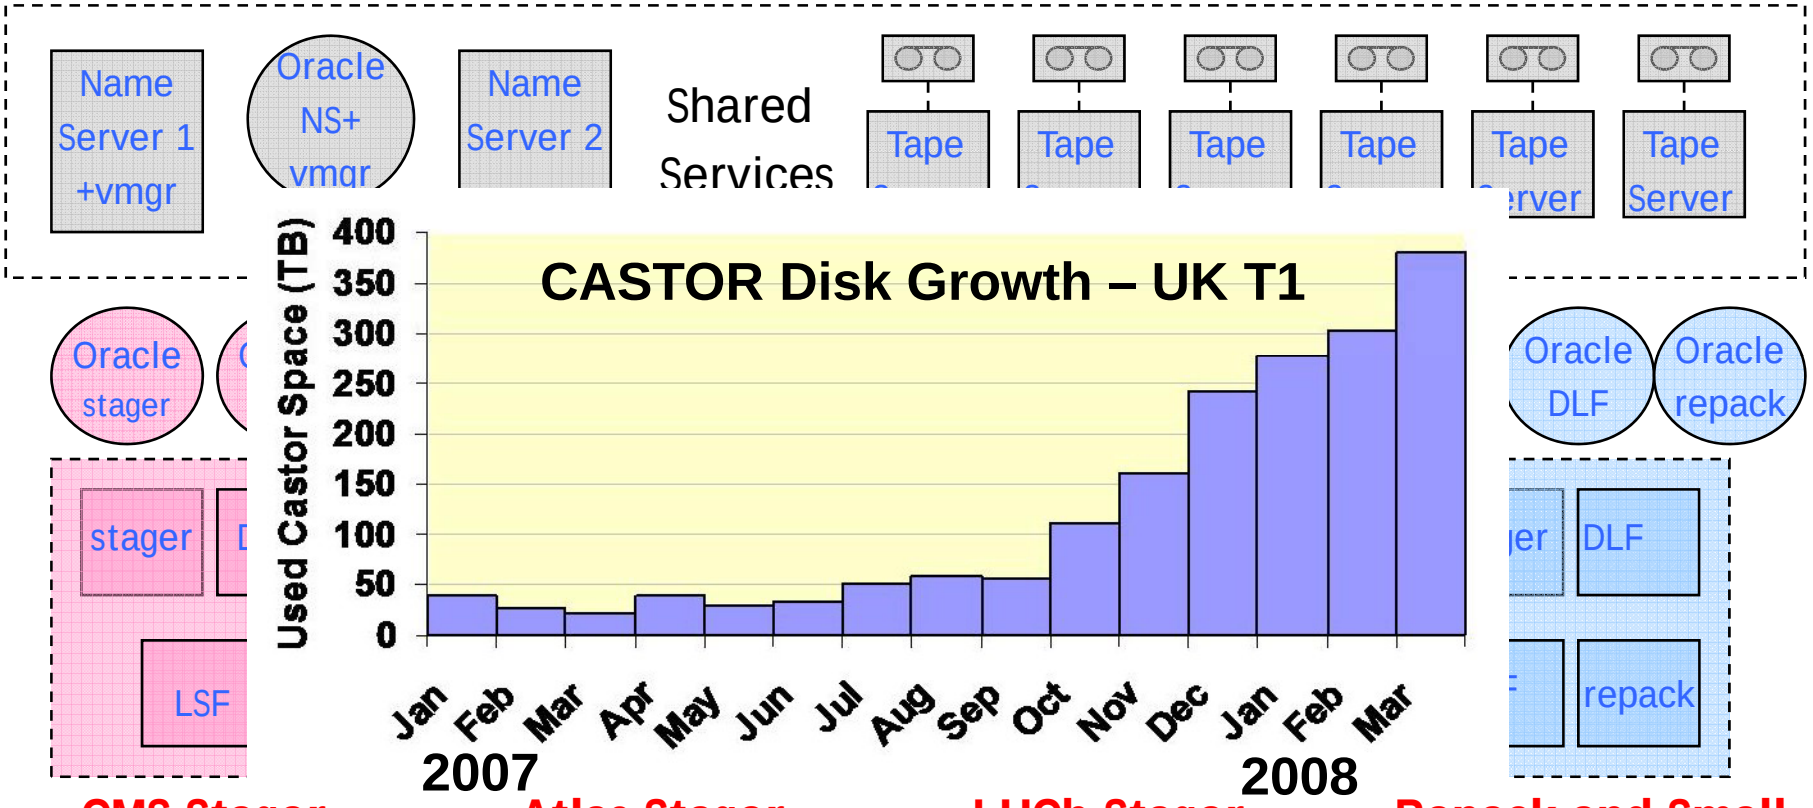

CMS Stager Instance

Diskservers

Atlas Stager Instance

Diskservers

LHCb Stager Instance

Diskservers

Repack and Small User Stager Instance

1 Diskserver

CMS – Lots of tuning!

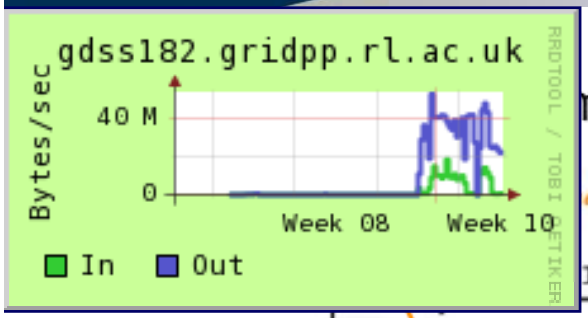

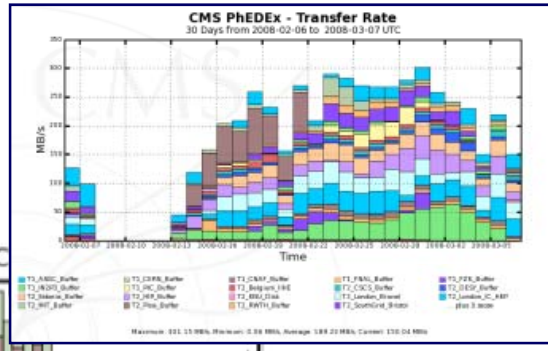

n on week-3 (all T1s together)

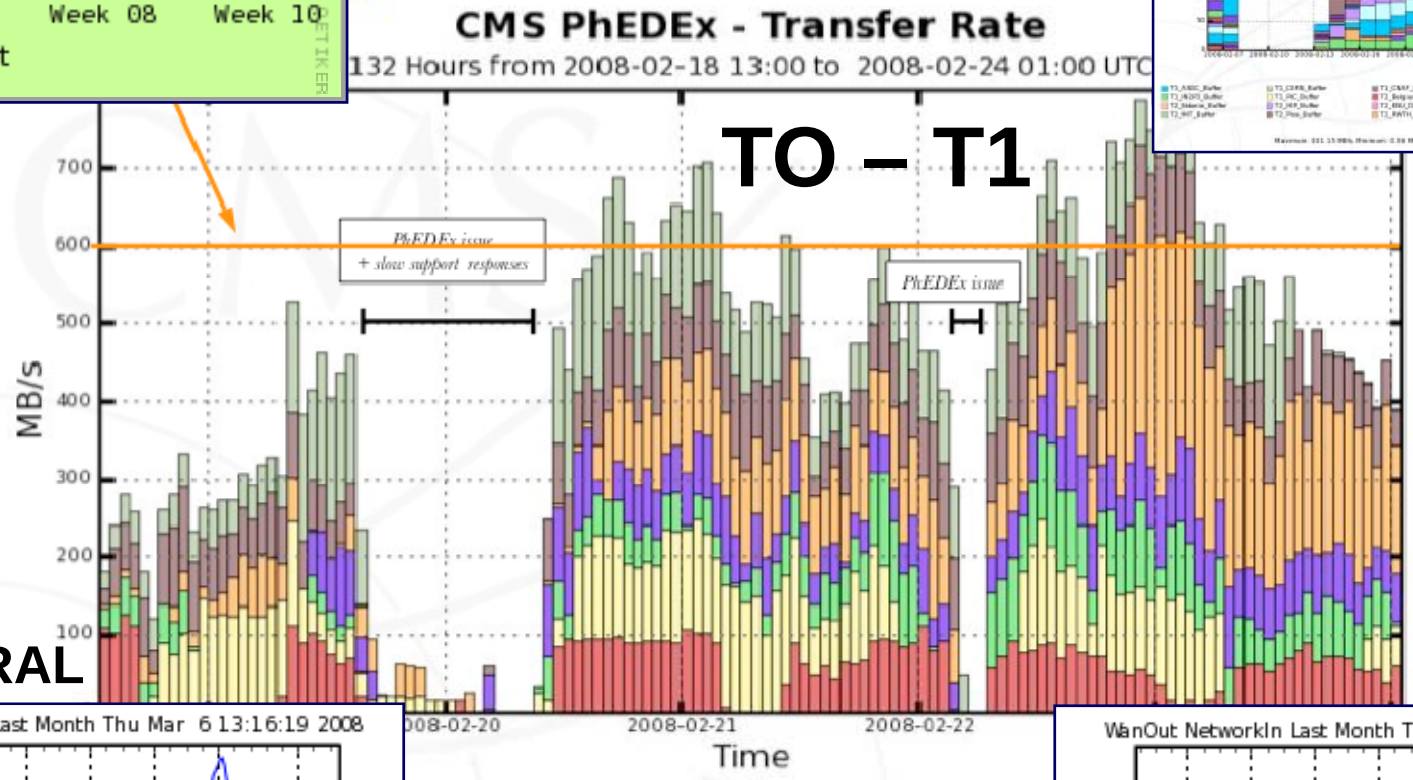

TO – T1

RAL

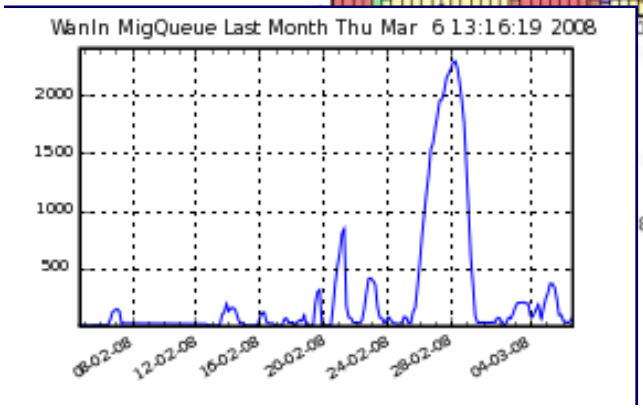

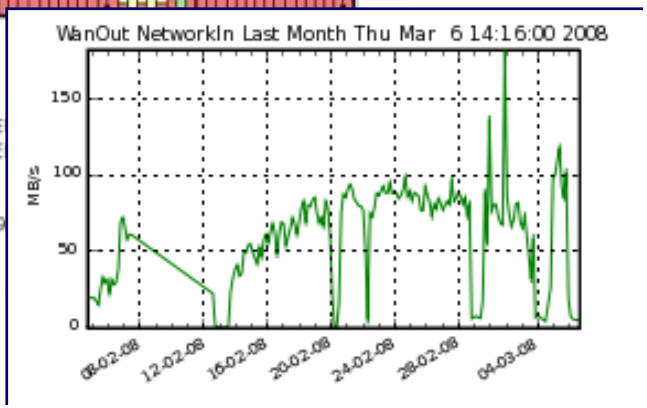

UK Tier 2 Centres

Science & Technology Facilities Council

GridPP UK Computing for Particle Physics

ScotGrid Scottish Grid Service

NorthGrid

LT2

SouthGrid

ScotGrid

Durham, Edinburgh, Glasgow

NorthGrid

Daresbury, Lancaster, Liverpool, Manchester, Sheffield

SouthGrid

Birmingham, Bristol, Cambridge, Oxford, RAL PPD, Warwick

London

Brunel, Imperial, QMUL, RHUL, UCL

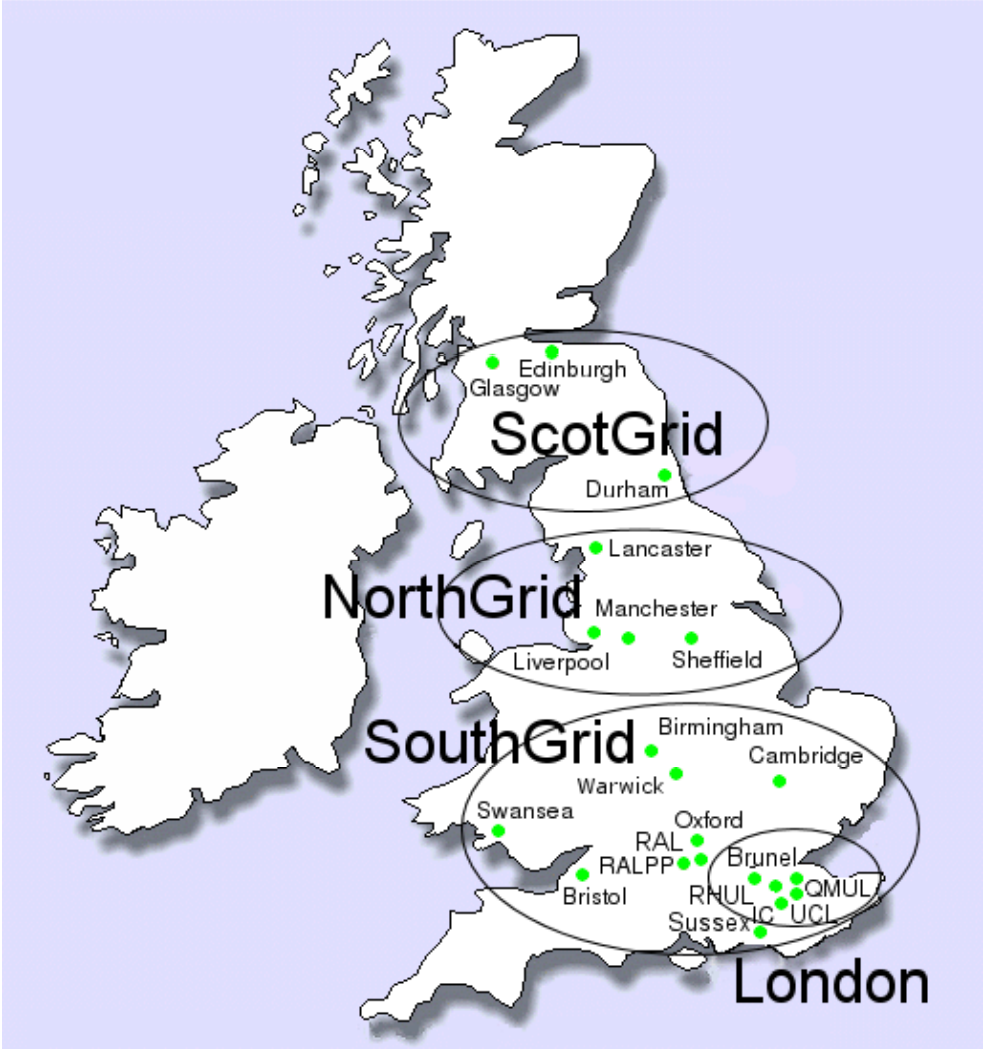

An Example - NorthGrid

Lancaster

Liverpool

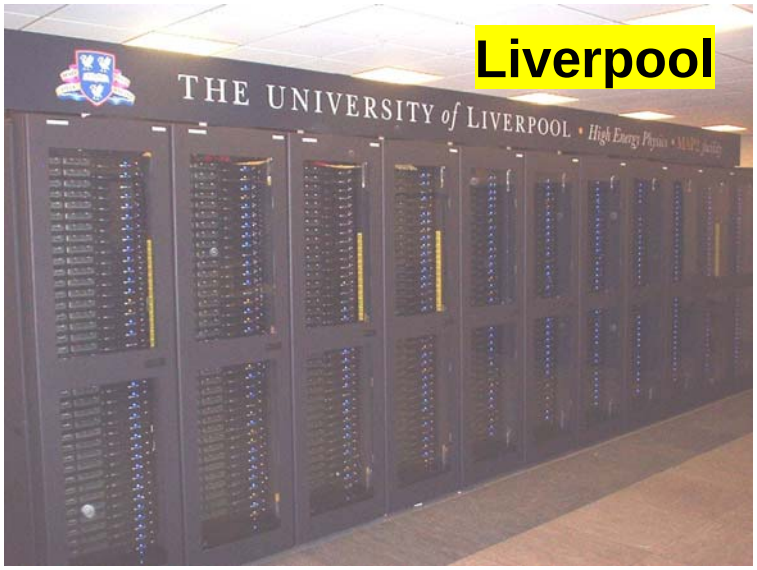

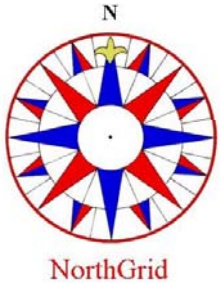

Manchester

Sheffield

Tier 2 Resources

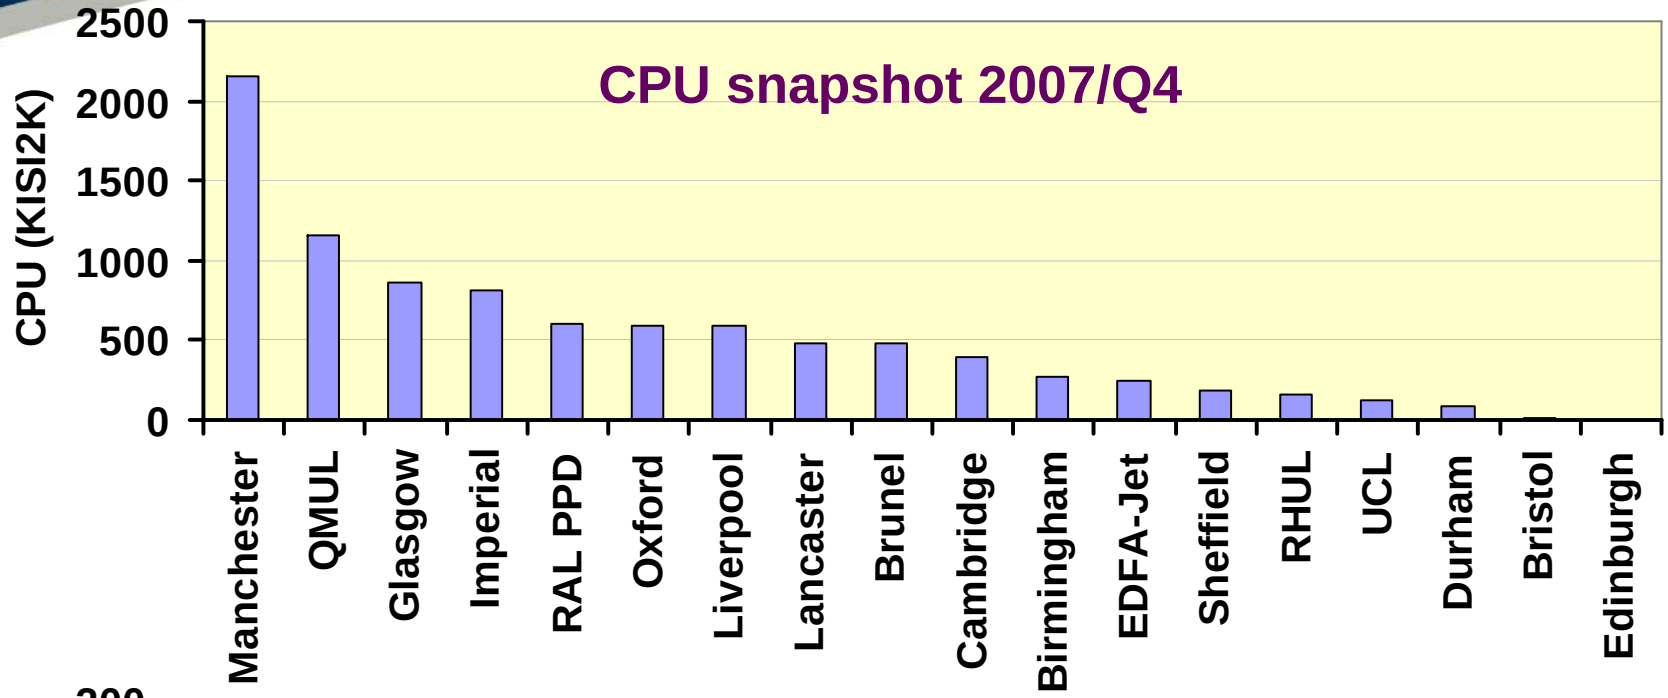

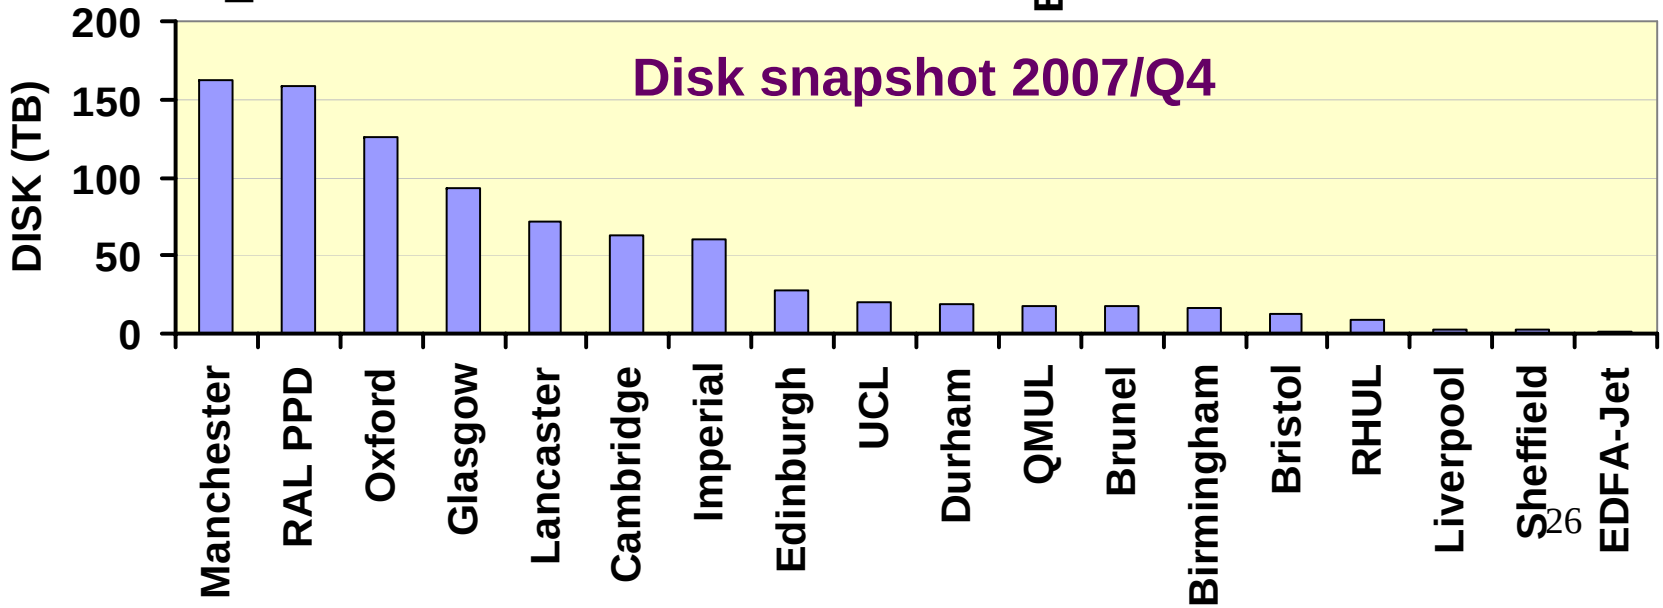

CPU by UK Region

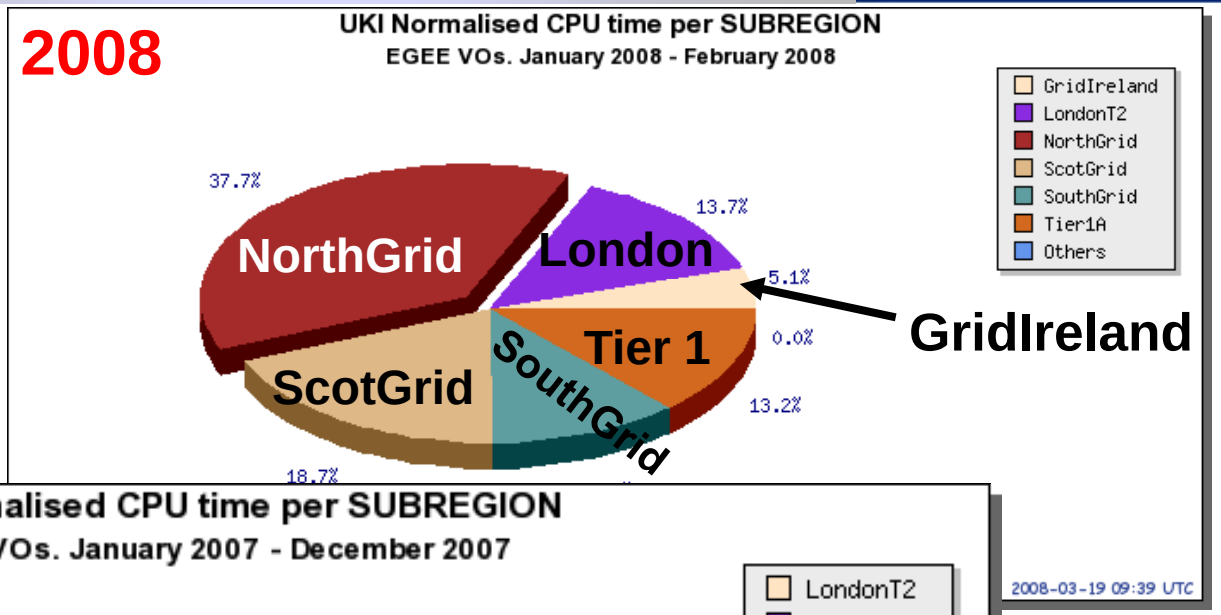

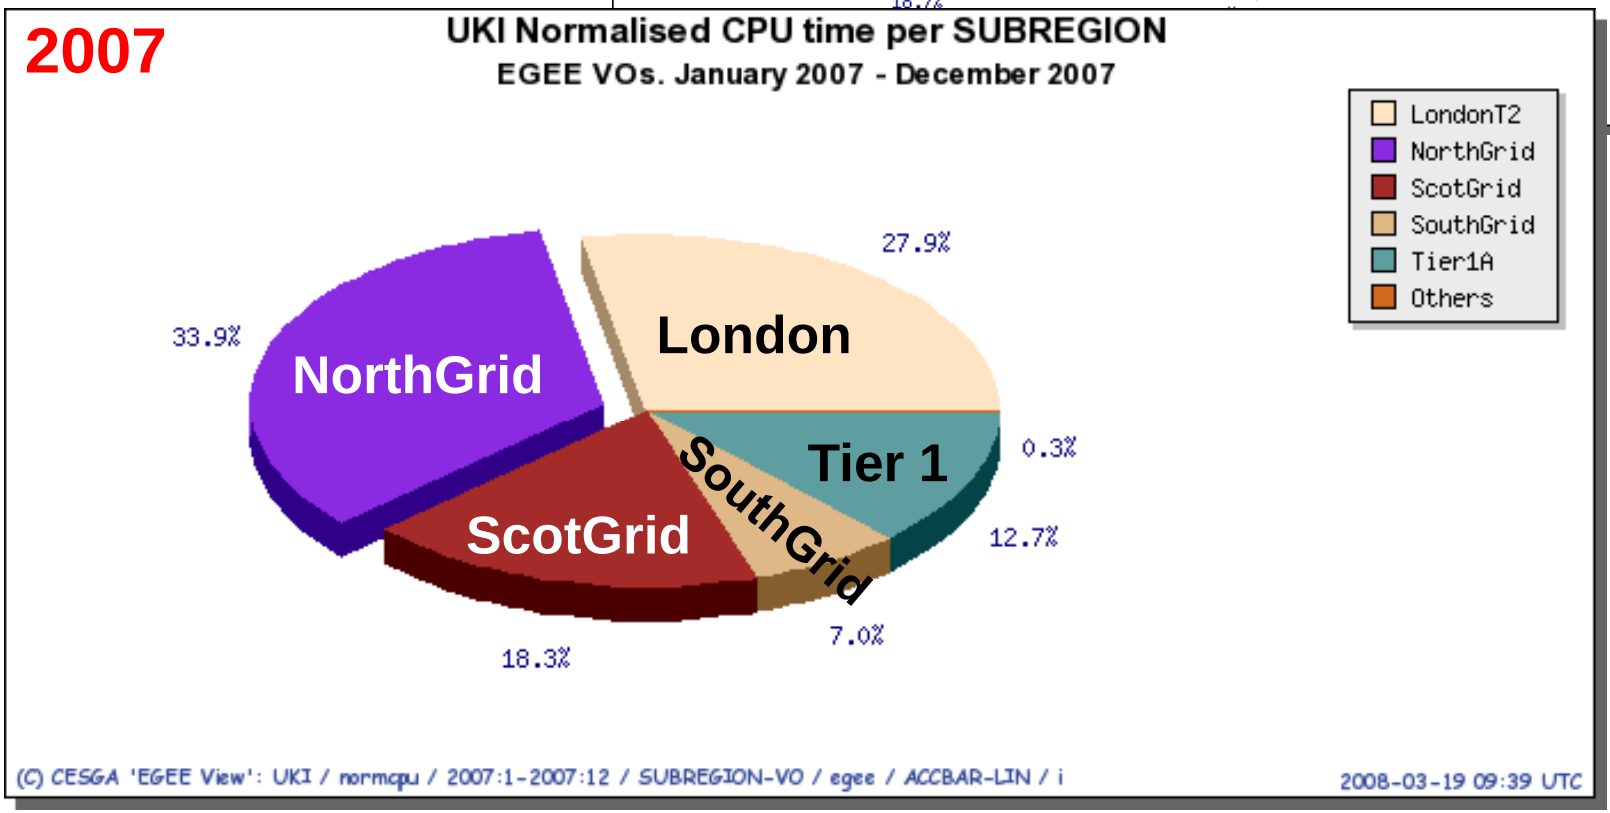

T0-T1-T2 Data Movement

Big challenge to make this efficient – esp. for analysis of GPD experiments.

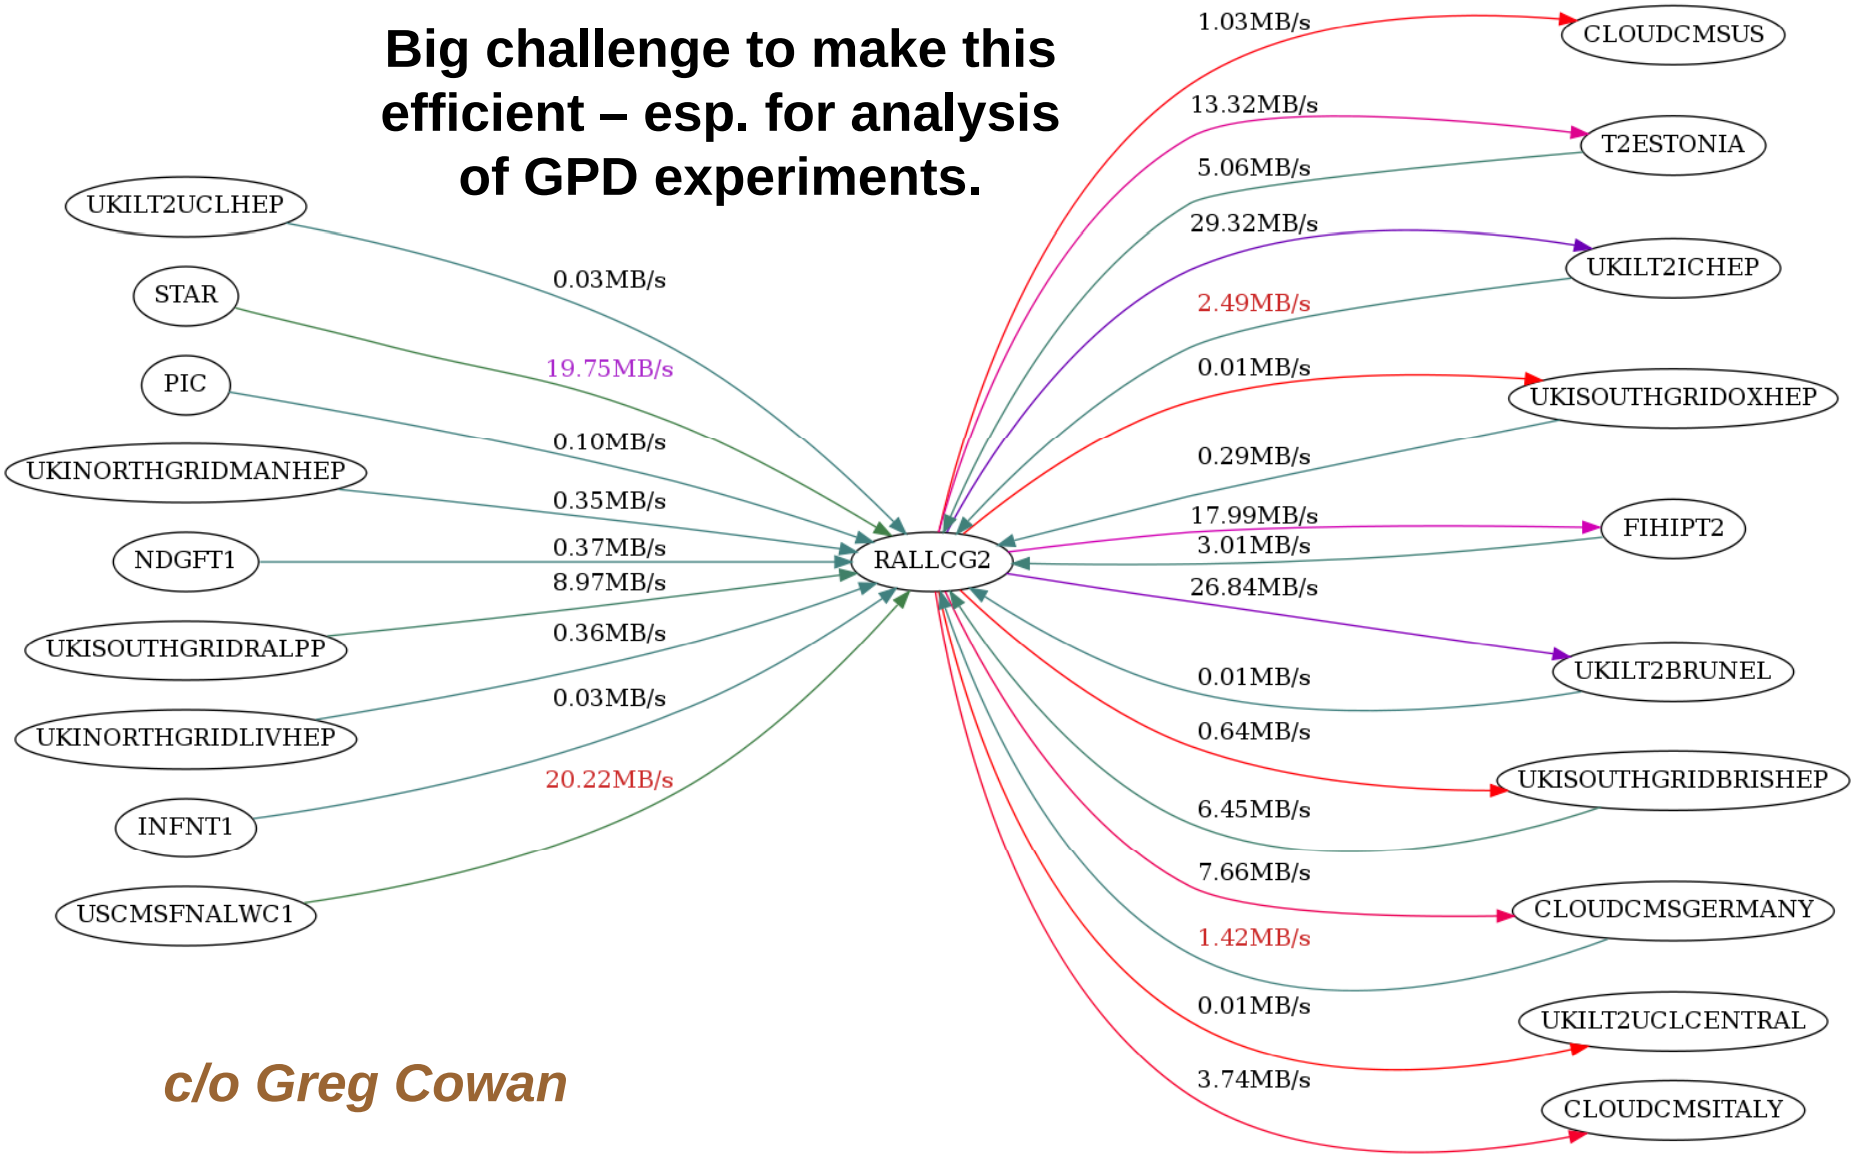

c/o Greg Cowan

RAL FTS transfer rates for VO=all from 2008-03-03 16;36;59 to 2008-03-04 16;36;59 (last day). Coloured rate implies jobs in queue, ready to be transferred.

Grid Efficiency

Percentage of successful ATLAS test jobs each day – c/o Steve Lloyd

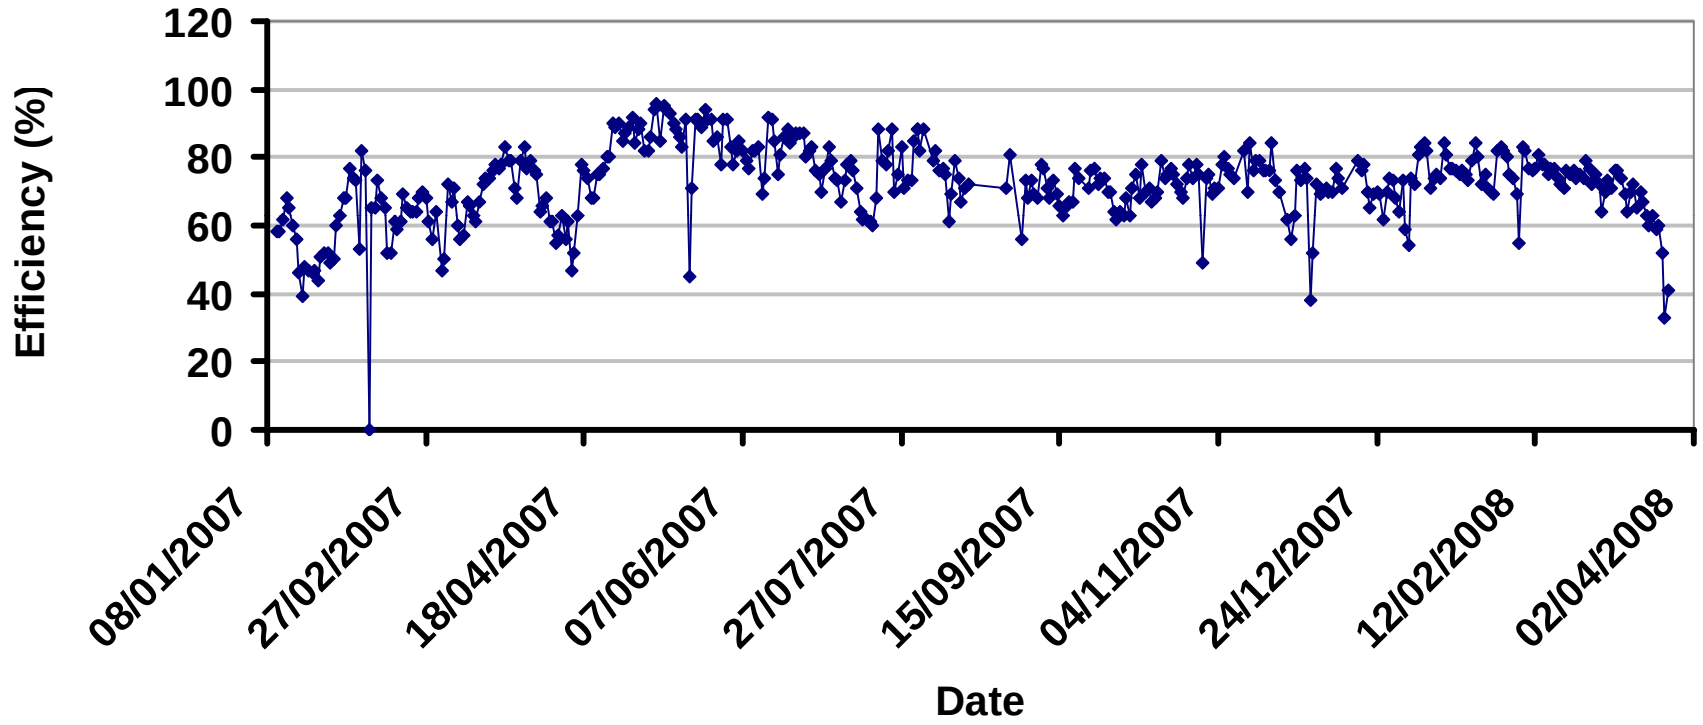

Averaged over all GridPP sites

CPU Efficiency

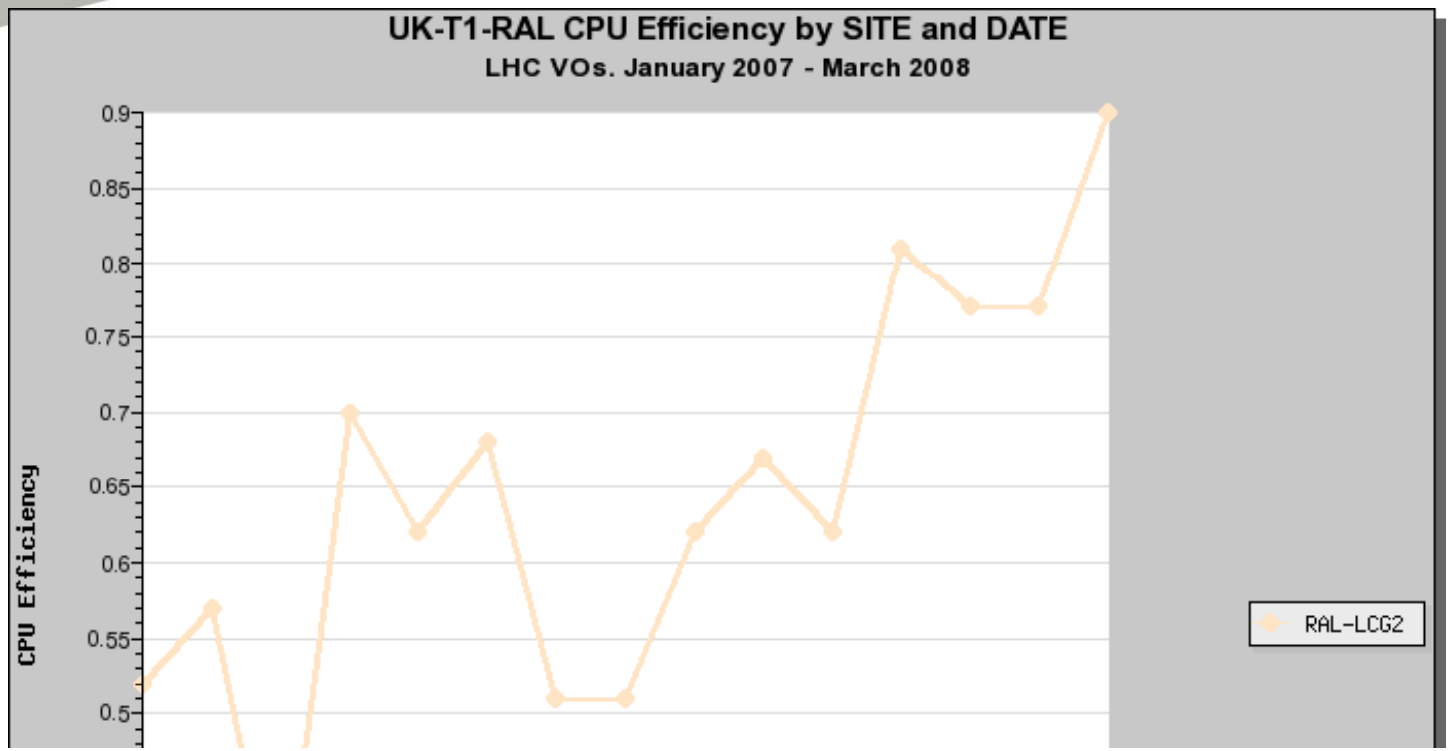

CPU Efficiency (%) by VO and DATE

VO	Jan 07	Feb 07	Mar 07	Apr 07	May 07	Jun 07	Jul 07	Aug 07	Sep 07	Oct 07	Nov 07	Dec 07	Jan 08	Feb 08	Mar 08
alice	79.4	70.5	70.4	81.1	72.4	97.5	101.0	71.2	98.0	98.1	97.7	102.9	94.4	65.4	43.1
atlas	58.3	52.4	62.8	83.4	82.2	84.5	92.6	74.8	54.6	67.0	64.4	80.9	95.0	75.5	96.3
cms	86.5	72.8	23.4	63.9	21.7	14.7	0.1	7.9	82.0	65.3	67.9	68.8	78.1	81.3	86.0
lhcb	37.7	42.2	63.8	94.9	45.6	70.1	82.7	16.5	53.3	47.9	19.9	65.3	18.9	80.7	71.5
Total	52.3	57.1	35.1	69.7	61.6	68.3	51.0	51.4	62.3	67.0	62.0	80.9	77.5	77.5	90.1

Efficiency = CPU Time/Wallclock Time

LHC Approaches!

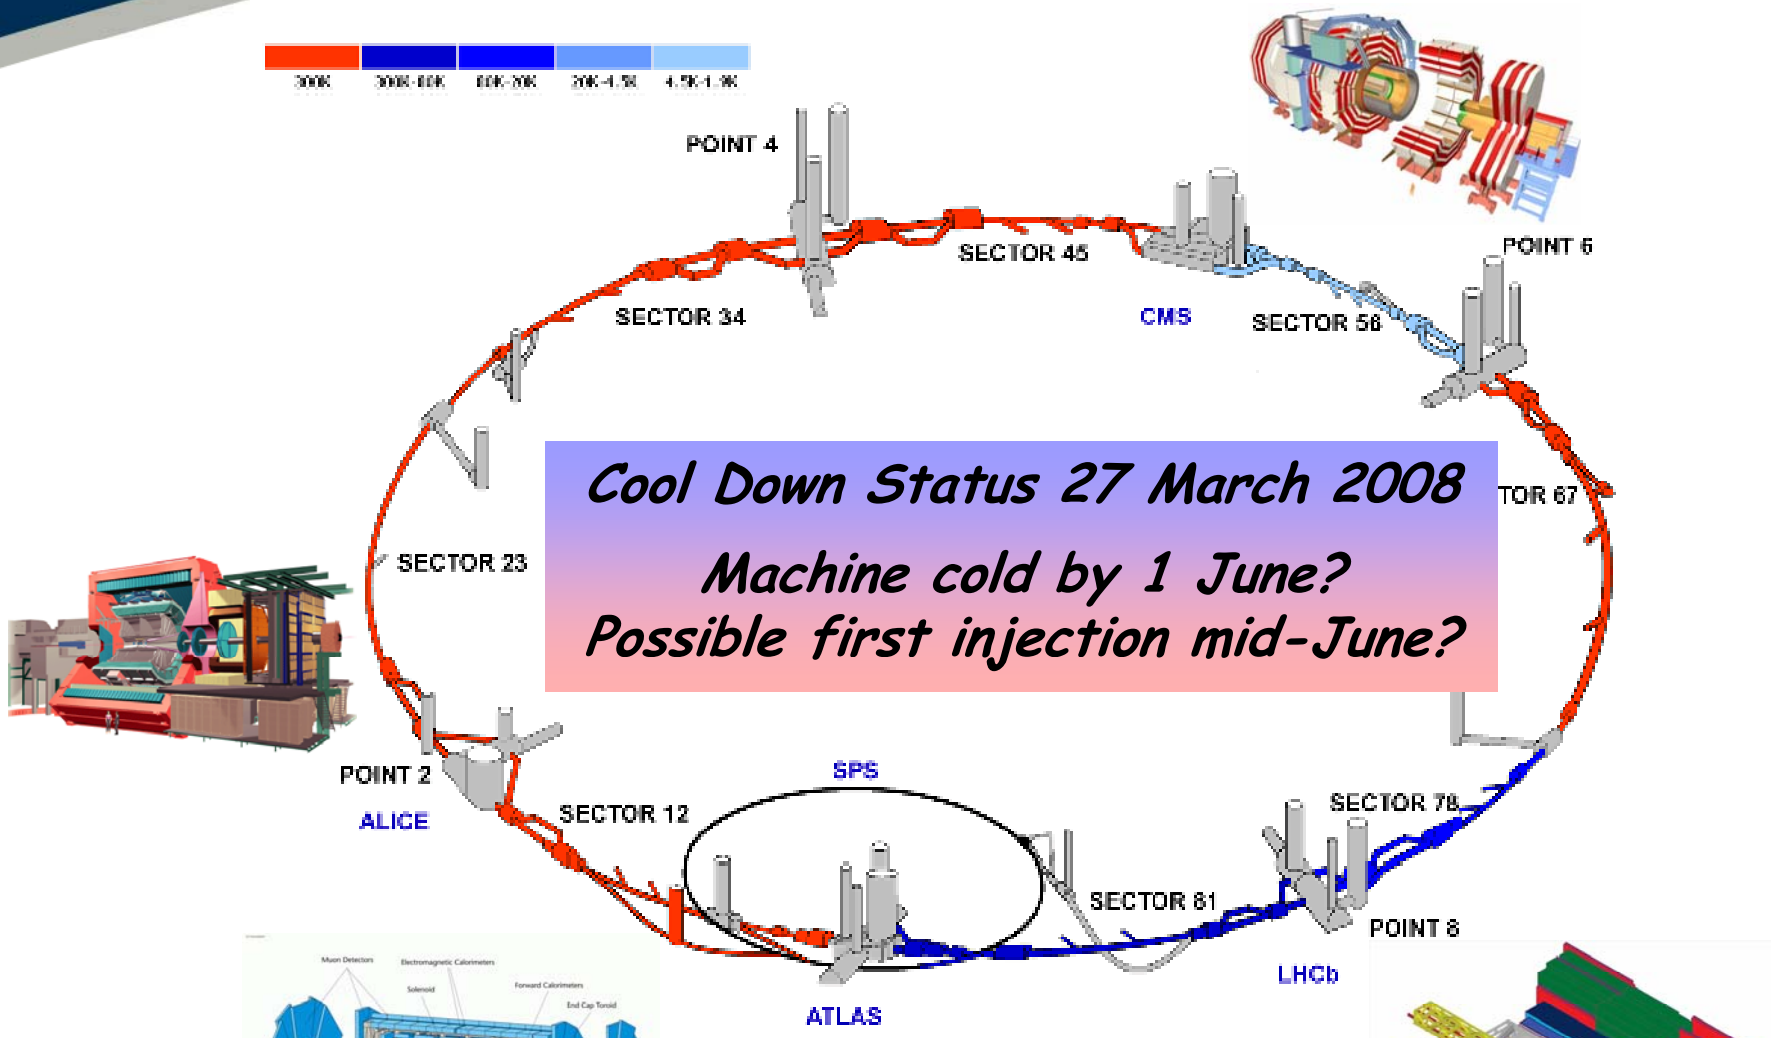

LHC Experiments - CCRC08

CCRC08 = *Common Computing Readiness Challenge*

Phase 1: 4 – 29 Feb 2008, Phase 2: 5 – 30 May 2008

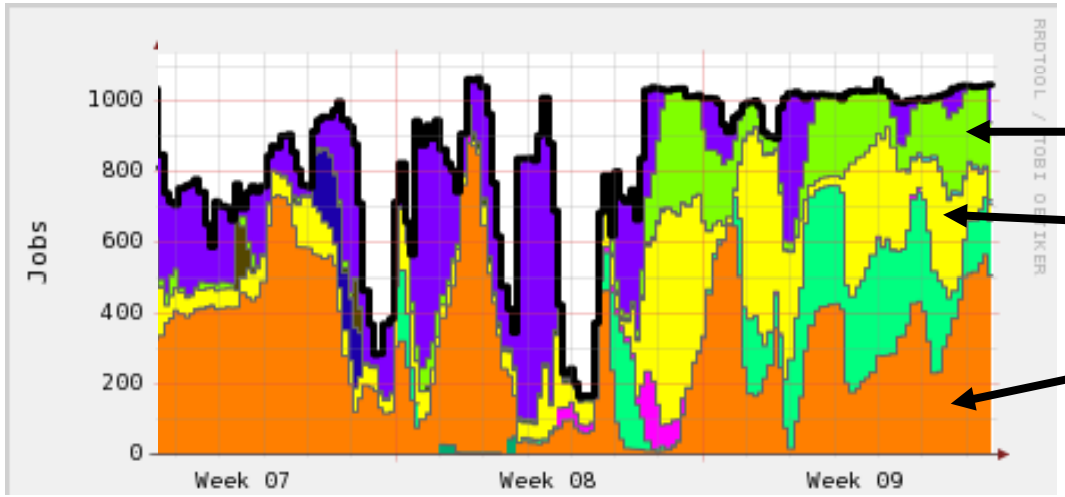

LHCb

CMS

ATLAS

UK Tier 1 Common running

All Sites Throughput

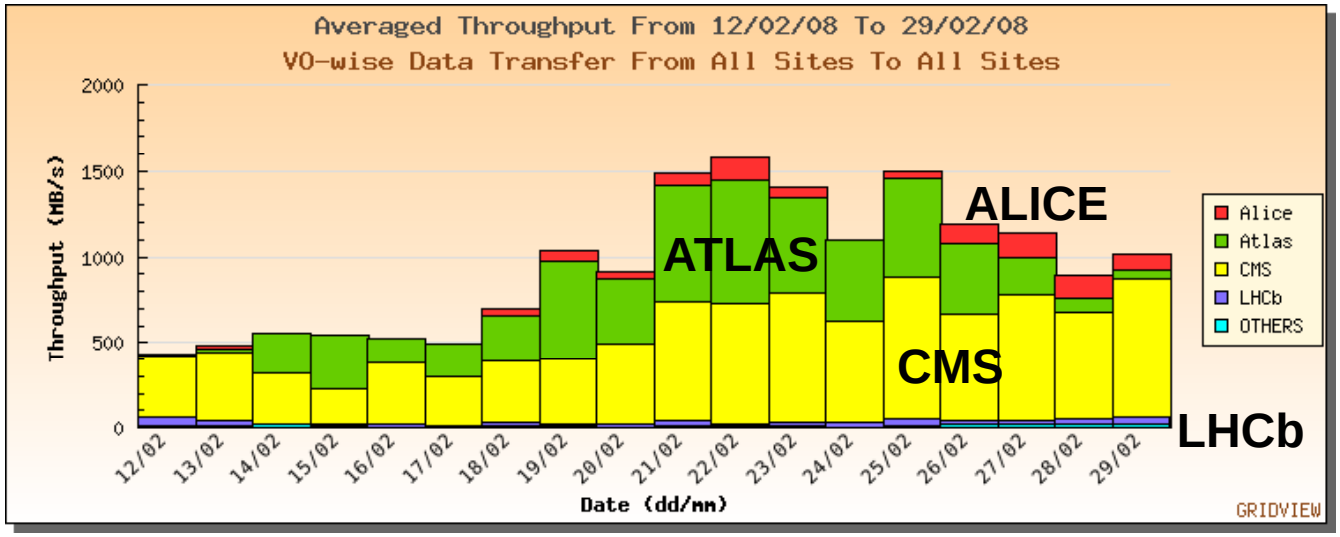

ATLAS 2008 Schedule

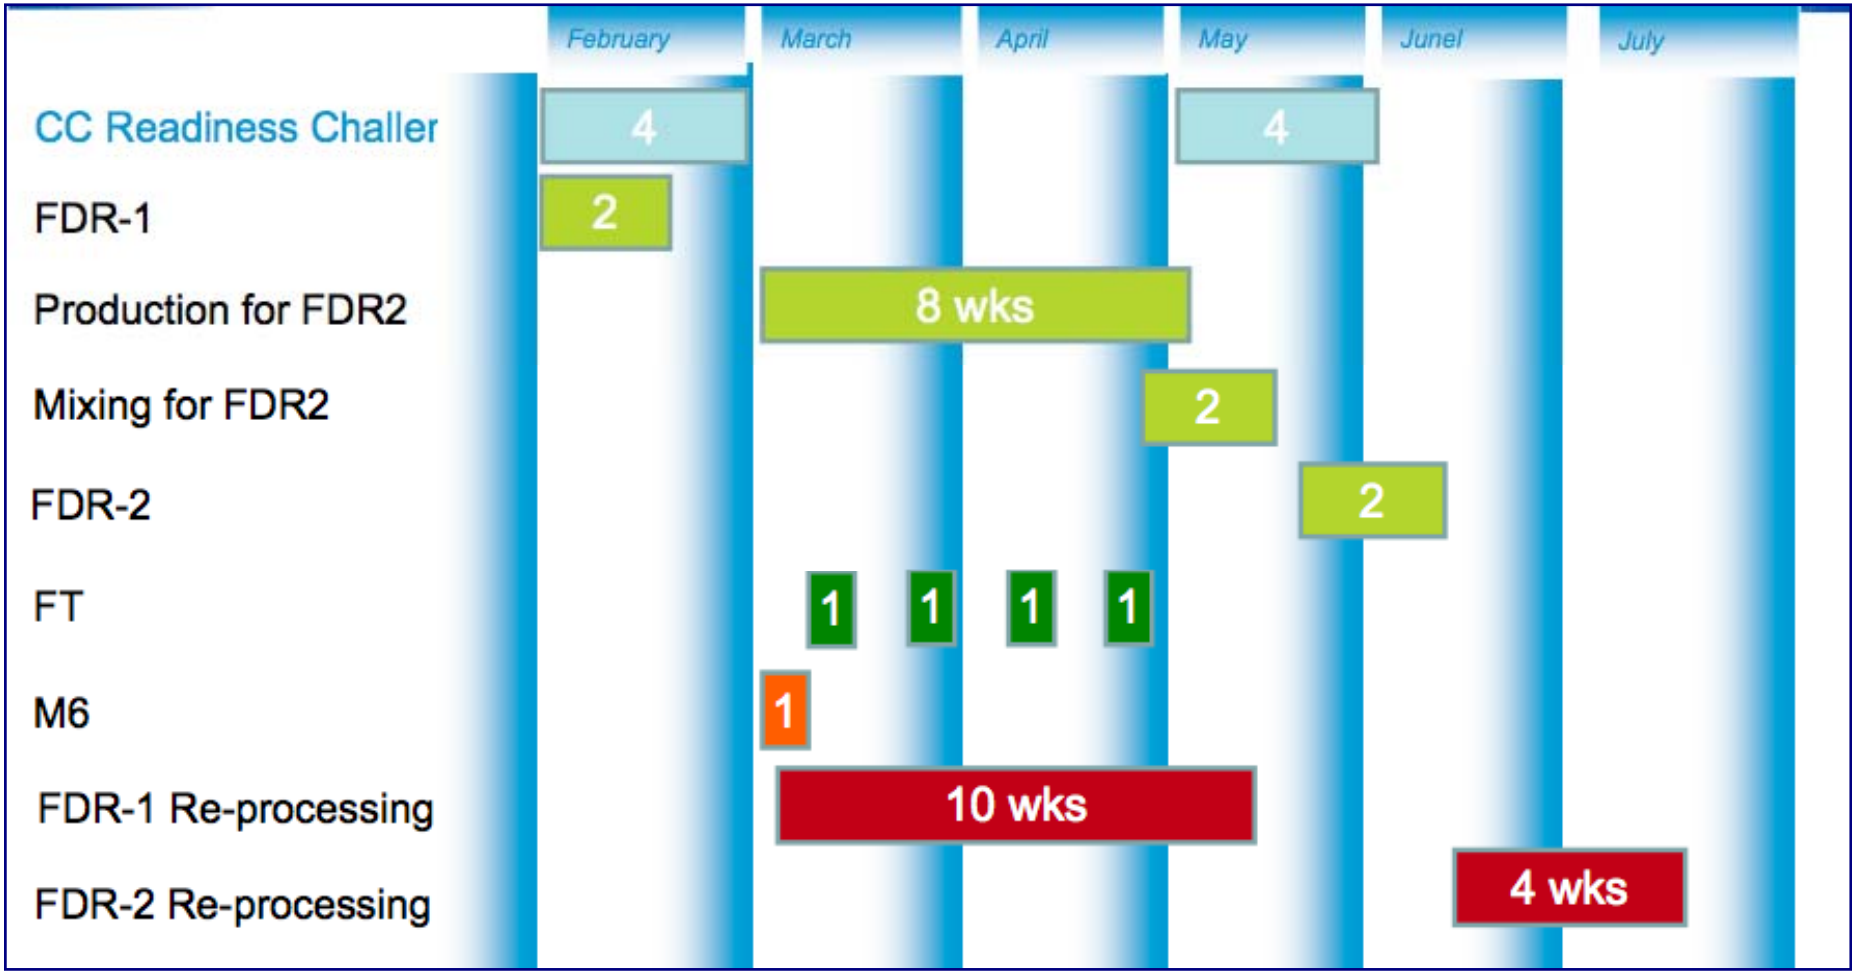

CMS Version

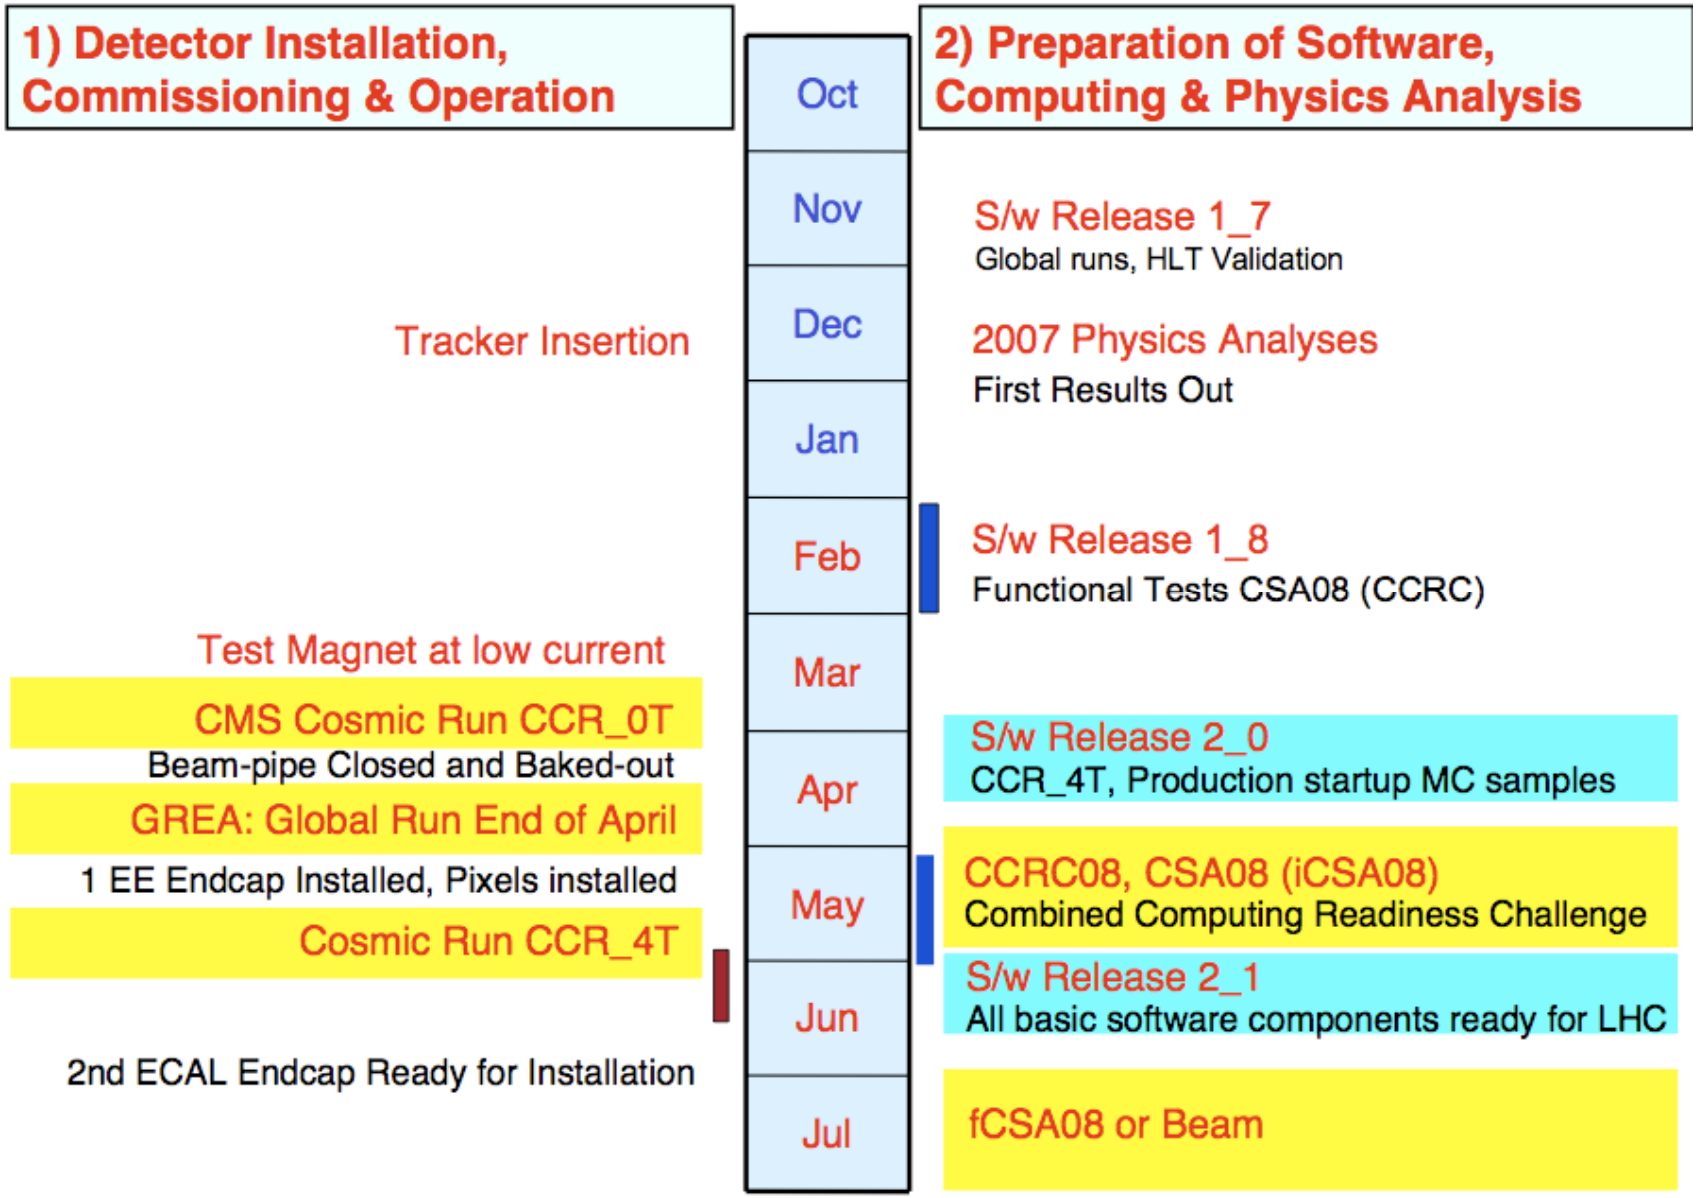
1) Detector Installation, Commissioning & Operation
2) Preparation of Software, Computing & Physics Analysis
Oct
Nov
Dec
Jan
Feb
Mar
Apr
May
Jun
Jul
S/w Release 1_7
Global runs, HLT Validation
Tracker Insertion
2007 Physics Analyses
First Results Out
S/w Release 1_8
Functional Tests CSA08 (CCRC)
Test Magnet at low current
CMS Cosmic Run CCR_0T
Beam-pipe Closed and Baked-out
S/w Release 2_0
CCR_4T, Production startup MC samples
GREA: Global Run End of April
1 EE Endcap Installed, Pixels installed
CCRC08, CSA08 (iCSA08)
Combined Computing Readiness Challenge
Cosmic Run CCR_4T
S/w Release 2_1
All basic software components ready for LHC
2nd ECAL Endcap Ready for Installation
fCSA08 or Beam

User Analysis - Ganga

Ganga – simplifies use of the Grid!

UK led project

- Ganga has built-in support for ATLAS and LHCb
- Component architecture allows customisation for other user groups

ack. K. Harrison

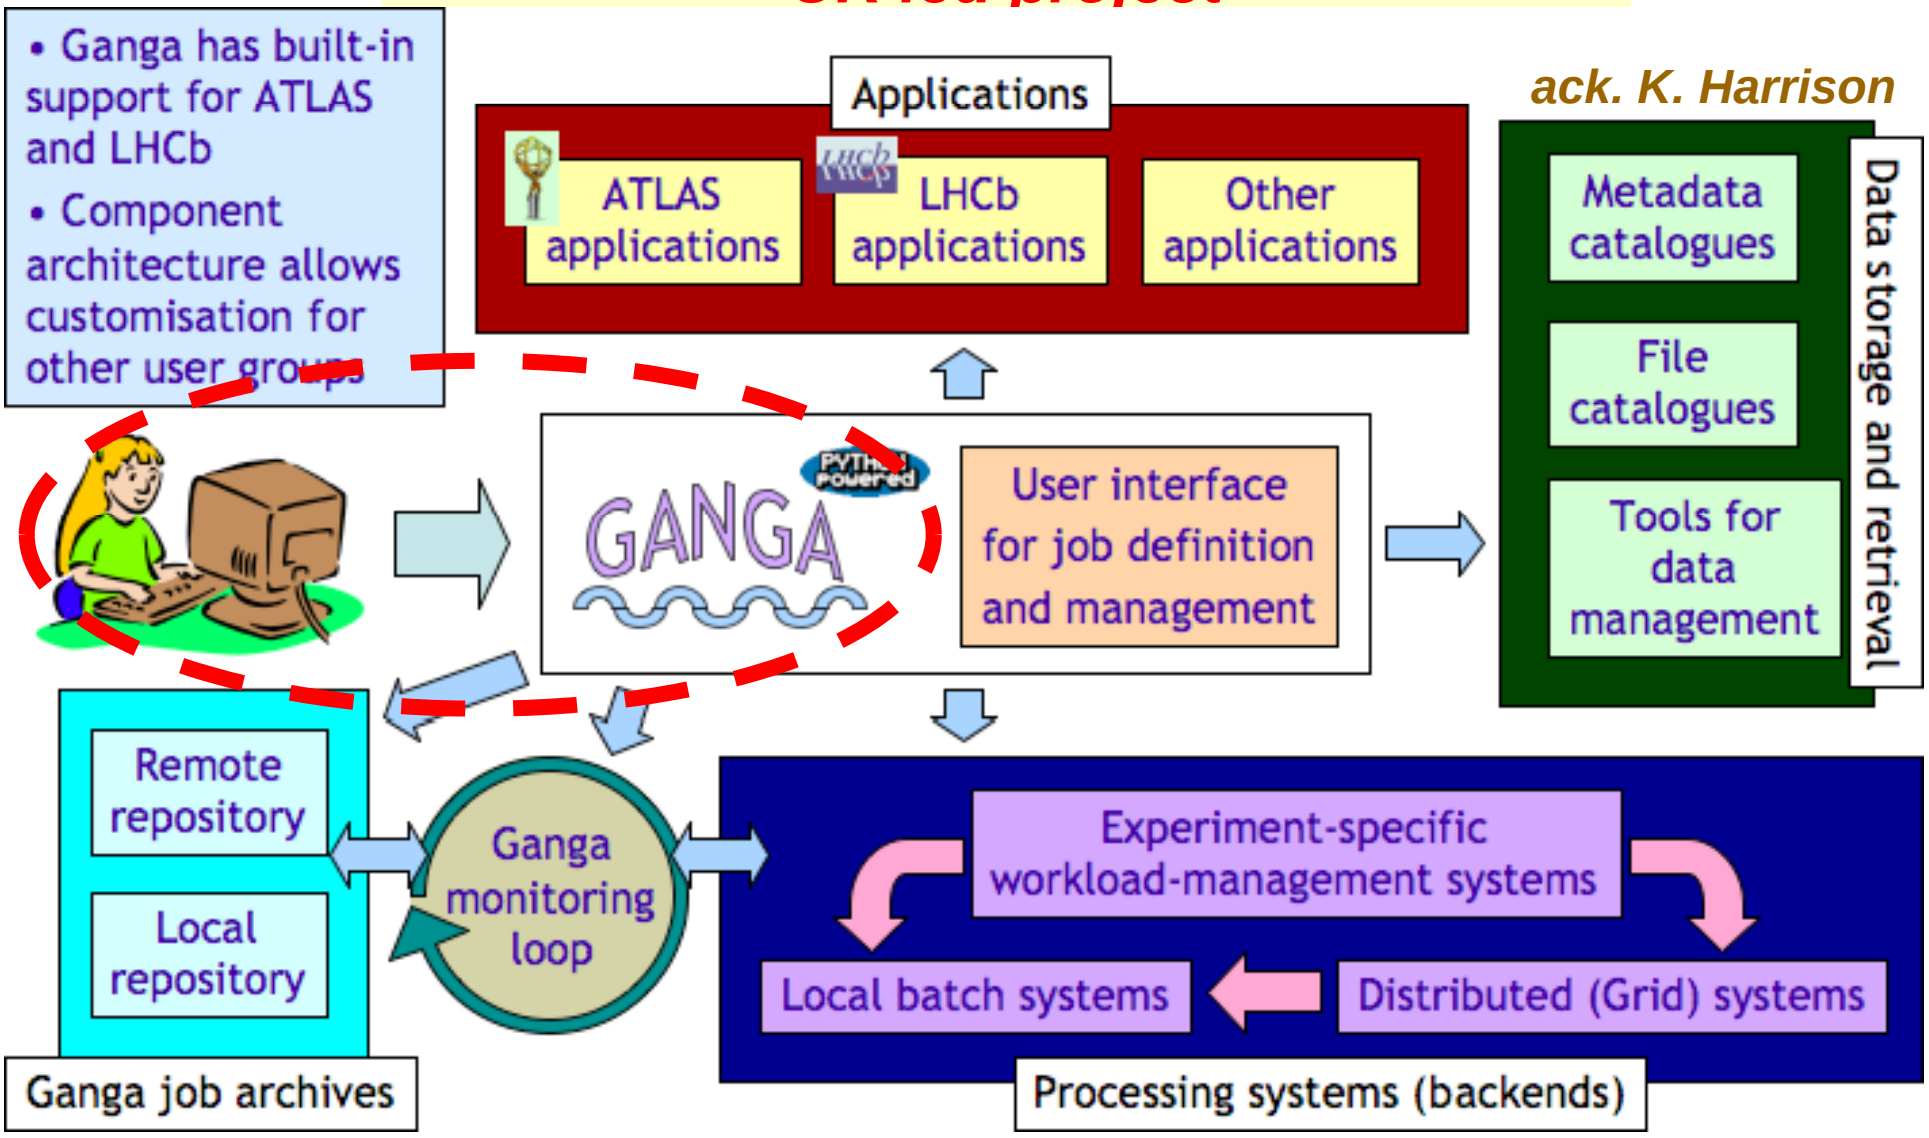

Conclusions

- **Great progress has been made in developing an operational Grid for the LHC experiments.**
- **Leading UK contribution - *thanks to a large number of people in T0/T1/T2 centres from 21 institutes and from the experiments.***
- **User analysis increasing & posing new challenges.**
- **GridPP3 started yesterday!**
- **Final ramp-up of resources for the start of LHC.**
- **Reliability/availability/efficiency needs to be improved over next few months - still a lot of work to do.**
- **Hype finally becomes reality? ….**

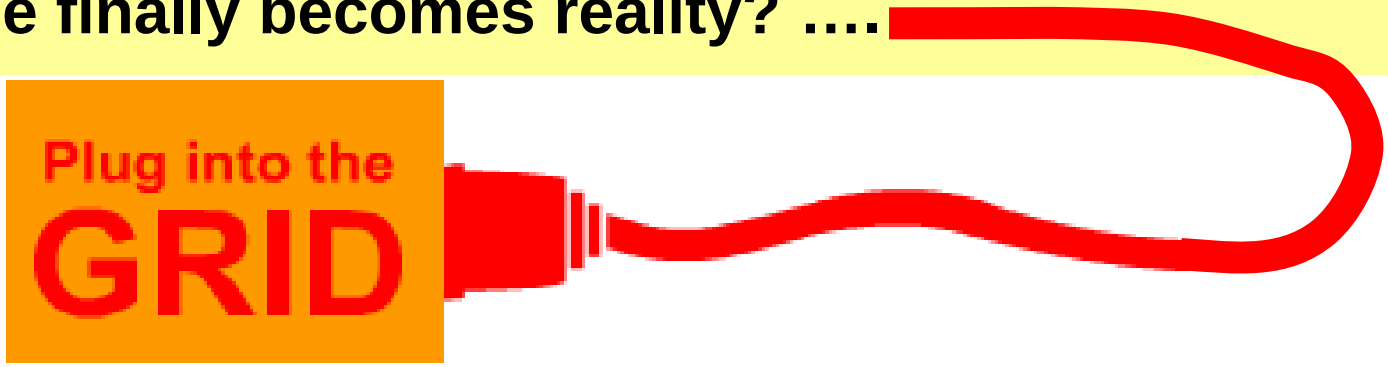

The End

...but also the start

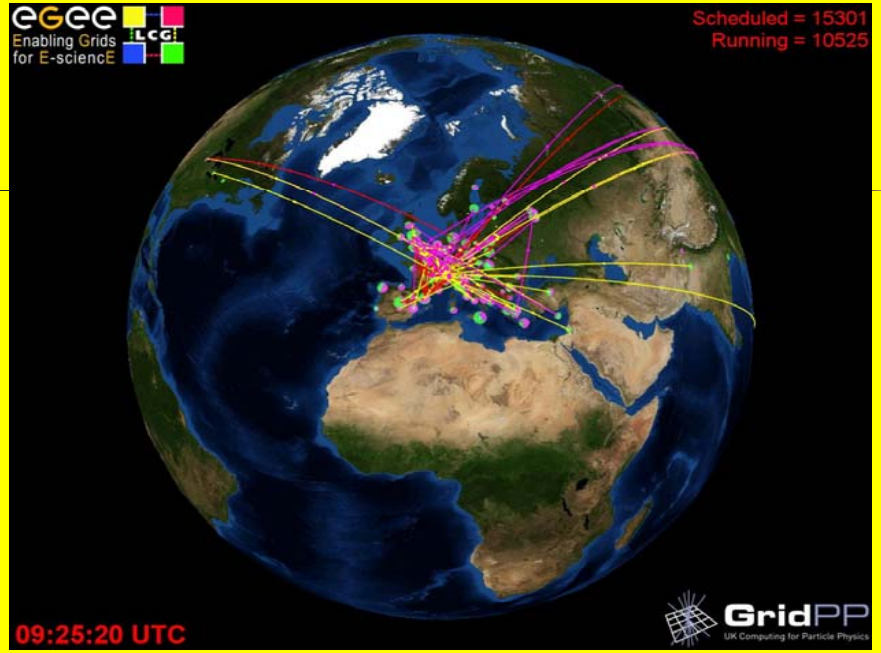

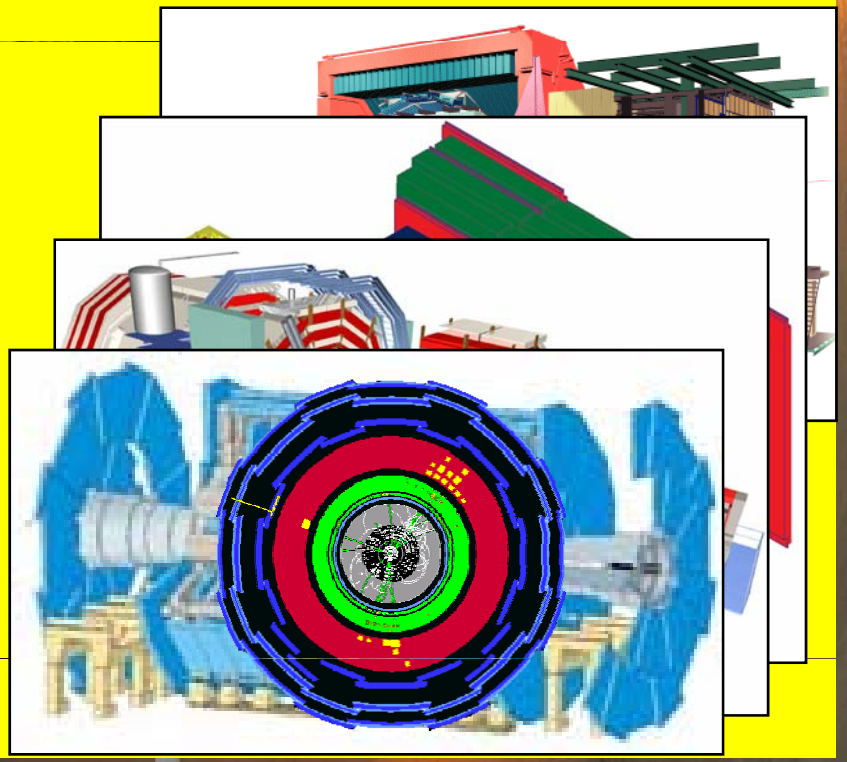

GridPP3