

Deployment

GridPP20 - Trinity College, Dublin
Tuesday 11th March 2008

Summary

- Please see previous talks!

... or perhaps the quick re-write of my talk can give a different perspective...

-> Consequence of the Britton-Doyle effect

Overview

- The DTEAM
- Resource situation
- Availability & reliability
- Recent news – including CCRC
- Middleware status
- Work still to do
- Summary

Who is in the DTEAM?

GridPP3: Deployment Team

- Production manager - Jeremy Coles
- London Grid coordinator - Duncan Rand
- SouthGrid coordinator - Pete Gronbech
- NorthGrid coordinator - Alessandra Forti
- ScotGrid coordinator - Graeme Stewart
- Tier-1 technical contact - Derek Ross
- Storage expert - Greig Cowan
- ATLAS technical contact - Frederic Brochu
- CMS technical contact - David Colling
- LHCb technical contact - Raja Nandakumar
- Security officer - Mingchao Ma
- Data management - Andrew Elwell
- Application support - Stephen Burke

Affiliated (means often attend)

- SRM/CASTOR - Jens Jensen
- Networking - Barney Garrett

Others on discussion list

- NGS - Matt Viljoen
- Security - David Kelsey
- Vulnerabilities - Linda Cornwall
- VOMS - Sergey
- Tier-1 on demand!

What sort of discussions/focus?


DTEAM agenda standing items

- Experiment problems & issues
- Weekly operations updates
- ROC updates (includes open tickets)
- Two discussion items
- *Team member updates*

Example recent topics

- Networking update
- Site issues (posix access and tests) [ATLAS]
- gLite transition strategy [PMB]
- 64-bit resources & SL5
- Site performance

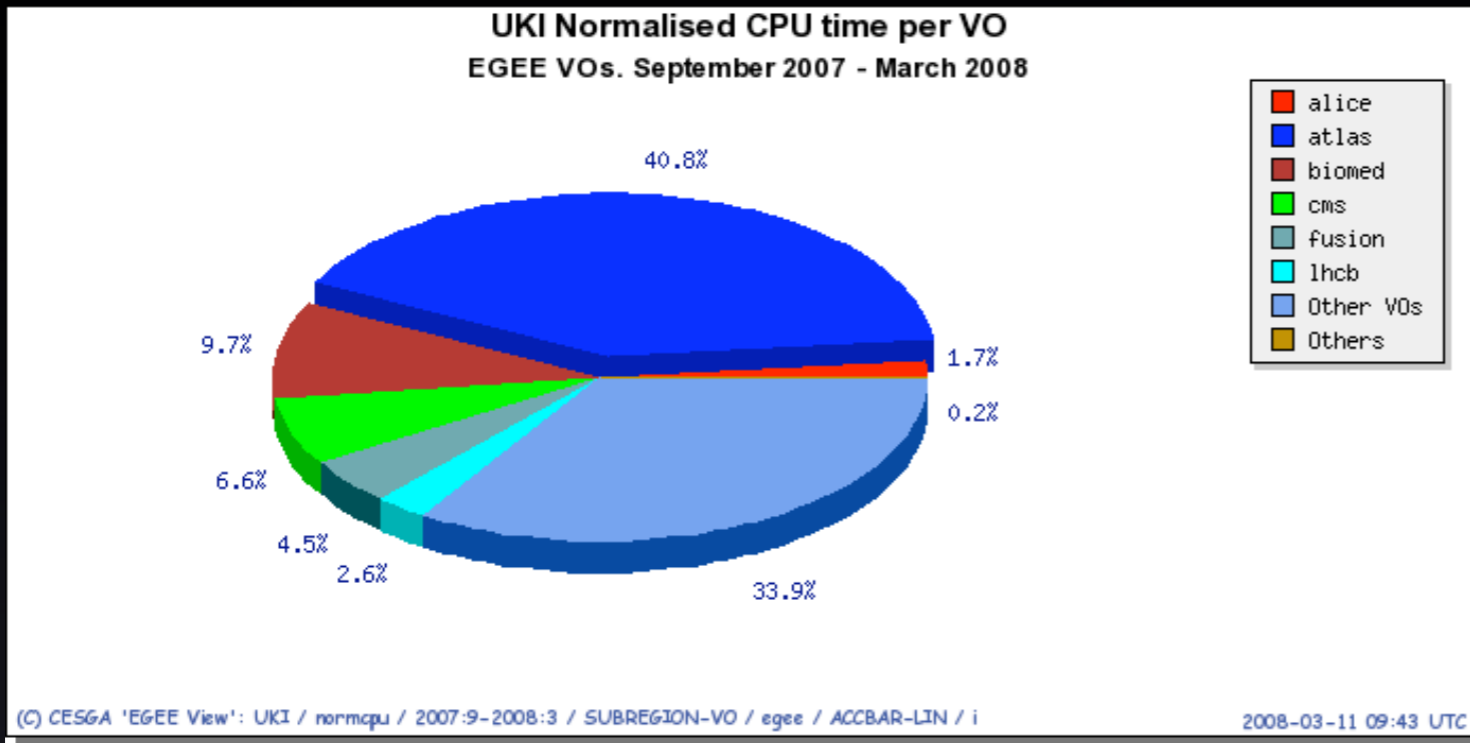
Review of T2 resources - CPU



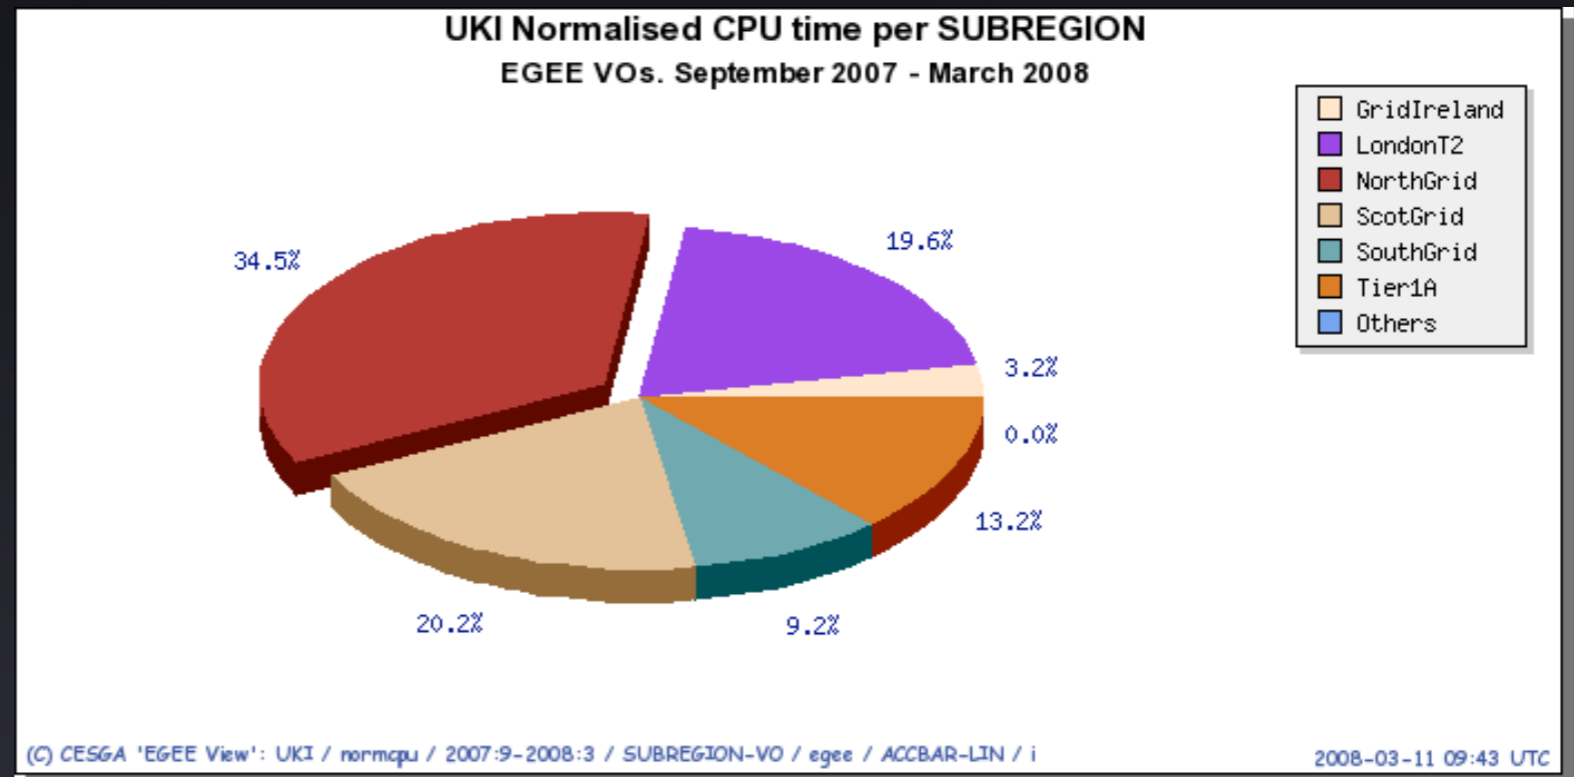
CPU snapshot from end Q4 2007

- CPU still looking good overall.
- Substantial new resources coming on at shared resource sites (Bristol; Edinburgh; UCL; Birmingham ...)


CPU usage



Saved!




Review of T2 resources - storage



- Disk available still behind projections/commitments
- Usage is increasing (note “primary” CCRC sites)
- CPU:Storage ratios still not close to experiment requirements

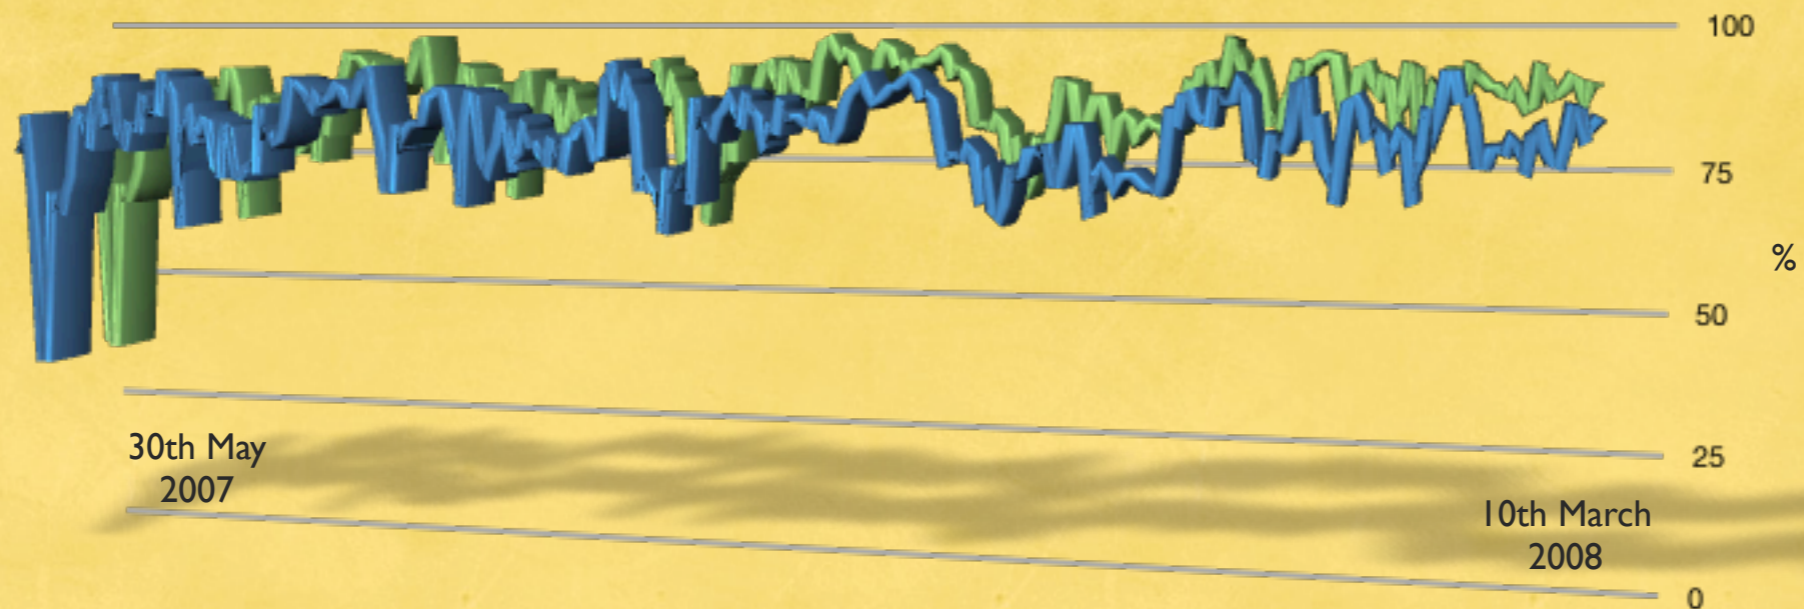
Disk usage

John Gordon plugs his laptop into the Grid!



Does this currently tell us anything interesting - other than the number of supported VOs?

Availability & Reliability



This plot shows the average (over all GridPP sites) availability and reliability over time.

Correlation as would be expected, but more recently this is less close than previously.

Focus on CCRC

- **ATLAS**: Setup of SRM endpoints (inc. spacetokens). FDR exports (success 2nd batch). Several FTS channel configuration problems. T1-T2 replication within cloud. Reprocessing at T1s. (note SRM-FTS-SRM issues). Achieved most milestones.
- **CMS**: Storage manager & T0 integration tests. T0 processing and archiving. (No high concurrent tape writing activity). Distributed data transfers. T0-T1. T1-T1. T1-T2. T2-T1 (targets generally met). T1 staging (data to buffer) and reprocessing.
- **LHCb**: Develop & test DIRAC3. Tested model. RAW data upload and distribution to T1s (mostly a success). Data reconstruction T1s+T0 (some WN upload issues). Used SRMv2.
- INSERT PLOT OF HIGH RATE... and then remove it as Dave already showed it!

Middleware issues (from CCRC)

- **lcg_utils** segfault when used with a Classic SE (no classic SE on the testbed...)
- **Long FQNS** (VO naming bug) – impacted sites hosting VOs with DNS-style VO names
- **FTS proxy corruption** – race condition when 2 clients simultaneously delegated
- FTS struggles with volatile static information
- Top level BDII's have short cache times (-> 4hrs)
- Long list for storage ... CASTOR (...); dCache (20 problems for example reserved space calc.; various copy fails; token passing; spaces unprotected...); DPM (default ACLs on dirs; slow file removal)

Recommendations/changes (CCRC)

- Avoid co-hosting BDII with other services that could generate a high load. New BDII with index set coming.
- Improve lcg_utils efficiency (stop repeated queries)
- Advise VOs about failover methods that (should not) lead to DOS
- Implement service discovery APIs - with limiters on query frequency

CCRC operations observations

- **Baseline services** defined but delivered late
- **New releases** still able to break previous set up
- **Communications** to be improved (e.g. GGUS latency)
- Not many sites attended the (short) daily update calls
- **WLCG metrics** not defined/measured – GridMap functionality improving. Expts. developing new site monitoring views too! Need global view for all 4 expt. workflows.
- **24x7 support** not in place at many sites
- Most “issues” in descending order (from elogs)
 - ⦿ Site configuration > SRM > FTS > Gridftp > dCache > Hardware
- Some problems still only noticed once m/w in production – e.g. FTS proxy corruption

Look out for ...

- **CREAM-CE**: Undergoing acceptance tests. Expect it in 8 weeks depending on YAIM module.
- **WMS on SL4**: 5 day tests started. PPS in 3 weeks?
- **glxec on WNs**: Tests in March. Batch system testing being prepared. 4 months
- ◎ Pilot job frameworks GDB working group mandated Jan 08. Review frameworks & define security requirements. Advise.
- **Job priorities**: System using VO-views. YAIM config for Maui/PBS available. Testing started. 3 months
- **AMGA-Oracle**: 4 weeks
- **FTS on SL4**: 2 months
- **SL5**: Depends on TCG priority. SL5 WN possibility for September.

Monitoring

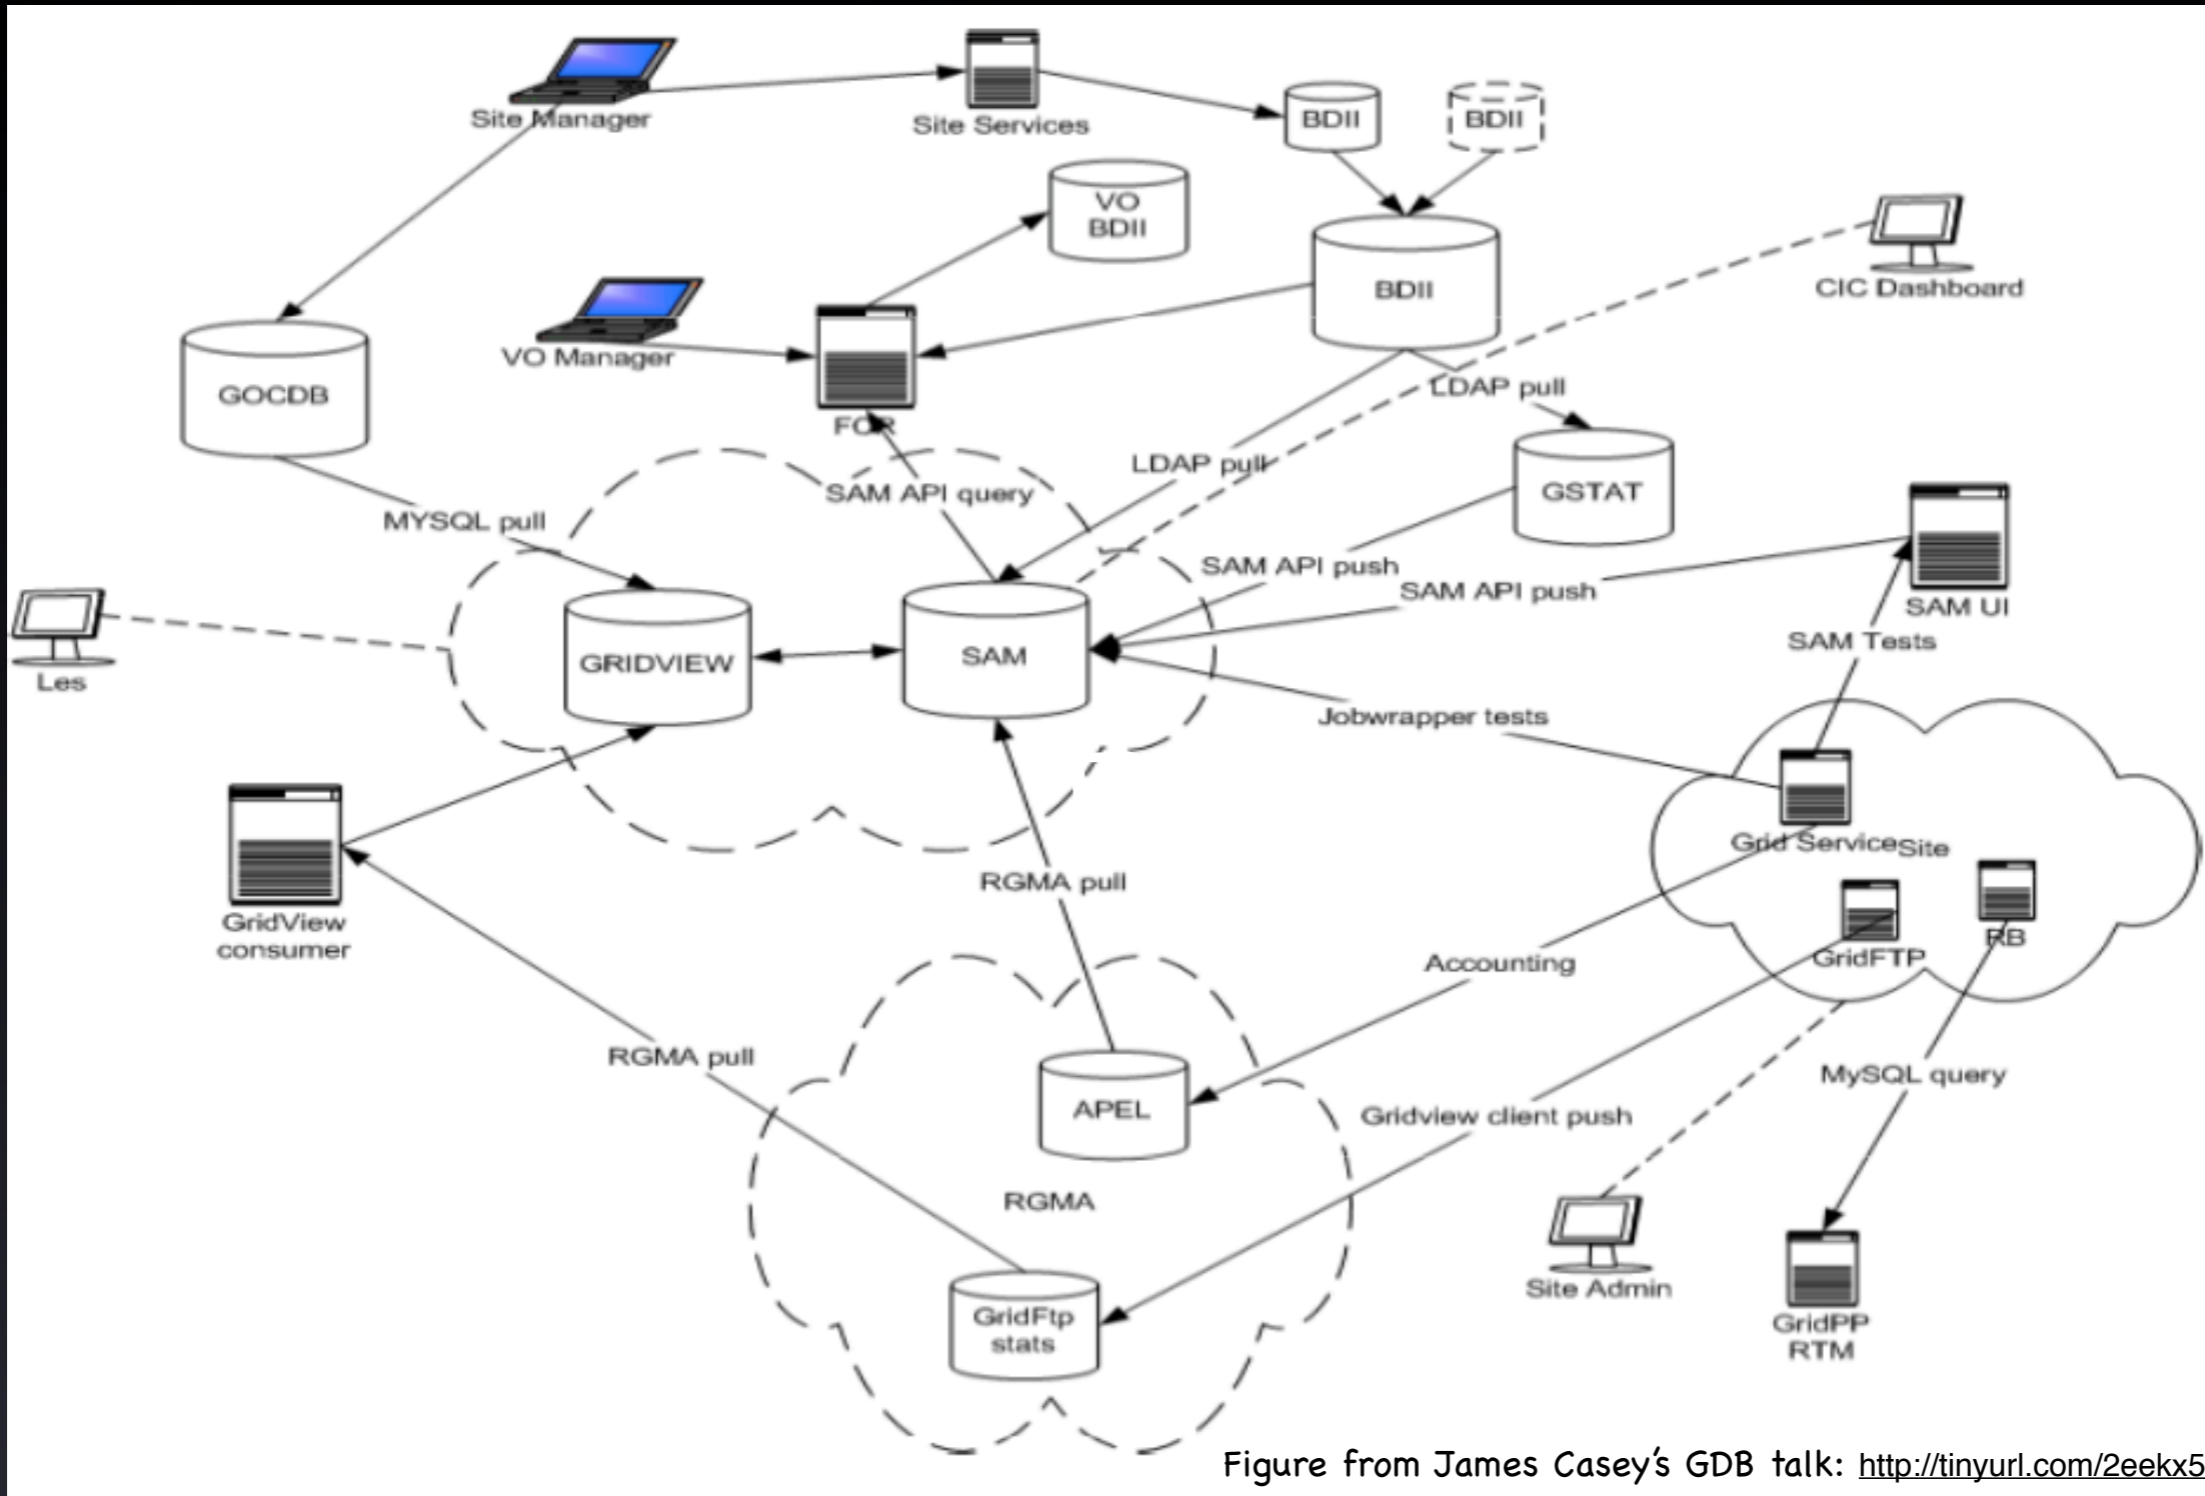














Figure from James Casey's GDB talk: <http://tinyurl.com/2eekx5>

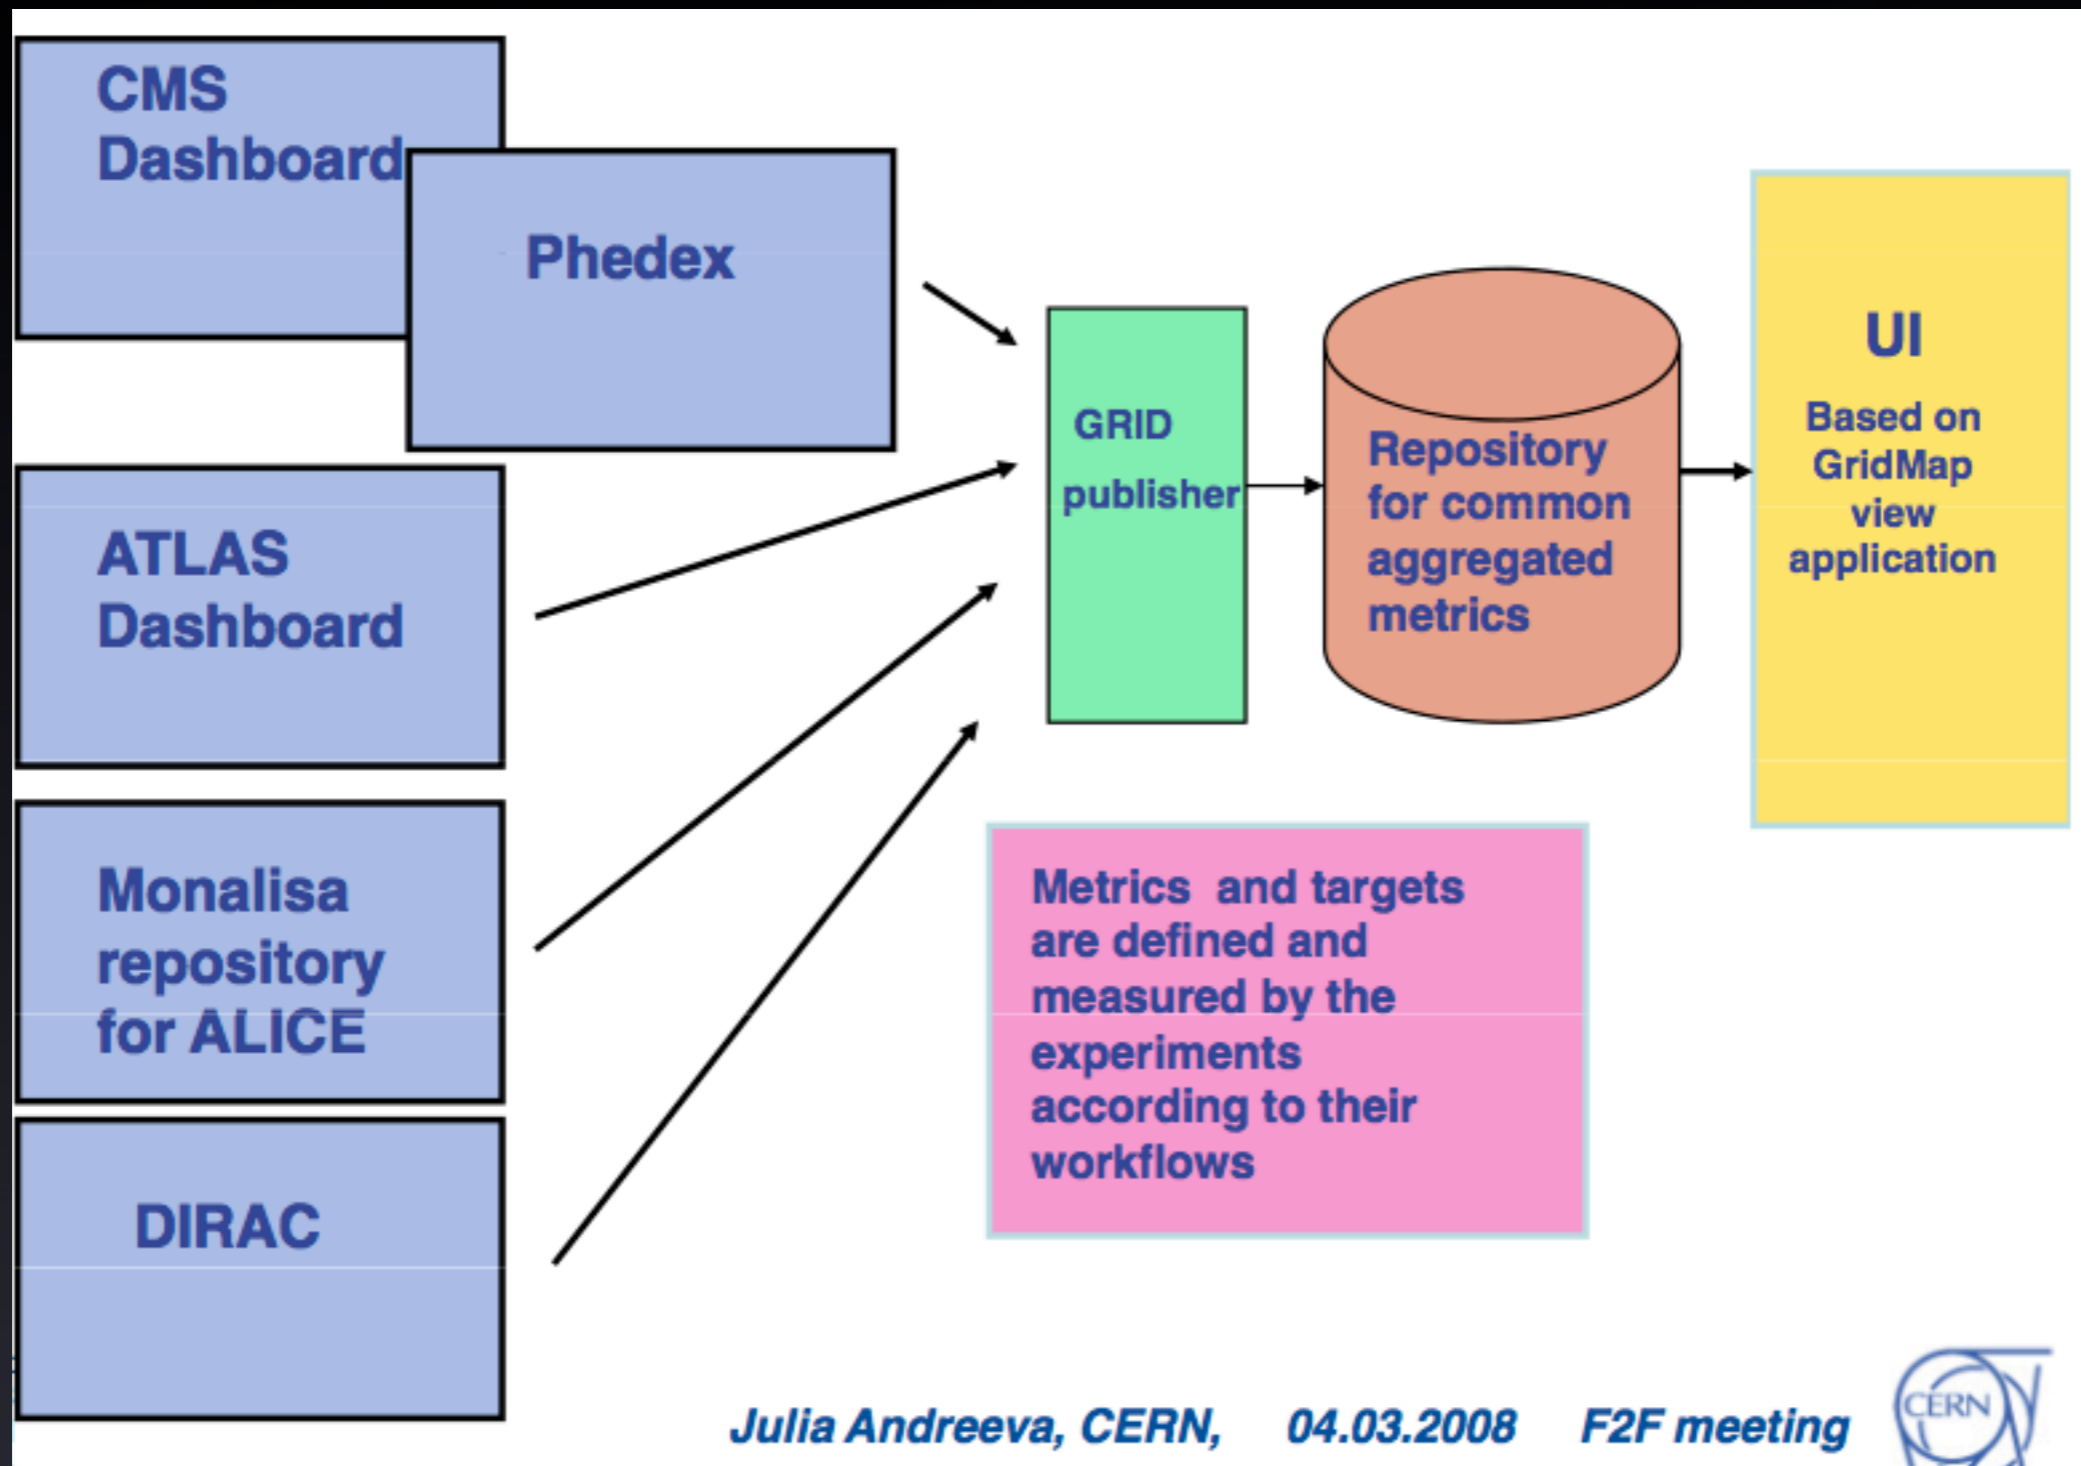
It is still a key area to be developed! Strategy is to move more to regions & provide tools for this... looking at ActiveMQ for messaging.

Monitoring - integration direction

Host ↑↓	Service ↑↓	Status ↑↓	Last Check ↑↓	Duration ↑↓	Attempt ↑↓	Status Information
ce107.cern.ch	CE-ATLAS-GngRbt-Atlas    	OK	03-04-2008 23:58:20	1d 21h 42m 37s	1/1	SAM status: ok
	CE-host-cert-valid-OPS   	OK	03-05-2008 09:15:12	14d 0h 39m 11s	1/1	SAM status: ok
	CE-lhcb-install-LHCb   	OK	03-05-2008 07:20:56	9d 7h 19m 15s	1/1	SAM status: ok
	CE-lhcb-job-Boole-LHCb   	OK	03-05-2008 07:21:02	9d 7h 19m 13s	1/1	SAM status: ok
	CE-lhcb-job-Brunel-LHCb   	OK	03-05-2008 07:21:06	9d 7h 19m 12s	1/1	SAM status: ok
	CE-lhcb-job-Gauss-LHCb   	OK	03-05-2008 07:21:01	9d 7h 19m 14s	1/1	SAM status: ok
	CE-lhcb-os-LHCb   	OK	03-05-2008 07:20:50	9d 7h 19m 18s	1/1	SAM status: ok
	CE-lhcb-queues-LHCb   	OK	03-05-2008 07:20:54	0d 4h 20m 13s	1/1	SAM status: ok
	CE-sft-brokerinfo-OPS    	OK	03-05-2008 09:43:07	14d 1h 27m 33s	1/1	SAM status: ok
	CE-sft-caver-Alice    	OK	02-18-2008 12:04:00	15d 23h 37m 7s	1/1	SAM status: ok
	CE-sft-caver-Atlas    	OK	03-05-2008 00:13:37	14d 1h 4m 30s	1/1	SAM status: ok
	CE-sft-caver-OPS    	OK	03-05-2008 09:43:13	14d 1h 27m 18s	1/1	SAM status: ok
	CE-sft-csh-OPS    	OK	03-05-2008 09:43:07	14d 1h 27m 31s	1/1	SAM status: ok
	CE-sft-job-Alice    	OK	02-18-2008 12:53:47	15d 22h 47m 20s	1/1	SAM status: ok
	CE-sft-job-Atlas    	CRITICAL	03-05-2008 07:38:21	0d 10h 6m 0s	1/1	SAM status: error
	CE-sft-job-CMS    	OK	03-05-2008 09:18:26	0d 2h 22m 41s	1/1	SAM status: ok
	CE-sft-job-OPS    	OK	03-05-2008 10:18:42	0d 1h 22m 25s	1/1	SAM status: ok
	CE-sft-lcg-rm-OPS    	OK	03-05-2008 09:43:39	5d 19h 23m 48s	1/1	SAM status: ok
	CE-sft-softver-OPS    	OK	03-05-2008 09:43:07	14d 1h 27m 33s	1/1	SAM status: ok
	CE-sft-vo-swdir-Alice    	OK	02-18-2008 12:03:55	15d 23h 37m 12s	1/1	SAM status: ok
	CE-sft-vo-swdir-Atlas    	OK	03-05-2008 00:13:31	14d 1h 4m 36s	1/1	SAM status: ok
	CE-sft-vo-tag-Atlas    	CRITICAL	03-05-2008 00:13:32	14d 1h 4m 35s	1/1	SAM status: error

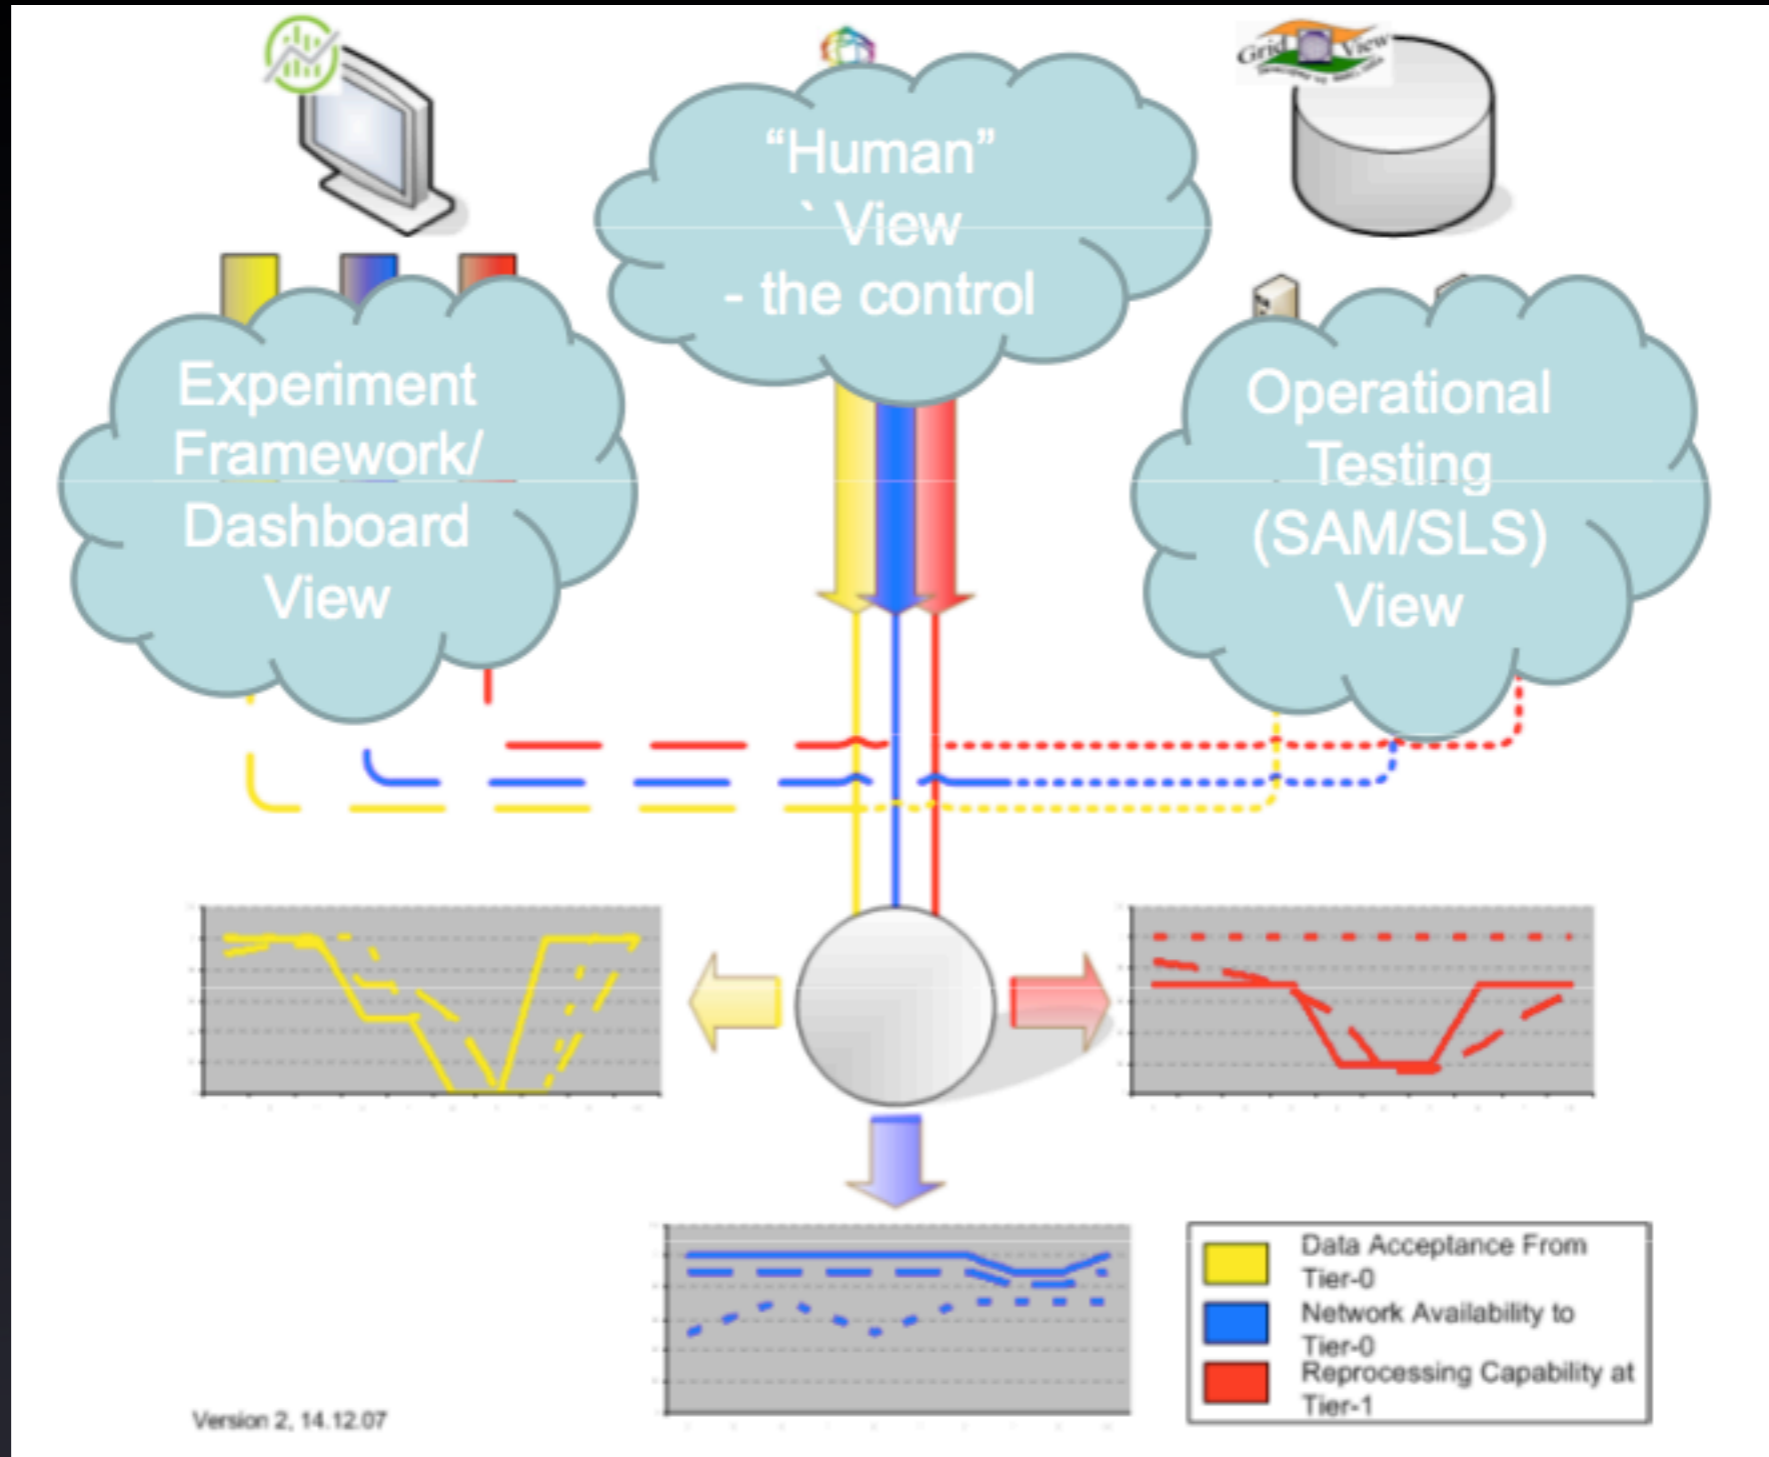
I. View SAM results within site monitoring

Experiment workflows



See Julia Andreeva's talk: <http://tinyurl.com/2hl8mh>

The overall view



James Casey to the rescue again!

<http://tinyurl.com/yqsl83>

Summary

- The DTEAM definition changes slightly in GridPP3
- The T2 resource situation is much as at GridPP19
- Most sites have improved their reliability & availability; a few have brought the averages down
- CCRC part 1 has run with general success
- The results have been well reviewed – recommendations and improvements made in several areas
- There are many areas still in progress and in particular monitoring