



**GridPP**  
UK Computing for Particle Physics

# GridPP Structures and Status Report

Jeremy Coles  
j.coles@rl.ac.uk

<http://www.gridpp.ac.uk>



- GridPP background

- GridPP components involved in Service Challenge preparations

- Sites involved in SC3

- SC3 deployment status

- Questions and issues

- Summary



**GridPP**  
UK Computing for Particle Physics

# GridPP background

## GridPP

GridPP - A UK Computing Grid for Particle Physics

19 UK Universities, CCLRC (RAL & Daresbury) and CERN

Funded by the Particle Physics and Astronomy Research Council (PPARC)

GridPP1 - Sept. 2001-2004  
£17m "From Web to Grid"

GridPP2 - Sept. 2004-2007  
£16(+1)m "From Prototype to Production"






# GridPP2 Components

A. Management, Travel, Operations

F. LHC Computing Grid Project (LCG Phase 2) [review]

B. Middleware Security Network Development

E. Tier-1/A Deployment: Hardware, System Management, Experiment Support



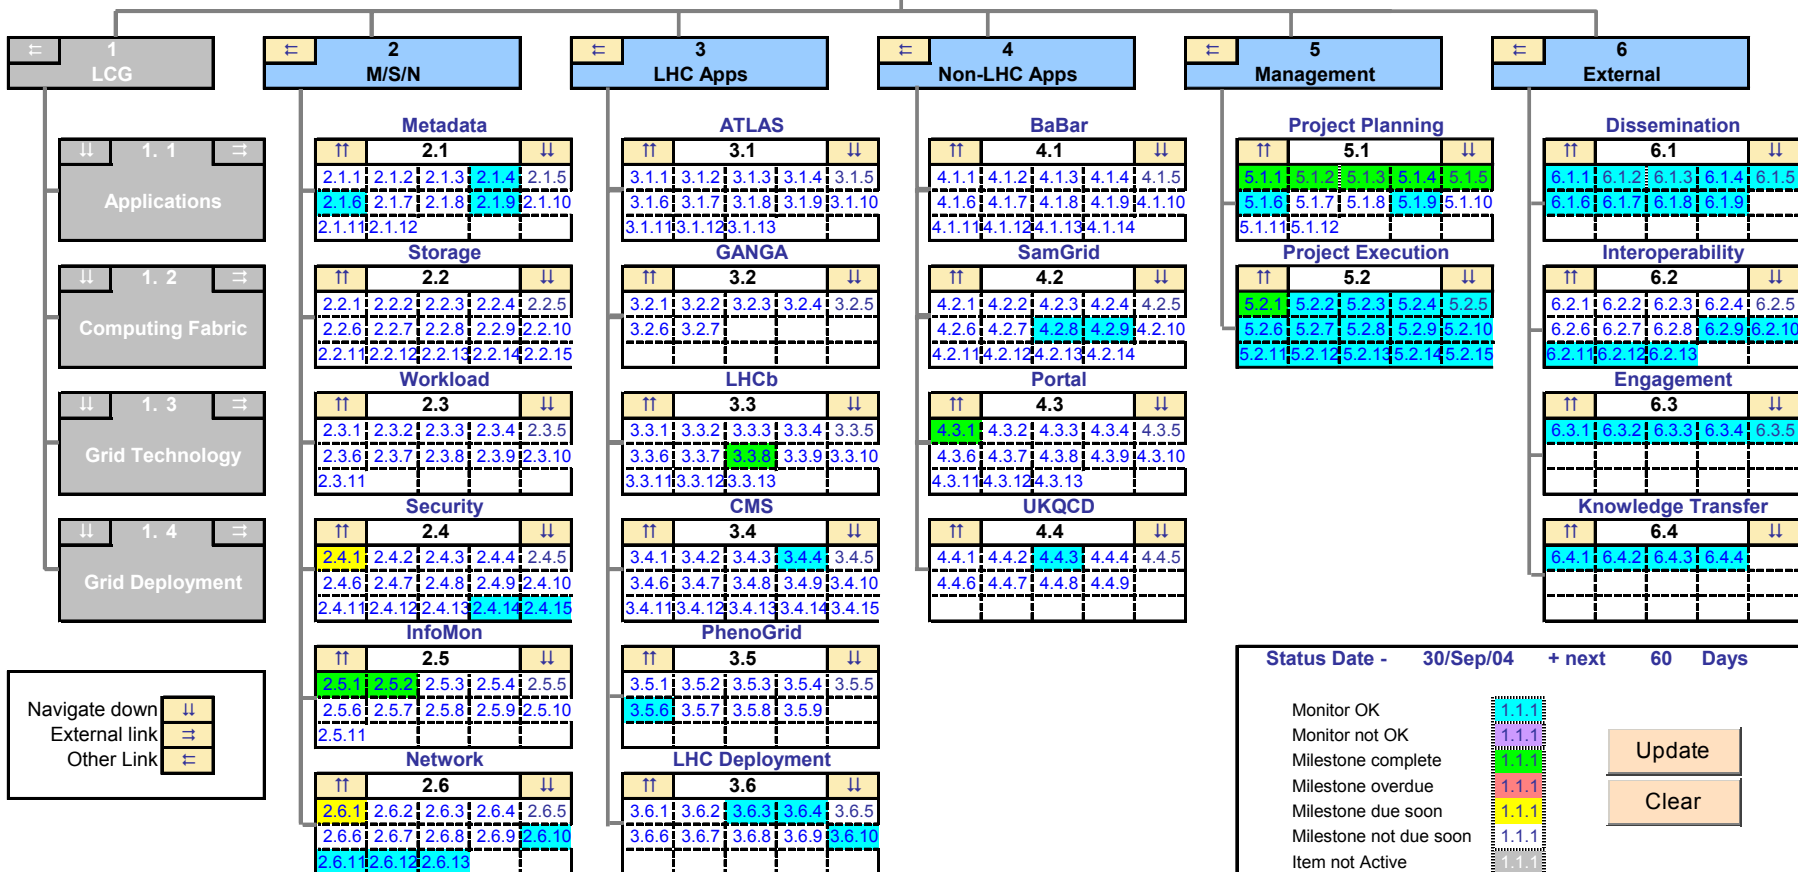
C. Grid Application Development LHC and US Experiments + Lattice QCD + Phenomenology

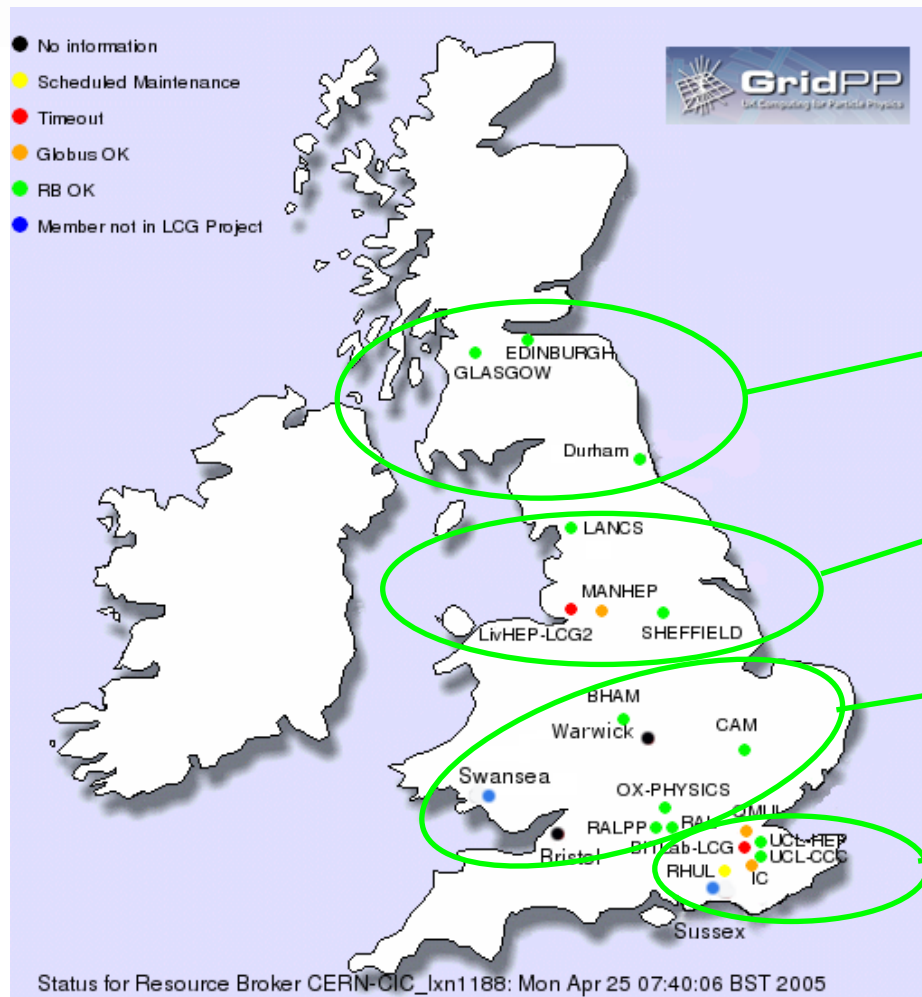
D. Tier-2 Deployment: 4 Regional Centres - M/S/N support and System Management



**GridPP2 Goal: To develop and deploy a large scale production quality grid in the UK for the use of the Particle Physics community**

Production Grid Milestones																	Production Grid Metrics																
0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	0.10	0.11	0.12	0.13	0.14	0.15	0.16	0.17	0.100	0.101	0.102	0.103	0.104	0.105	0.106	0.107	0.108	0.109	0.110	0.111	0.112	0.113	0.114	0.115	0.116
0.18	0.19	0.20	0.21	0.22	0.23	0.24	0.25	0.26	0.27	0.28	0.29	0.30	0.31	0.32	0.33	0.34	0.117	0.118	0.119	0.120	0.121	0.122	0.123	0.124	0.125	0.126	0.127	0.128	0.129	0.130	0.131	0.132	0.133
0.35	0.36	0.37	0.38	0.39	0.40	0.41	0.42	0.43	0.44	0.45	0.46	0.47	0.48	0.49	0.50	0.51	0.134	0.135	0.136	0.137	0.138	0.139	0.140	0.141	0.142	0.143	0.144	0.145	0.146	0.147	0.148	0.149	0.150
0.52	0.53	0.54	0.55	0.56	0.57												0.151	0.152	0.153	0.154	0.155	0.156											





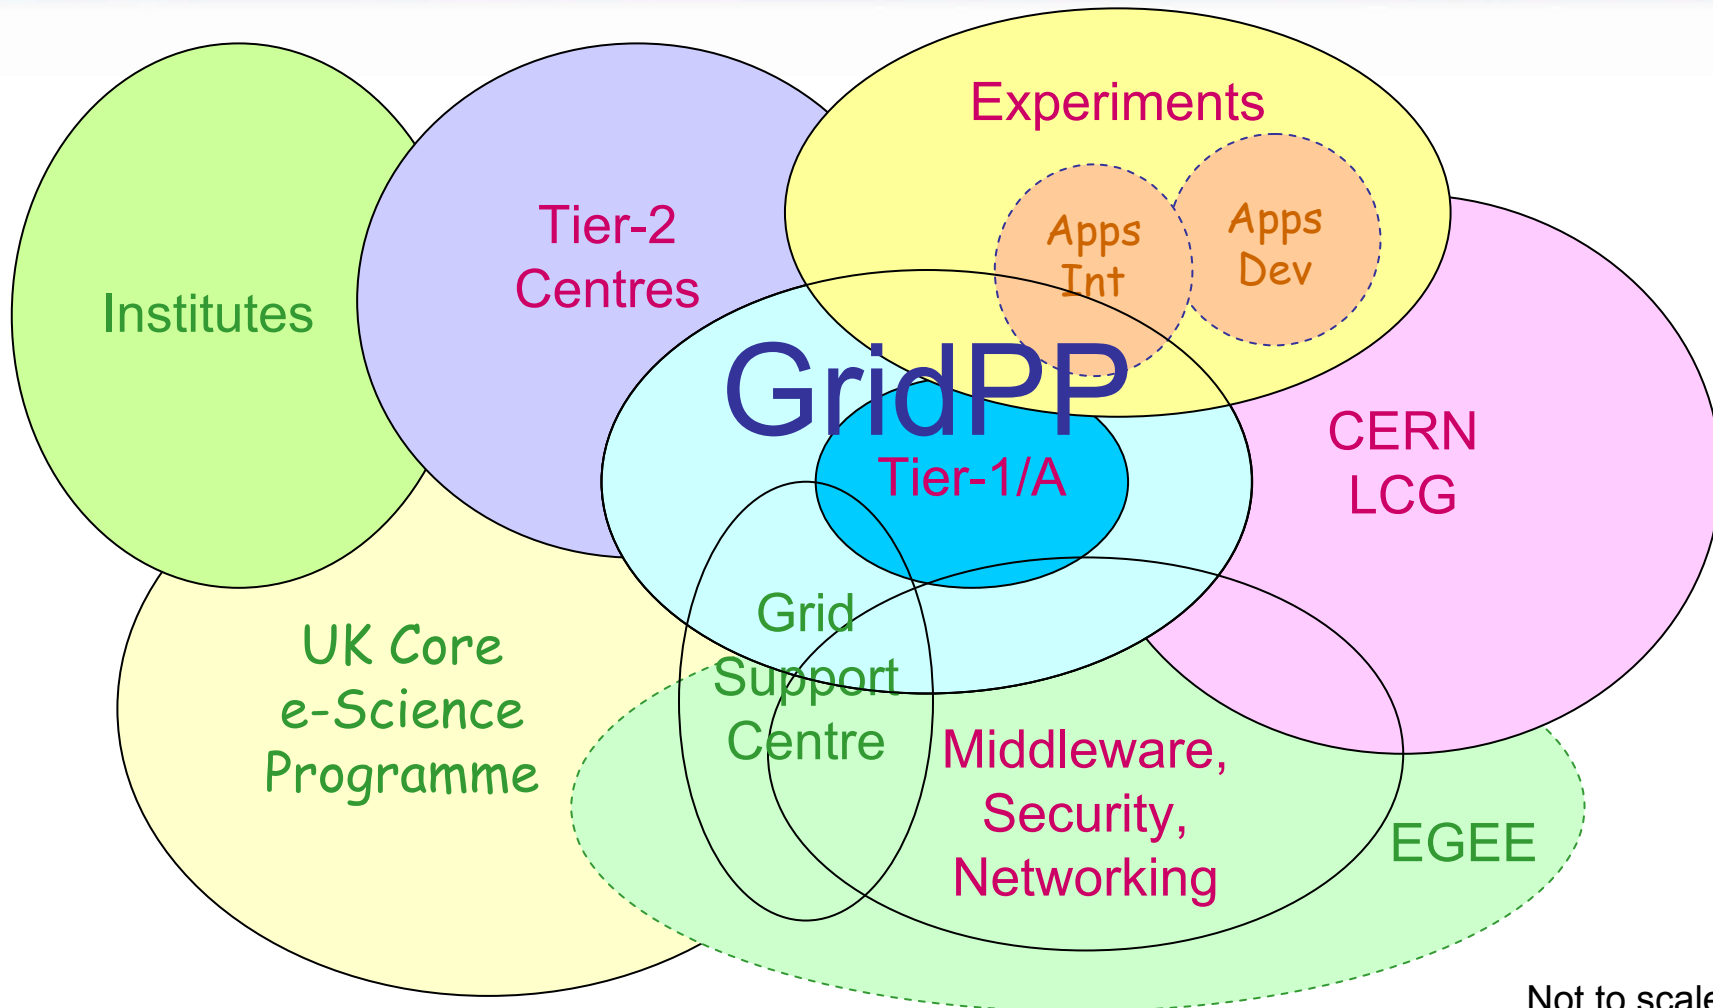
Tier-2s

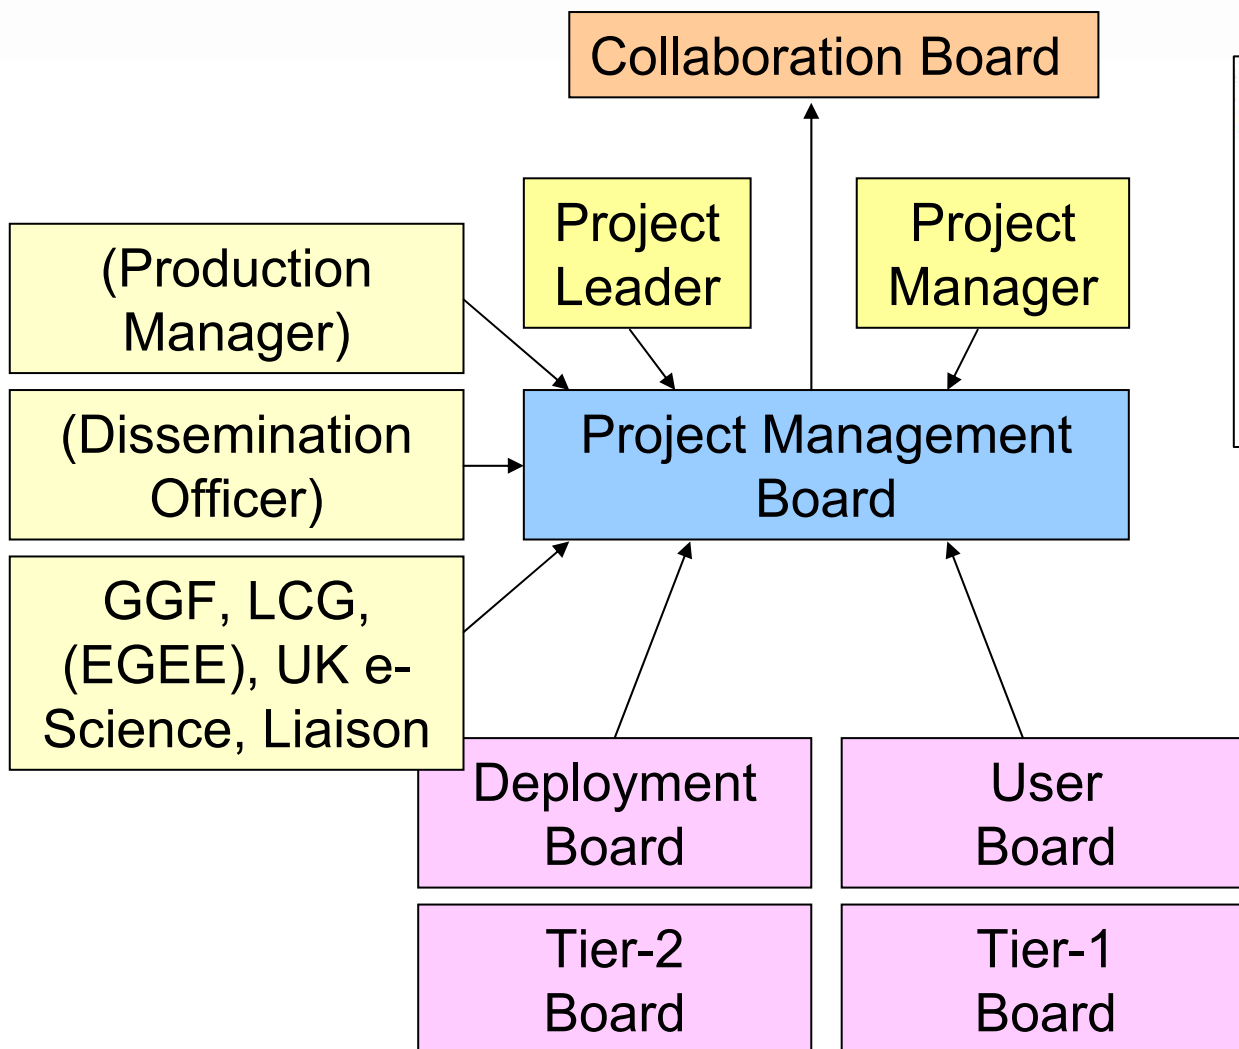
ScotGrid

NorthGrid

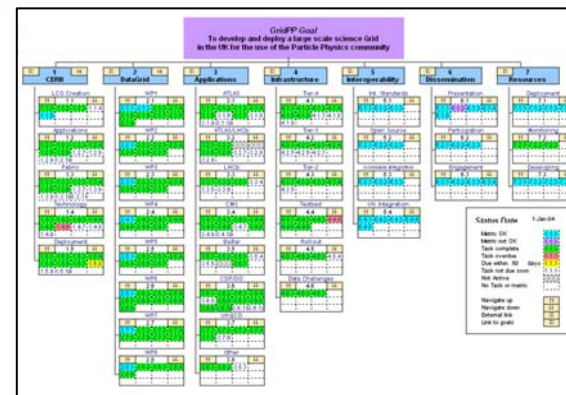
SouthGrid

London Tier-2





## Project Map



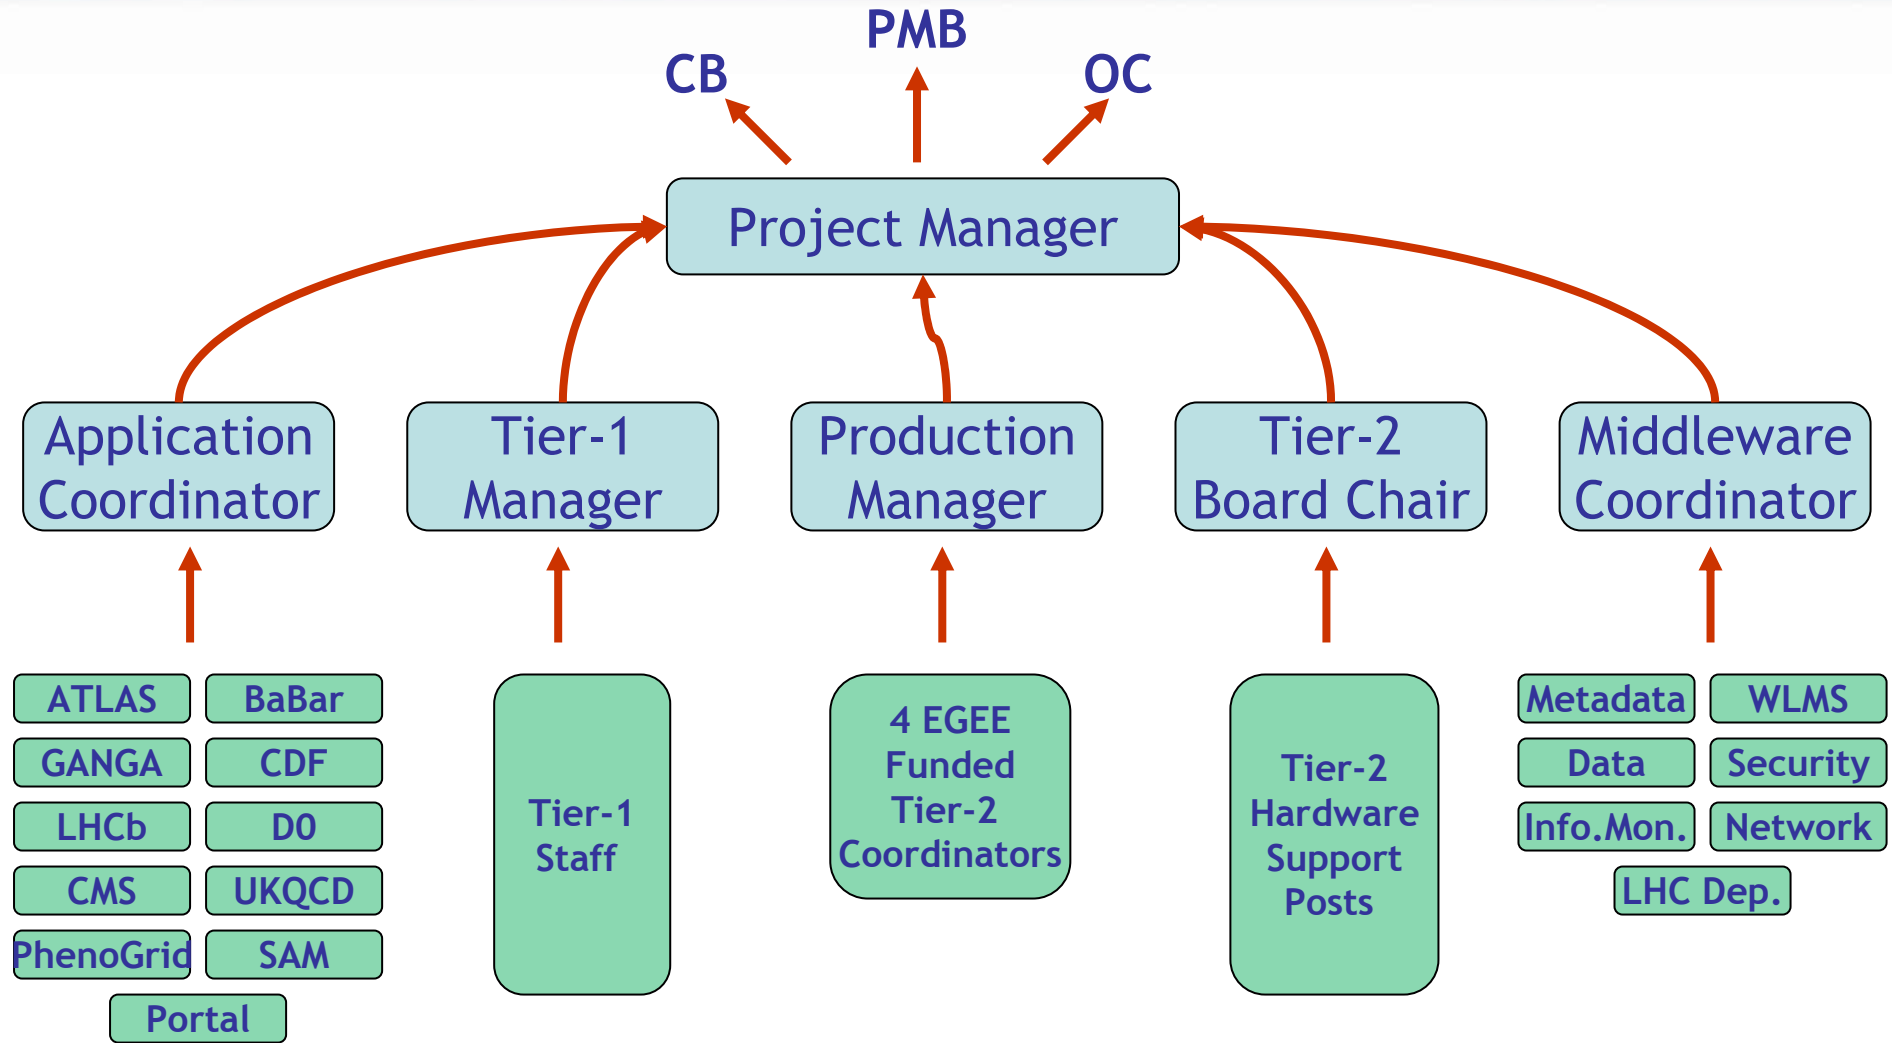
## Risk Register

ID	Name	GridPP		LCG		EDG		Agas		Infrastr.		Interop.		
		(1)M	Risk	(1)M	Risk	(1)M	Risk	(1)M	Risk	(1)M	Risk	(1)M	Risk	
R1	Recruitment/retention difficulties			2	2	1	1	4	2	1	1	3	3	
R2	Sudden loss of key staff			1	4	1	3	3	1	3	1	3	3	
R3	Minimal Contingency			1	2	2	2	4	2	4	2	2	2	
R4	GridPP deliverables late			1	3	3	2	4	2	3	2	2	2	
R5	Sub-components not delivered to project	2	4	2	2	1	2	4	2	3	4	2	2	
R6	Non take-up of project results			2	1	2	3	4	1	4	4	4	4	
R7	Change in project scope			1	1	1	1	1	1	1	1	1	1	
R8	Bad publicity	2	2	4	1	3	1	3	3	3	3	3	3	
R9	No publicity	2	1	2	2	2	2	2	2	2	2	2	2	
R10	External software dependence	1	3	3	3	3	1	4	4	2	4	4	4	
R11	Lack of monitoring of staff			1	2	2	2	4	1	3	1	2	2	
R12	Withdrawal of an experiment	1	4	4	1	4	1	2	2	2	2	2	2	
R13	Lack of cooperation between Tier centres			1	2	2	2	2	2	2	2	2	2	
R14	Scalability problems			1	3	2	2	2	2	2	2	2	2	
R15	Software maintainability problems			1	3	2	2	2	2	2	2	2	2	
R16	Technology shifts	1	3	1	2	4	1	3	3	3	3	3	3	
R17	Evolution of research	1	2	2	1	1	3	3	3	3	3	3	3	
R18	Lack of funding to meet LCG PH-1 goals			1	1	1	1	1	1	1	1	1	1	
R19	Adequate persistency solution not ready							1	3	3	3	3	3	
R20	Conflicting software requirements							4	1	4	2	2	4	
R21	Tier-A hardware fails to meet requirements							1	3	3	3	3	3	
R22	Other Hardware fails to meet requirements							1	1	1	1	1	1	
R23	Hardware physical risk (large scale)							1	4	4	4	4	4	
R24	Hardware physical risk (small scale)							2	2	2	2	2	2	
R25	Hardware procurement problems							2	2	2	2	2	2	
R26	LAN Bottlenecks							2	1	2	2	2	2	
R27	Tier-2 organisation fails							2	2	2	2	2	2	
R28	Tier-2 hardware not used as planned							2	1	2	2	2	2	
R29	SYSMAN effort inadequate							2	2	2	2	2	2	
R30	Firewalls interfere with Grid							1	1	1	1	1	1	
R31	Inability to establish trust relationships							2	2	2	2	2	2	
R32	Security inadequate to operate Grid							1	3	3	3	3	3	
R33	ECDF does not establish standards											1	3	3
R34	Minimal open source code development											2	2	4
R35	Failure of international cooperation							1	1	4	4	4	4	
R36	e-Science and GridPP divergence											2	2	4
R37	Institutes do not embrace Grid							1	3	3	3	3	3	3
R38	Grid is not stable enough for use							1	3	3	3	3	3	3
R39	Delay of the LHC							2	2	4	4	4	4	4
R40	Lack of future funding							1	4	4	4	4	4	4
R41	Network backbone failure							1	4	4	4	4	4	4
R42	Network backbone bottleneck							1	2	2	2	2	2	2





# Reporting Lines





- 1. LCG “expert”
- 2. dCache deployment
- 3. Hardware support

- 1. Hardware support post
- 2. Tier-2 coordinator (LCG)

- 1. Network advice/support
- 2. SRM deployment advice/support

Tier-1  
Manager

Production  
Manager

Tier-1  
Staff

4 EGEE  
Funded  
Tier-2  
Coordinators

Tier-2  
Hardware  
Support  
Posts

Network group

Storage group

# Activities to help deployment

- RAL storage workshop - review of dCache model, deployment and issues
- Biweekly network teleconferences
- Biweekly storage group teleconferences. GridPP storage group members available to visit sites.
- UK SC3 teleconferences - to cover site problems and updates
- SC3 “tutorial” last week for grounding in FTS, LFC and review of DPM and dCache

Agendas for most of these can be found here:

<http://agenda.cern.ch/displayLevel.php?fid=338>



# Sites involved in SC3

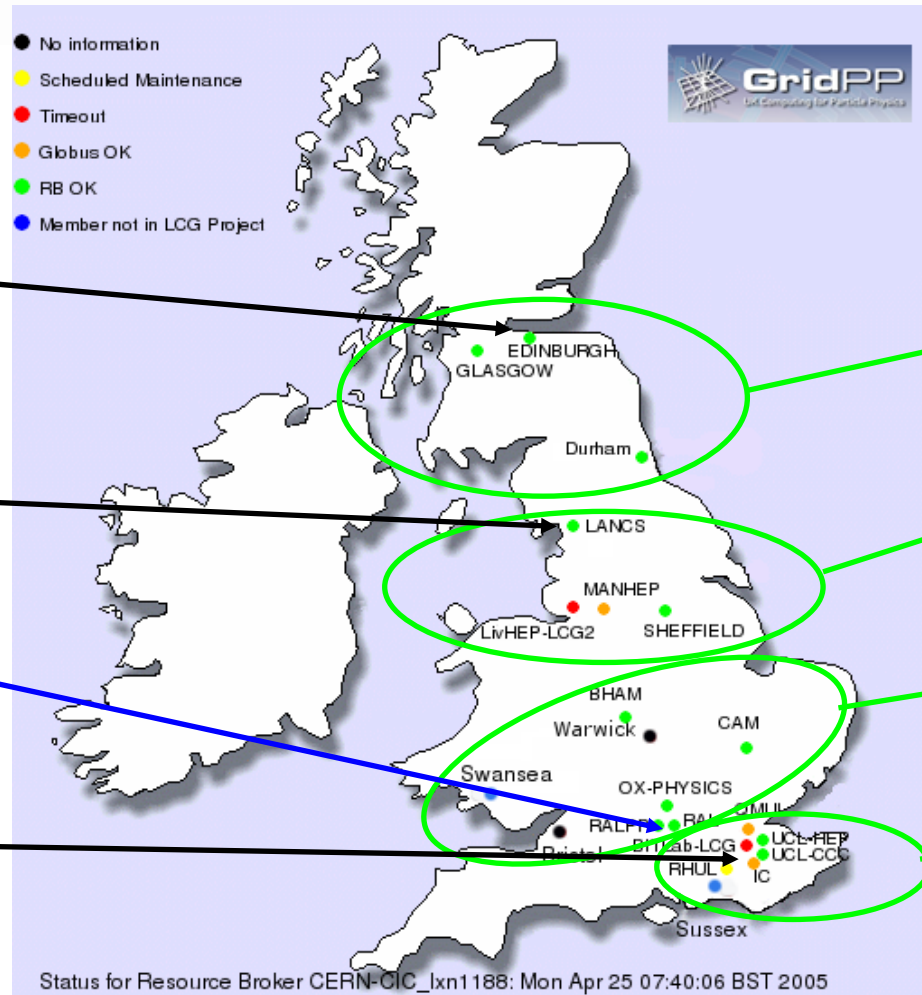
## SC3 sites

LHCb

ATLAS

Tier-1

CMS



## Tier-2s

ScotGrid

NorthGrid

SouthGrid

London Tier-2



## Connectivity

- RAL to CERN 2\*1 Gbits/s via UKLight and Netherlight
- Lancaster to RAL 1\*1 Gbits/s via UKLight  
(fallback: SuperJanet 4 production network)
- Imperial to RAL SuperJanet 4 production network
- Edinburgh to RAL SuperJanet 4 production network

## Status

- 1) Lancaster to RAL UKLight connection requested 6<sup>th</sup> May 2005
- 2) UKLight access to Lancaster available now
- 3) Additional 2\*1 Gbit/s interface card required at RAL

## Issues

- Awaiting timescale from UKERNA on 1 & 3
- IP level configuration discussions (Lancaster-RAL) just started
- Merging SC3 production traffic and UKLight traffic raises several new problems
- Underlying network provision WILL change between now and LHC start-up



- Currently merging SC3 dCache with production dCache
  - Not yet clear how to extend production dCache to allow good bandwidth to UKLIGHT, the farms and SJ4
  - dCache expected to be available for Tier-2 site tests around 3<sup>rd</sup> June
- Had problems with dual attachment of SC2 dCache. A fix has been released but we have not yet tried it.
  - Implications for site network setup
- CERN link undergoing “health check” this week. During SC2 the performance of the link was not great
- Need to work on scheduling (assistance/attention) transfer tests with T2 sites. Tests should complete by 17<sup>th</sup> June.
- Still unclear on exactly what services need to be deployed - there is increasing concern that we will not be able to deploy in time!



<b>Date</b>	<b>Status</b>
Friday 20 <sup>th</sup> May	Network provisioned in final configuration
Friday 3 <sup>rd</sup> June	Network confirmed to be stable and able to deliver at full capacity. SRM ready for use
Friday 17 <sup>th</sup> June	Tier-1 SRM tested end-to-end with CERN. Full data rate capability demonstrated. Load balancing and tuning completed to Tier-1 satisfaction
Friday 1 <sup>st</sup> July	Completed integration tests with CERN using FTS. Certification complete. Tier-1 ready for SC3



## HARDWARE

- GridPP frontend machines upgraded to 2.4.0. **Limited CPU available (3 machines - 5CPUs). Still waiting for more details about requirements from LHCb**
- 22TB datastore deployed but a **raid array problem leads to half the storage being currently unavailable - IBM investigating**
- Disk server: IBM xSeries 440 with eight 1.9 GHz Xeon processors, 32Gb RAM
- GB copper Ethernet to GB 3com switch

## SOFTWARE

- dCache head and pool nodes rebuilt with Scientific Linux 3.0.4. dCache now installed on admin node using apt-get. **Partial install on pool node.** On advice will restart using YAIM method.
- **What software does LHCb need to be installed?**





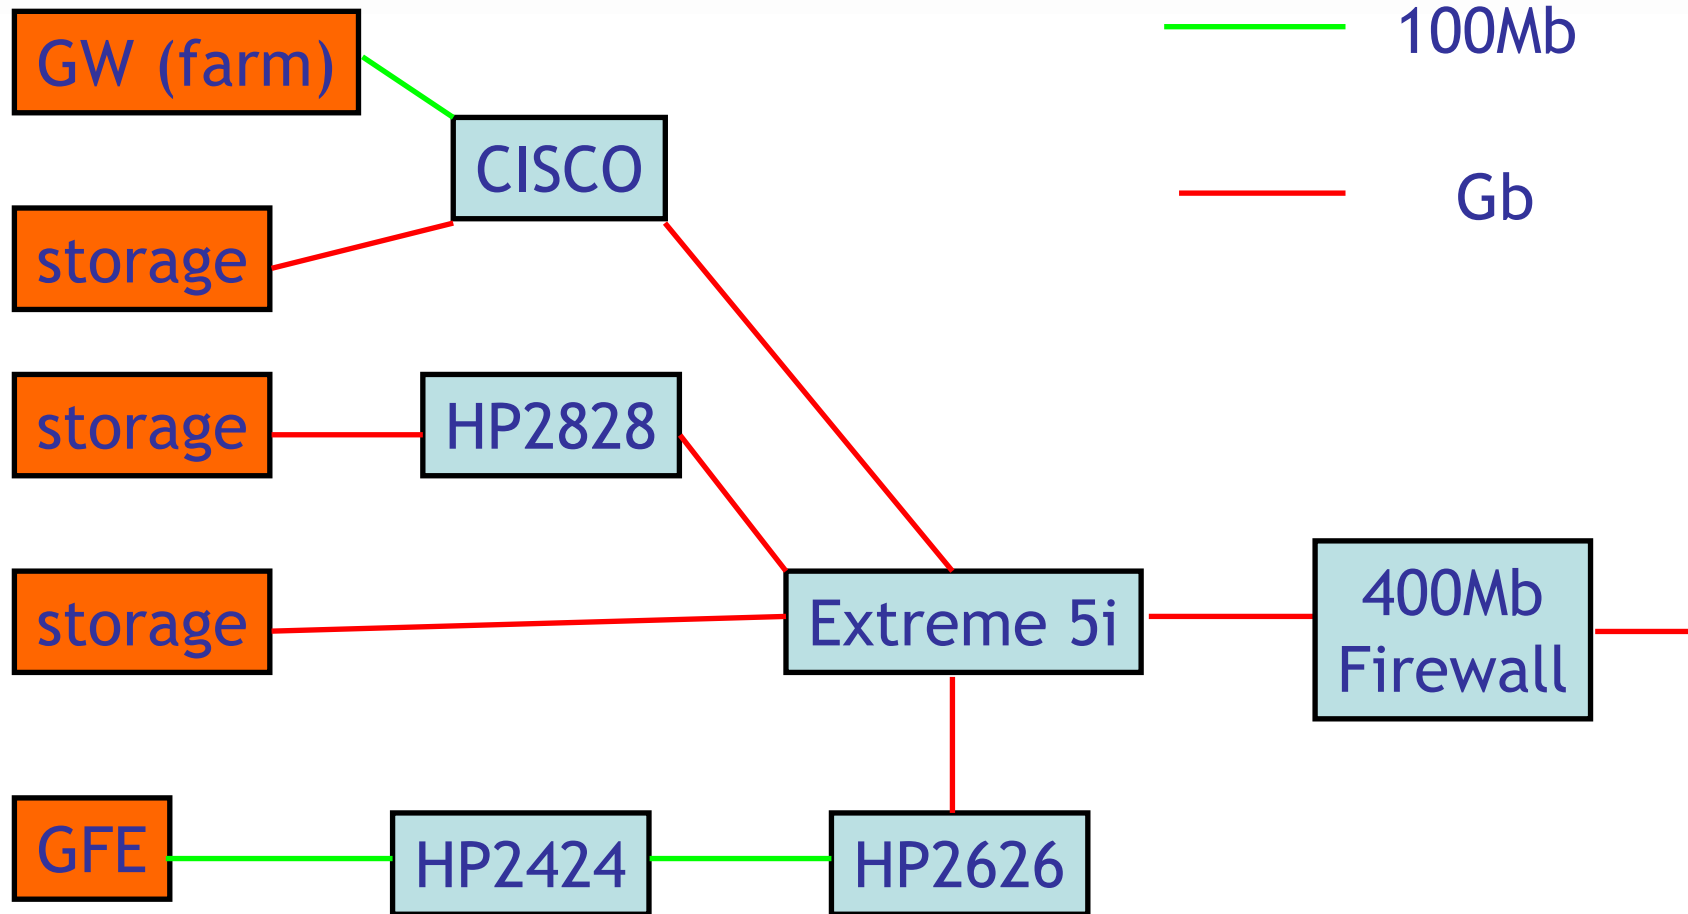
## HARDWARE

- 1.5TB CMS dedicated storage available for SC3
- 1.5TB (shared) shared storage can be made partially available
- May purchase more if needed - **what is required!?**
- Farm on 100Mb connection. 1Gb connection to all servers.
  - 400Mb firewall may be a bottleneck.
  - Two separate firewall boxes available
  - Could upgrade to 1Gb firewall - **what is needed?**



## SOFTWARE

- dCache SRM now installed on test nodes
  - 2 main problems both related to installation scripts
- This week starting installation on a dedicated production server
  - Issues: Which ports need to be open through the firewall, use of pools as VO filesystem quotas and optimisation of local network topology
- CMS software
  - Phedex deployed, but not tested
  - PubDB in process of deployment
  - Some reconfiguration of farm was needed to install CMS analysis software
  - What else is necessary? **More details from CMS needed!**





## HARDWARE

- Farm now at LCG 2.4.0.
- 6 I/O servers with 2x6TB RAID5 arrays each currently has 100 Mb/s connectivity
- Request submitted to UKERNA for dedicated lightpath to RAL. Currently expect to use production network.
- T2 Designing system so that no changes needed for transition from throughput phase to service phase.
  - Except possible IPerf and/or transfers of “fake data” as backup of connection and bandwidth tests in throughput phase

## SOFTWARE

- dCache deployed onto test machines. Current plans have production dCache available in early July!
- **Currently evaluating what ATLAS software and services needed overall** and their deployment
- Waiting for FTS release.



Task	Start Date	End Date	Resource
Optimisation of network	Mon 18/09/06	Fri 13/10/06	BD,MD,BG,NP, <b>RAL</b>
Test of data transfer rates (CERN-UK)	Mon 16/10/06	Fri 01/12/06	BD,MD, <b>ATLAS, RAL</b>
<b>Provision of end-to-end conn. (T1-T2)</b>			
Integrate Don Quixote tools into SC infrastructure at LAN	Mon 19/09/05	Fri 30/09/05	BD
Provision of memory-to-memory conn. (RAL-LAN)	Tue 29/03/05	Fri 13/05/05	UKERNA,BD,BG,NP, <b>RAL</b>
Provision and Commission of LAN h/w	Tue 29/03/05	Fri 10/06/05	BD,BG,NP
Installation of LAN dCache SRM	Mon 13/06/05	Fri 01/07/05	MD,BD
Test basic data movement (RAL-LAN)	Mon 04/07/05	Fri 29/07/05	BD,MD, <b>ATLAS, RAL</b>
Review of bottlenecks and required actions	Mon 01/08/05	Fri 16/09/05	BD
[SC3 – Service Phase]			
Review of bottlenecks and required actions	Mon 21/11/05	Fri 27/01/06	BD
Optimisation of network	Mon 30/01/06	Fri 31/03/06	BD,MD,BG,NP
Test of data transfer rates (RAL-LAN)	Mon 03/04/06	Fri 28/04/06	BD,MD,
Review of bottlenecks and required actions	Mon 01/05/06	Fri 26/05/06	BD



# What we know!

Task	Start Date	End Date
Tier-1 network in final configuration	Started	20/5/05
Tier-1 Network confirmed ready. SRM ready.	Started	1/6/05
Tier-1 – test dCache dual attachment?		
Production dCache installation at Edinburgh	Started	30/05/05
Production dCache installation at Lancaster		TBC
Production dCache installation at Imperial	Started	30/05/05
T1-T2 dCache testing		17/06/05
Tier-1 SRM tested end-to-end with CERN. Load balanced and tuned.		17/06/05
Install local LFC at T1		15/06/05
Install local LFC at T2s		15/06/05
FTS server and client installed at T1	ASAP?	??
FTS client installed at Lancaster, Edinburgh and Imperial	ASAP?	??



# What we know!

Task	Start Date	End Date
CERN-T1-T2 tests	24/07/05	
Experiment required catalogues installed??	WHAT???	
Integration tests with CERN FTS		1/07/05
Integration tests T1-T2 completed with FTS		
Light-path to Lancaster provisioned	6/05/05	??
Additional 2*1Gb/s interface card in place at RAL	6/05/05	TBC
Experiment software (agents and daemons)		
3D	01/10/05?	



# What we do not know!

- How much storage is required at each site?
- How many CPUs should sites provide?
- Which additional experiment specific services need deploying for September?  
E.g. Databases (COOL - ATLAS)
- What additional grid middleware is needed and when will it need to be available (FTS, LFC)
- What are the full set of pre-requisites before SC3 can start (perhaps SC3 is now SC3-I and SC3-II)
  - Monte Carlo generation and pre-distribution
  - Metadata catalogue preparation
- What defines success for the Tier-1, Tier-2s, the experiments and LCG!





- UK participating sites have continued to make progress for SC3
- Lack of clear requirements is still causing some problems in planning and thus deployment
- Engagement with the experiments varies between sites
- Friday's "tutorial" day was considered very useful but we need more!
- We do not have enough information to plan an effective deployment (dates, components, testing schedule etc.) and there is growing concern as to whether we will be able to meet SC3 needs if they are defined late
- The clash of the LCG 2.5.0 release and start of the Service Challenge has been noted as a potential problem.