



# CCRC – Review of February run based on agreed metrics

**Jamie Shiers**

~ ~ ~

**WLCG Collaboration Workshop, 22 April 2008**

# Objectives

- Primary objective (**next**) was to demonstrate that we (sites, experiments) could run together at 2008 production scale
  - **This includes testing all “functional blocks”:**
    - **Experiment to CERN MSS; CERN to Tier1; Tier1 to Tier2s etc.**
- Two challenge phases were foreseen:
  1. **February :** not all 2008 resources in place – still adapting to new versions of some services (e.g. SRM – **later**) & experiment s/w
  2. **May:** all 2008 resources in place – full 2008 workload, all aspects of experiments' production chains
-  N.B. failure to meet target(s) would necessarily result in discussions on **de-scoping!**
- Fortunately, the results suggest that this is not needed, although **much** still needs to be done before, and during, May!



# CCRC'08 – Motivation and Goals

What if:

- LHC is operating and experiments take data?
- All experiments want to use the computing infrastructure simultaneously?
- The data rates and volumes to be handled at the Tier0, the Tier1 and Tier2 centers are the sum of ALICE, ATLAS, CMS and LHCb as specified in the experiments computing model
- Each experiment has done data challenges, computing challenges, tests, dress rehearsals, .... at a schedule defined by the experiment
- This will stop: we will no longer be the master of our schedule...  
.... Once LHC starts to operate.
- We need to prepare for this ... together ....

A combined challenge by all Experiments should be used to demonstrate the readiness of the WLCG Computing infrastructure before start of data taking at a scale comparable to the data taking in 2008.

This should be done well in advance of the start of data taking on order to identify flaws, bottlenecks and allow to fix those.

We must do this challenge as WLCG collaboration: Centers and Experiments

# How We Measured Our Success

- Agreed up-front on specific targets and metrics – these were 3-fold and helped integrate different aspects of the service ([CCRC'08 wiki](#)):
  1. Explicit “**scaling factors**” set by the experiments for each functional block: discussed in detail together with sites to ensure that the necessary resources and configuration were in place;
  2. Targets for the lists of “**critical services**” defined by the experiments – those essential for their production, with an analysis of the impact of service degradation or interruption (**discussed at previous OBs...**)
  3. WLCG “**Memorandum of Understanding**” (MoU) targets – services to be provided by sites, target availability, time to intervene / resolve problems ...
- ↳ **Clearly some rationalization of these would be useful – significant but not complete overlap**

# What Did We Achieve? (High Level)

- Even **before** the official start date of the February challenge, the exercise had proven an extremely **useful** focusing exercise, in helping understand missing and / or weak aspects of the service and in identifying **pragmatic** solutions
- Although later than desirable, the **main bugs** in the middleware were fixed (just) in time and many sites upgraded to these versions
- The deployment, configuration and usage of SRM v2.2 went **better** than had predicted, with a noticeable **improvement** during the month
- Despite the high workload, we also **demonstrated** (most importantly) that **we can support** this work with the available manpower, although essentially no remaining effort for longer-term work (**more later...**)
- If we can do the same in May – when the bar is placed much higher – we will be in a **good position** for this year's data taking
- However, there are certainly **significant concerns** around the available manpower at all sites – not only today, but also in the longer term, when funding is unclear (**e.g. post EGEE III**)

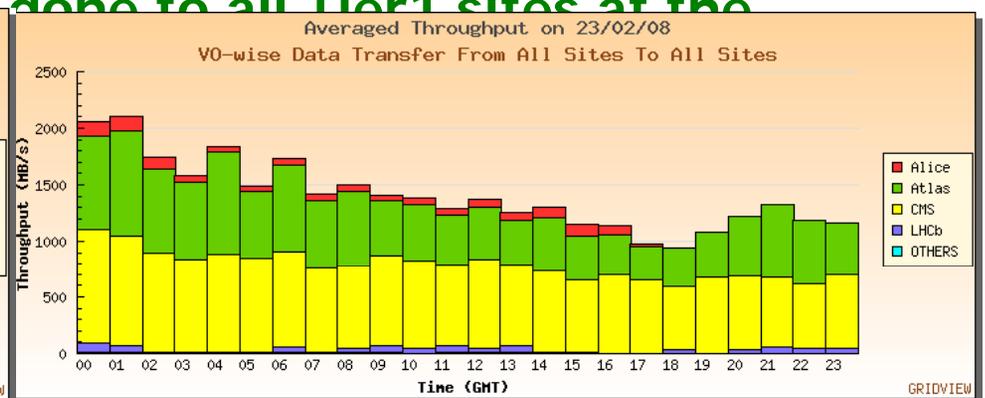
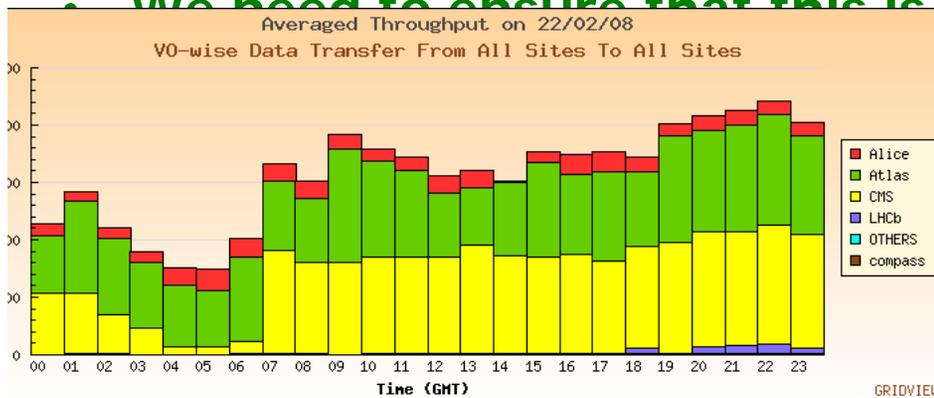
# CCRC'08 Preparations...

- Monthly Face-to-Face meetings held since time of “kick-off” during WLCG Collaboration workshop in Victoria, Canada
- Fortnightly con-calls with Asia-Pacific sites started in January 2008
- Weekly planning con-calls ← suspended during February: restart?
- **Daily “operations” meetings @ 15:00 started mid-January**
- ↳ **Quite successful in defining scope of challenge, required services, setup & configuration at sites...**
- **Communication – including the tools we have – remains a difficult problem... but...**
- **Feedback from sites regarding the information they require, plus “adoption” of common way of presenting information (modelled on LHCb) all help**
- ☺ **We were arguably (much) better prepared than for any previous challenge**
- **There are clearly some lessons for the future – both the May CCRC'08 challenge as well as longer term**

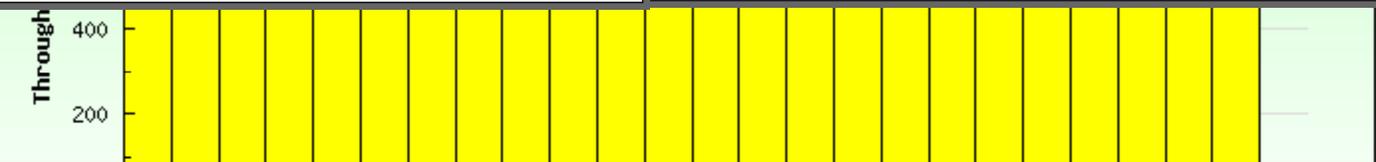
# Tier0 – Tier1 Data Export

- We need to sustain 2008-scale exports for at least ATLAS & CMS for at least two weeks
  - The short experience that we have is not enough to conclude that this is a solved problem [ **experience to early march suggest this is now OK!** ]
  - **The overall system still appears to be too fragile – sensitive to ‘rogue users’ (what does this mean?) and / or DB de-tuning**
  - **(Further) improvements in reporting, problem tracking & post-mortems needed to streamline this area**

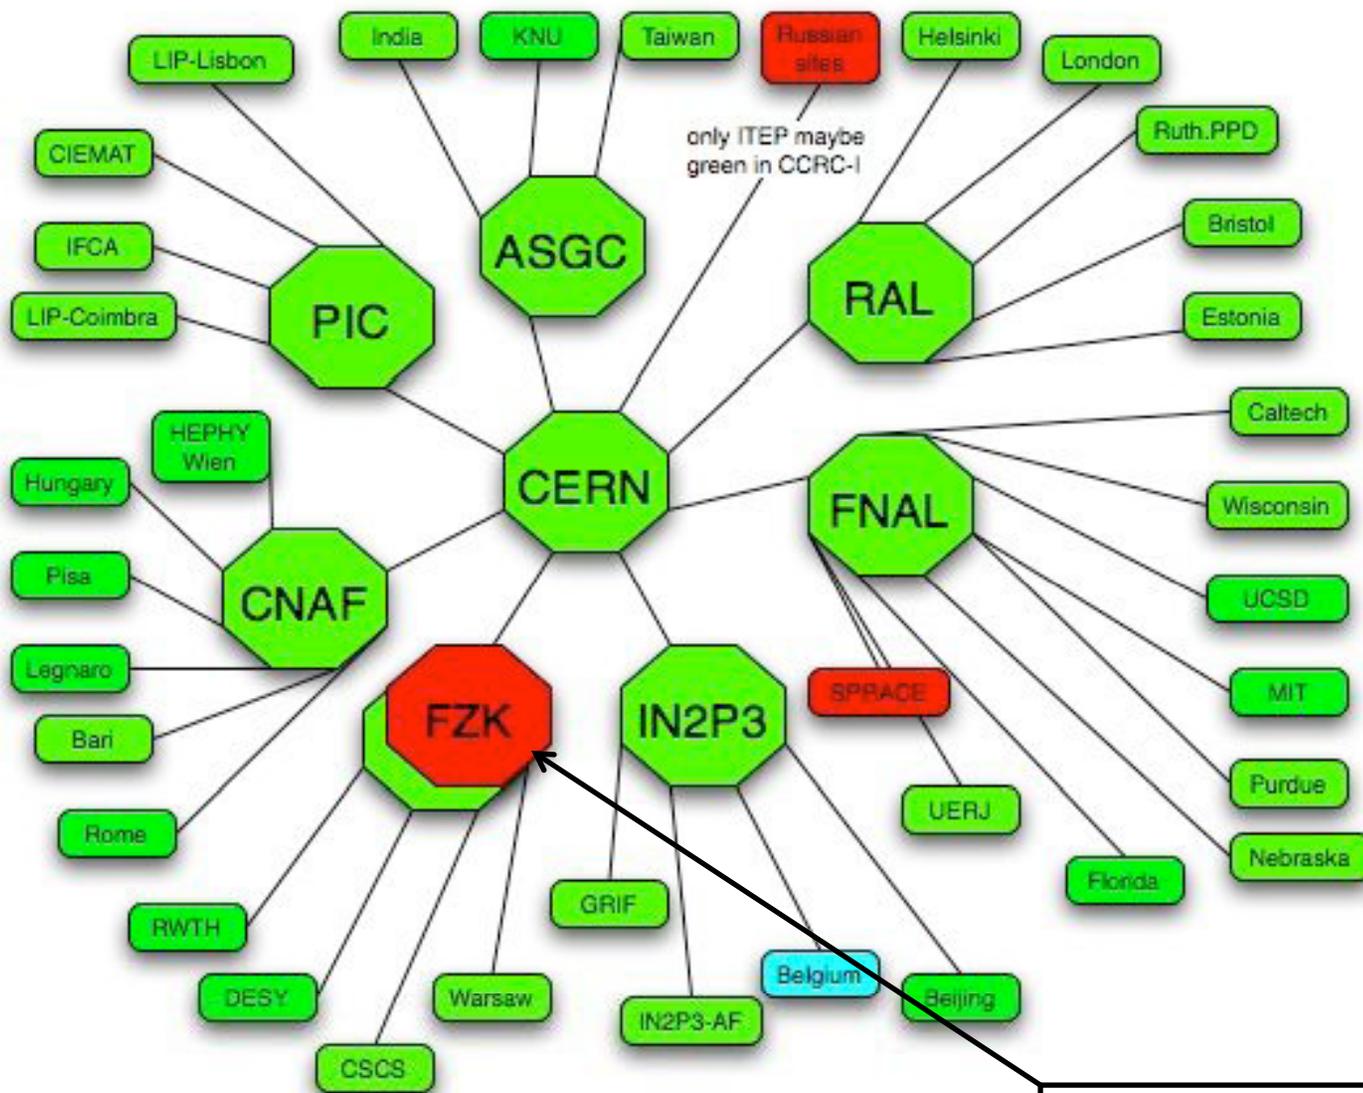
**We need to ensure that this is done to all Tier1 sites at the**



- Re-adjust
- e.g.



**Mid-February checkpoint (meeting with LHCC referees)**



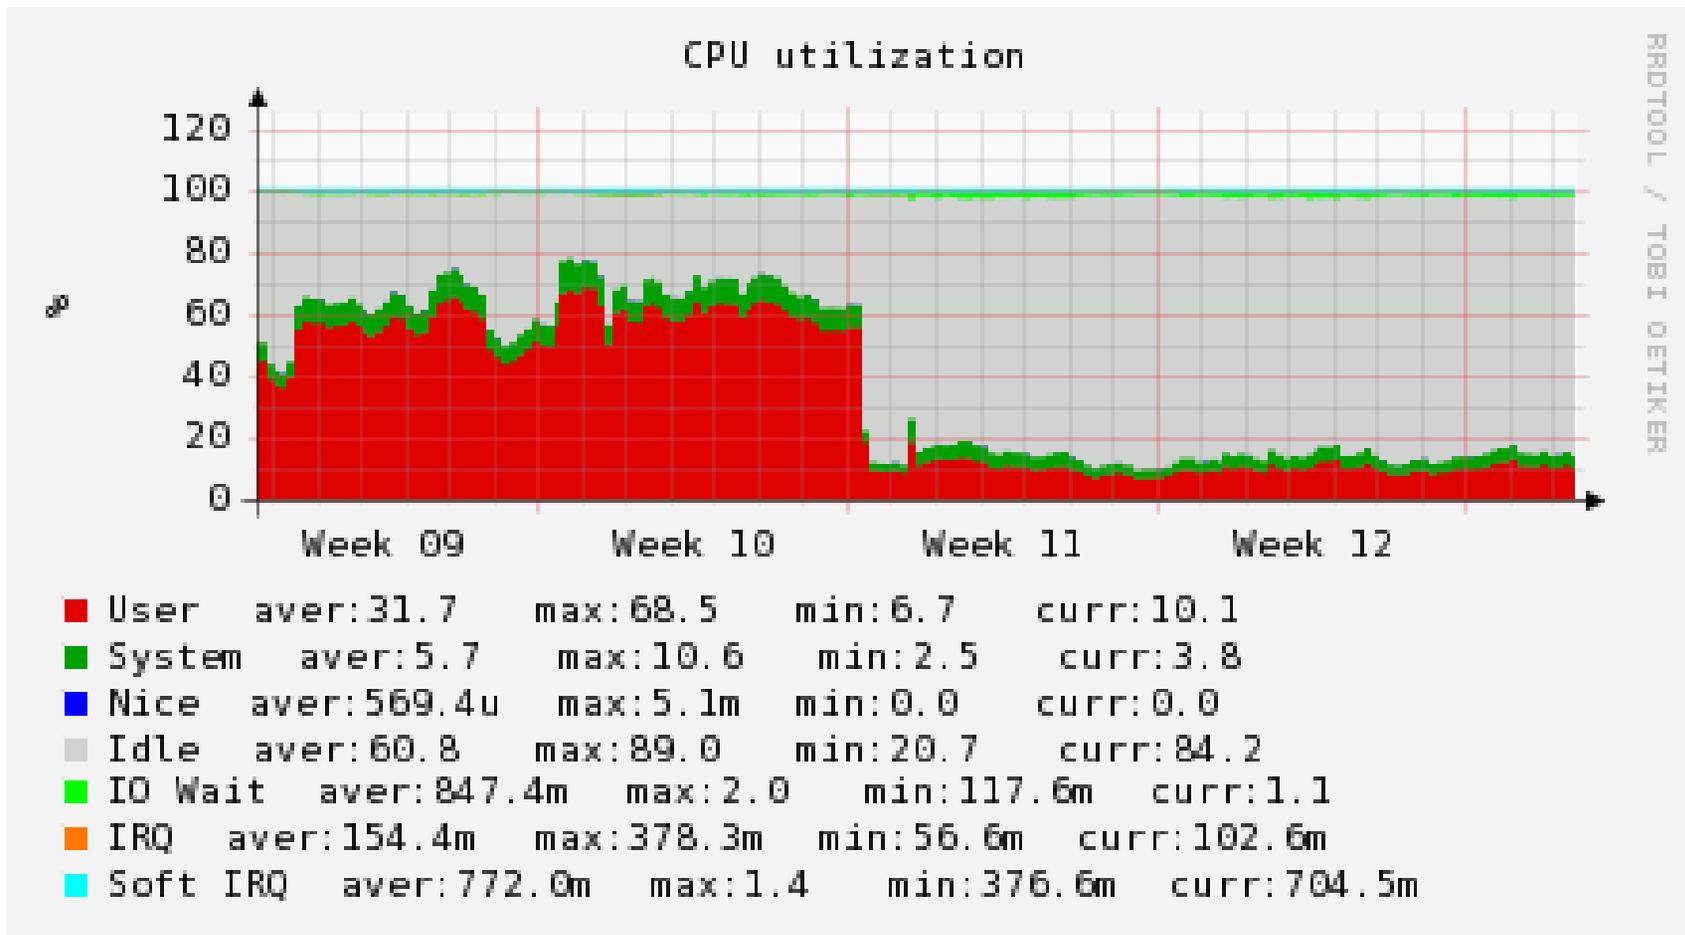
NOTE: only T0-T1 and regional T1-T2 links are depicted here (+ Russian sites)

- SRMv2 ready in week-1
- SRMv2 ready in week-2
- SRMv2 ready in week-3/4
- SRMv2 not 100% ready within CCRC-I time window

See slide note!

# CERN CE Load & Improvements

- 💣 Early on in February, CERN LCG CEs regularly had 'high load' alarm – scaling issue for the future???



# Database Issues

- Databases behind **many** of the most critical services listed by the experiments – “best effort” out of hours at CERN!
  - No **change** in Tier0 load seen during February – is the workload representative of 2008 data-taking or will this only be exercised in May? **Later??**
  - Oracle Streams **bug** triggered by (non-replicated) table compression for PVSS – **one week** for patch to arrive!
  - Interest in moving to Oracle 10.2.0.4 prior to May run of CCRC'08 **if** validated in time by experiments & WLCG – to be discussed at April F2F meetings
    - Motivation: avoid **long delays** (see above) in receiving any needed bug fixes – avoid back-porting to 10.2.0.3 – but is there enough time for realistic testing???
    - Remember “(shared) cached cursor” syndrome?
- **Looking unlikely at this stage...**

# LFC @ WLCG Tier1 Sites

Site (ISO standard)	Experiments (ATLAS, LHCb)	Comments
CA-TRIUMF	ATLAS	MySQL, plan to migrate to Oracle in 2009.
DE-KIT	ATLAS, LHCb	Oracle for LHCb, MySQL for others
ES-PIC	ATLAS, LHCb	
FR-CCIN2P3	ATLAS, LHCb	Oracle
IT-INFN-CNAF	ATLAS, LHCb	Oracle
NDGF	ATLAS	<i>r/s?</i> (LFC on <b>PostgreSQL(?)</b> or MySQL)
NL-T1	ATLAS, LHCb	Oracle (+MySQL for non-HEP VOs)
TW-ASGC	ATLAS	Oracle
UK-RAL	ATLAS, LHCb	MySQL → Oracle in 2-3 weeks (ATLAS)
US-BNL	ATLAS	LRC. Evaluating LFC. b/e decision later.
US-FNAL	N/A	N/A

Discussions on resilience to T1 failure in ATLAS clouds 1

# Service Observations (1/2)

- We must **standardize** and **clarify** the operator/experiment communications lines at Tier0 and Tier1.
- The management board milestones of providing 24x7 support and implementing agreed experiment VO-box Service Level Agreements need to be completed as soon as possible.
- As expected there were many **teething** problems in the first two weeks as SRMv2 endpoints were setup (over 160) and early bugs found after which the SRMv2 deployment worked **generally well**.
- Missing functionalities in the data management layers have been exposed (CCRC'08 "Storage Solutions Working Group" was closely linked to the February activities) and follow-up planning is in place.
- **Few** problems found with middleware – only 1 **serious** ([FTS proxy delegation corruption](#)), with work-around in place.
- Tier1s proved fairly reliable: follow-up on their tape operations worked is required.

# Critical Service Follow-up

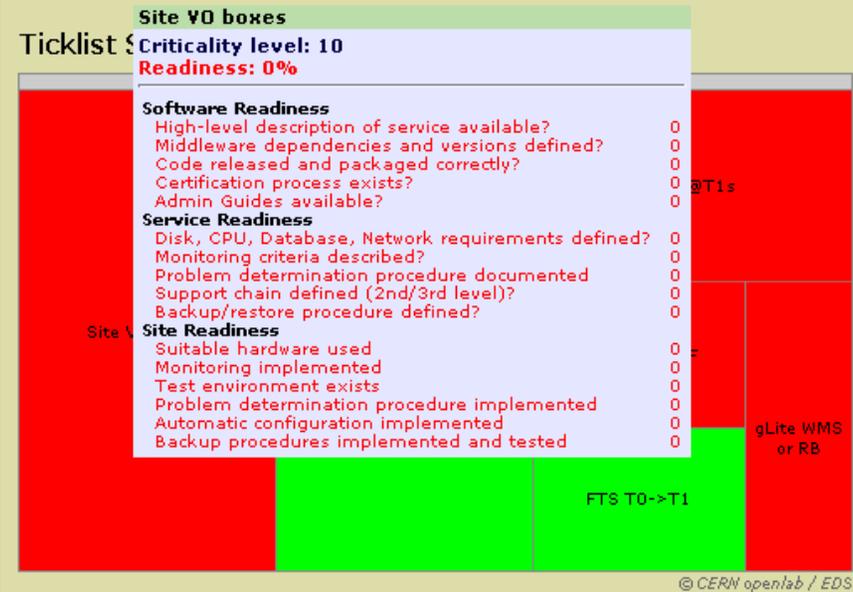
- Targets (not commitments) proposed for Tier0 services
  - Similar targets requested for Tier1s/Tier2s
  - Experience from first week of CCRC'08 suggests targets for **problem resolution** should not be too high (if ~achievable)
    - The MoU lists targets for responding to problems (12 hours for T1s)
  - ¿ Tier1s: 95% of problems resolved <1 working day ?
  - ¿ Tier2s: 90% of problems resolved < 1 working day ?
- **Post-mortem triggered when targets not met!**

Time Interval	Issue (Tier0 Services)	Target
End 2008	Consistent use of all WLCG Service Standards	100%
30'	Operator response to alarm / call to x5011 / <b>mailing list</b>	99%
1 hour	Operator response to alarm / call to x5011 / <b>mailing list</b>	100%
4 hours	Expert intervention in response to above	95%
8 hours	Problem resolved	90%
24 hours	Problem resolved	99%

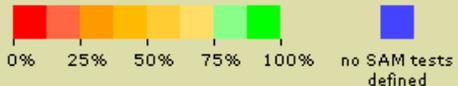
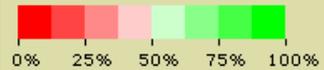
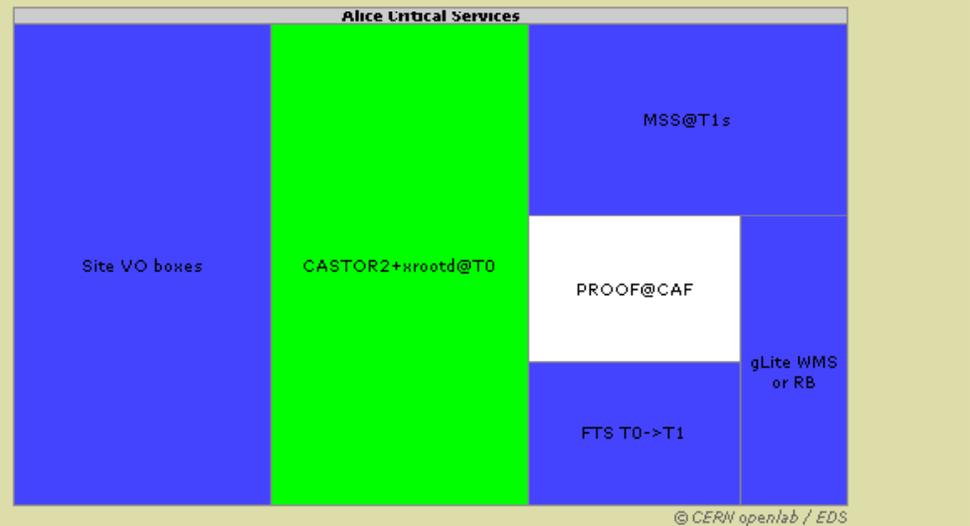
# Service Observations (2/2)

- Some particular experiment problems seen at the WLCG level:
  - **ALICE:** Only one Tier1 (FZK) was fully ready, NL-T1 several days later than the last 3 only on the last day (RAL being setup in March)
  - **ATLAS:** Creation of physics mix data sample took much longer than expected and a reduced sample had to be used.
  - **CMS:** Inter-Tier1 performance not as good as expected.
  - **LHCb:** New version of Dirac had teething problems – 1 week delay.
- Only two inter-experiment interferences were logged: FTS congestion at GRIF caused by competing ATLAS and CMS SEs (solved by implementing sub-site channels) and degradation of CMS exports to PIC by ATLAS filling the FTS request queue with retries.
- We must collect and analyze the various metrics measurements.
- ☺ **The electronic log and daily operations meetings proved very useful and will continue.**
- ☹ **Not many Tier1s attend the daily phone conference and we need to find out how to make it more useful [ for them! ]**
- Overall a good learning experience and positive result. Activities will continue from now on with the May run acting as a focus point.

# WLCG CCRC'08 Critical Services "GridMap"



## Test Status (live data)



Alice Atlas CMS LHCb

CERN-PROD TRIUMF-LCG2 IN2P3-CC FZK-LCG2 INFN-T1 SARA-MATRIX NDGF-T1 pic Taiwan  
BNL-LCG2 USCMS-FNAL-WC1

# Experiment View

## In Order of Appearance (March F2F...)

- CMS
  - [ Very ] Detailed presentation of up-front metrics per functional block
  - 100% success not reported, but well understood status
- ATLAS:
  - CCRC was a very useful exercise for ATLAS
  - Achieved most milestones in spite of external dependencies
  - It's difficult to serve the Detector, Physics and IT community!
- ALICE:
  - For ALICE, the CCRC exercise has fulfilled its purpose
    - Focus on data management
    - Brings all experiments together
    - Controlled tests, organization
- LHCb:
  - Initial phase of CCRC'08 was dedicated to development and testing of DIRAC3
  - CCRC'08 now running smoothly
    - Online->T0 and T0-T1 transfers on the whole a success
    - Some issues with reconstruction activity and data upload from the WNs
    - Investigating with Tier-1s recent problem of determining file sizes using gfal
    - Quick turnaround for reported problems

# Well, How Did We Do?

- Remember that prior to CCRC'08 we:
  - a) Were not confident that we were / would be able to support all aspects of all experiments simultaneously
  - b) Had discussed possible fall-backs if this were not demonstrated
    - **The only conceivable "fall-back" was de-scoping...**
- Now we are reasonably confident of the former
- Do we need to retain the latter as an option?
- Despite being rather late with a number of components (not desirable), things settled down reasonably well
- Given the much higher "bar" for May, need to be well prepared!

# Main Lessons Learned

☺ **Generally, things worked reasonably well...**

➤ **Still improvements in communication are needed!**

- Tools still need to be streamlined (e.g. elog-books / GGUS), and reporting automated

¿ **Service dashboards – should be in place before May...**

- F2Fs and other meetings working well in this direction!

• **Pre-established metrics extremely valuable!**

- As well as careful preparation and extensive communication!

↳ **Now continuous production mode –this will continue – as will today's infrastructure & meetings**

# Recommendations

- ✓ **To improve communications with Tier2s and the DB community, 2 new mailing lists have been setup, as well as regular con-calls with Asia-Pacific sites (time zones...)**
- Follow-up on the lists of “**Critical Services**” must continue, implementing not only the appropriate monitoring, but also ensuring that the WLCG “standards” are followed for **D**esign, **I**mplementation, **D**eployment and **O**peration
- **Clarify reporting and problem escalation lines (e.g. operator call-out triggered by named experts, ...) and introduce (light-weight) post-mortems when MoU targets not met**
- **We must continue to improve on open & transparent reporting, as well as further automations in monitoring, logging & accounting**
- ↳ **We should foresee “data taking readiness” challenges in future years – probably with a similar schedule to this year – to ensure that full chain (new resources, new versions of experiment + AA s/w, middleware, storage-ware) is ready**

# CCRC'08 Summary from February

- The preparations for this challenge have proceeded (largely) smoothly – we have both learnt and advanced a lot simply through these combined efforts
  - As a focusing activity, CCRC'08 has already been very useful
  - We will learn a lot about our overall readiness for 2008 data taking
  - We are also learning a lot about how to run smooth production services in a more sustainable manner than previous challenges
- 💣 **It is still very manpower intensive and schedules remain extremely tight: full 2008 readiness still to be shown!**
- 👉 **More reliable – as well as automated – reporting needed**
- Maximize the usage of up-coming F2Fs (March, April) as well as WLCG Collaboration workshop to fully profit from these exercises
- **June on: continuous production mode (all experiments, all sites), including tracking / fixing problems as they occur**

# Summary

- The February run of CCRC'08 was largely **successful** and introduced the important element of up-front and **measurable** metrics for many aspects of the service
  - We still have a lot to do to **demonstrate** full 2008 readiness – May and beyond will be a busy time with no let up in production services prior to then
  - We have developed a way of planning and operating the service that works **well** – re-enforce this and build on it incrementally
- **This is it – the WLCG Production Service!**