Grid Today, Clouds on the Horizon

Jamie Shiers, Jamie.Shiers@cern.ch [a]

[a] *CERN, 1211 Geneva 23, Switzerland*

Abstract

By the time of CCP 2008, the largest scientific machine in the world - the Large Hadron Collider - had been cooled down as scheduled to its operational temperature of below 2 degrees Kelvin and injection tests were starting. Collisions of proton beams at 5 + 5 TeV were expected within one to two months of the initial tests, with data taking at design energy (7 + 7 TeV) foreseen for 2009.

In order to process the data from this world machine, we have put our "Higgs in one basket" - that of Grid computing (1). After many years of preparation, 2008 saw a final "Common Computing Readiness Challenge" (CCRC08) - aimed at demonstrating full readiness for 2008 data taking, processing and analysis. By definition, this relied on a world-wide production Grid infrastructure.

But change - as always - is on the horizon. The current funding model for Grids - which in Europe has been through 3 generations of EGEE projects, together with related projects in other parts of the world, including South America - is evolving towards a long-term, sustainable e-infrastructure, like the European Grid Initiative (EGI) (2). At the same time, potentially new paradigms, such as that of "Cloud Computing" are emerging.

This paper summarizes the results of CCRC08 and discusses the potential impact of future Grid funding on both regional and international application communities. It contrasts Grid and Cloud computing models from both technical and sociological points of view. Finally, it discusses the requirements from production application communities, in terms of stability and continuity in the medium to long term.

Key words: Grid; Cloud; Distributed Computing; Resilience

1. Introduction

Operating at below 2 degrees Kelvin, the world's largest scientific machine to date – the Large Hadron Collider (LHC) at CERN – is also the coolest. The date for first circulating beam in this machine – targetted for September 10th 2008 – was announced in a press release (3) that coincided with the CCP 2008 conference. Whilst the physics not only of the machine but also that to be studied by the massive detectors that will take data from the collisions of particles accelerated by the machine is of great interest, this paper focusses on the computational systems that have been put in place to perform the offline production and analysis of these data.

An estimated 15PB of data will be collected per year and processed using a world–wide grid that is built on the production grids provided primarily by the Enabling Grids for E-sciencE (EGEE) (4) project in Europe, the Open Science Grid (OSG) (5) in the Americas, with additional resources being provided in the Nordic countries (6) and elsewhere.

During the past year a series of challenges have been carried out designed to show the readiness of these offline computing systems. Termed the Common Computing Readiness Challenge 2008 (CCRC'08 for short), the name already hints that future such challenges may be required. Indeed, as the computing infrastructure continues to expand and evolve, it is our intent to carry out such a data taking readiness challenge at the beginning of each year. With the LHC entering operation this year, future challenges will inevitably be of a somewhat different nature as they will have to accomodate re-processing of the previous year's (later years') data that will need to be performed in parallel.

The importance of CCRC'08 is indicated by the simple fact that one year ago we had still not demonstrated our ability to support all the needed activities of all four experiments (ATLAS, CMS, LHCb and ALICE) concurrently. Failure to do so during CCRC'08 would have inevitably have led to painful discussions on de–scoping. Fortunately, as is described in more detail below, early results were sufficiently encouraging that this possibility was swiftly dismissed.

Aside from the demonstration of readiness for the demands of 2008 data taking and processing, this exercise was also the most demanding proof of worldwide production grids to date. This in itself is an important achievement and – with today's increased emphasis on alternate models, such as commercial "cloud computing", calls for a cost-benefit analysis of the various approaches used or considered today in this domain.

2. The Common Computing Readiness Challenge

The Common Computing Readiness Challenge was designed to show that "we" – the set of four experiments, the several hundred sites and the WLCG service (1) – are ready to handle the needs of 2008 data taking and processing. Previously, each experiment had performed numerous tests, with further "generic" tests showing that individual components were ready. We had not, however, demonstrated all four experiments together with all production data flows. Furthermore, a number of key components were deployed relatively late (distributed database services) or else there were new versions (data / storage management) with significantly changed functionality. Finally, 2008 resources were only required to be in production by the beginning of April – at the bare minimum, a new scale test would have been required (and will be needed in future years).

The experiments data processing is often referred to as "functional blocks". These correspond to the activities performed at the different sites (for an overview of the WLCG computing model, including the roles of the various Tiers, see (1)), such as raw data recording, first pass reconstruction, distribution to the Tier1s, reprocessing, simulation, analysis and so forth. For each such functional block a set of metrics per experiment and per site were agreed up-front. Similarly, for all of the services identified by the experiments as critical to their production, targets for intervention and problem resolution were set. Finally, the WLCG Memorandum of Understanding lists the services and service levels that each site must provide. Whilst similar to the above, these represent an additional set of metrics that were measured using service availability monitoring (e.g. service availability as a percentage of wall–clock time measured on an annual basis).

It was intended that these metrics be reviewed after the challenge to see if they were
– measurable;
– specific;
– achievable.

Not only were there concerns about possible de–scoping, but also the late deployment of the new storage management software (Storage Resource Manager (SRM) v2.2) was felt to be a significant risk factor. The targets eventually agreed for production deployment were the end of 2007 at the Tier0 and the "main" (self–defining) Tier1s and at all Tier2s by the end of February. These targets were largely met with close to 200 sites having deployed SRM v2.2 services by the end of February 2008.

Two phases of the challenge were forseen:
– February – essentially an integration test with not all 2008 resources deployed;
– May – with the full 2008 resources plus resolution of any issues found in February.

The February run went rather smoothly – thanks at least in part to the thorough preparation and follow–up, with monthly face–to–face planning meetings, weekly conference calls and finally daily operations calls. These daily operations calls continue to this day and are expected to be a permanent feature of LHC operations. The experiments also run sometimes twice–daily experiment operations meetings, as well as their own computing operations workshops – sometimes referred to as "jamborees".

Whilst there was less overlap of the experiments' activities than desirable during February, data export tests well exceeded the 2008 requirements and the overall service ran sufficiently smoothly (a few problems reported per experiment per day) to remove the need for discussions on de–scoping.

For a variety of reasons, not all sites managed to meet their commitments for deploying the full

2008 resources in production by 1st April. (This itself rang alarm bells – hardware acquisition, installation and commissioning is a lengthy process and one cannot risk "just–in–time" purchasing, even if on paper this will lead to better performance per currency unit.) However, this run managed not only to meet or exceed all of the scaling targets but also significantly improved on the overall service reliability. The number of problems reported were less and the time taken to resolve them significantly shorter. These were largely configuration problems – something that should be avoidable with proper configuration management.

In both the February and May runs the number of middleware bugs was extremely low (1–2) and the middleware release process functioned smoothly and efficiently. The same can be said of many other areas of the service, although there are lingering concerns that not all "real" data taking scenarios were fully exercised and that the controlled nature of a test does not fully reflect the more chaotic environment that we will face in just weeks from the time of writing.

Furthermore, a small number of issues were (unsurprisingly) found with the SRM v2.2 implementations, but it has been decided that more production experience is required before attempting to design long–term solutions. For 2008, it is accepted that the small limitations can be worked–around without too much pain.

Finally – and despite the extensive work on ensuring that the underlying services are as resilient we can afford – the challenges exposed some weaknesses of certain computing models. This once again underlines the necessity of operations and application support teams working hand–in–hand with the application community. Generic middleware and / or support can only provide at best mediocre service levels and unoptimized performance.

However, these exercises and the many data and service challenges that preceded them have confirmed that we are able to handle all aspects of the service – from steady state through isolated problems to "show–stoppers" – the latter typically being resolved in a matter of days at worst. This gives us confidence that we are ready to *face* 2008 data taking, although we need to be ready for unforeseen problems and work through them using the well–established methodology that we have built up.

A detailed paper describing all of the results from CCRC'08 is in preparation and will be presented at the 2008 IEEE Nuclear Science Symposium.

May collisions commence!

3. Grid Applications – a Classification

Broadly speaking, applications that need to use a Grid can be categorized as follows:

- Provisioned: the application needs significant quasi–dedicated resources for long periods (years, if not decades – certainly long compared with the useful lifetime of specific computing resources, say 3 – 5 years). Examples of this category include High Energy Physics experiments, astrophysics / astronomy etc;
- Scheduled: large amounts of resources are required for periods much shorter than the above - perhaps days or weeks (again, measured against the yardstick of the lifetime of hardware resources). However, given the required scale, dedicated provisioning for such short periods is excluded. Examples in this category range from time-critical - such as disaster response - to non–critical, that can be scheduled well in advance;
- Opportunistic: the least demanding area which can effectively soak up any under–used resources that are available. Much less time critical that the two previous areas.

Given the ostensibly differing needs of these application areas, can a single shared Grid infrastructure meet their requirements? What are the pros and cons of such an approach, in particular what is the motivation for funding bodies to invest in such a shared infrastructure and why would application communities be willing to use it in this non-exclusive way? Perhaps the most persuasive argument comes from the second category - both those that are time–critical and those that are not. Consider, for example, a disaster response scenario. By definition one cannot know when such an event will take place, even if there might be times of increased risks for certain types of disaster. However, when such an event occurs, time is of the essence and any delay will have a corresponding effect in the response to any such problem. Governments and other organizations involved in disaster response cannot realistically be expected to have the necessary computing infrastructure on hot-standby in case of these hopefully rare events (in contrast to regularly used response teams, such as fire brigades and ambulances). Furthermore, unless the needed infrastructure is regularly exercised, it is likely that it will not be in a usable state on the hopefully rare occasion that it is needed. Further arguments in terms of economy can be found both in the areas of scheduled (non-time-critical) and opportunistic use. Finally, provisioned applications offer an excellent long-term load that can be guaranteed to ensure that the Grid infrastructure is permanently in full production status and evolves continuously to remain state-of-the-art and hence competitive.

It is clear that these applications do not require grids. However, we will argue that grids not only provide a cost–effective solution but also yield other benefits that are of importance to both science and society – surely a big plus.

4. The Need for Reliable Grid Services

In a typical data taking year, the LHC will operate for some 100 – 200 days. Following the operational cycle of previous accelerators, this data taking period may well be split up into "runs" of some 3 weeks interspersed by approximately one week of machine development and testing. Outside of the data taking period, re–processing of the data with improved calibration constants and algorithms, analysis and preparation for the following year's data taking will be carried out in parallel. This leaves very few opportunities for making changes. By deploying the services with a view to resilience, we are able to perform many typical interventions (kernel upgrades, security fixes, middleware upgrades, database patches, hardware replacement) with either zero user–visible downtime or else with very short interruptions. In some cases – e.g. when database schema are modified – we are unable to perform this without an interruption of some 30 minutes, but such interventions are normally limited to once per service per year. The exception today is in the area of storage management – this must clearly be improved and the techniques used for other services are equally valid for these.

However, more significant changes, such as a major change of architecture, are much harder to foresee. History tells us that these are inevitable. Indeed, a comparison with LEP – the previous collider housed in the same 27km tunnel – suggests that major changes can be expected within the coming few years. Some of the likely changes are described in more detail below.

The techniques used to achieve service reliability cover the following areas:
- design;
- implementation;
- deployment;
- operation.

and are described in more detail in (7).

Two mindsets that have proven particularly important in this respect, described in more detail in the above paper, are:
- Think service - a service is far more than a middleware release in a production repository
- Think Grid - a Grid is the ultimate distributed computing system (so far). A change to a service deployed at a given site or site(s) may well have an impact far wider than the local community and must be planned and announced accordingly.

5. All change!

Today's computing enviroments are anything but stable. Hardware is constantly being replaced and newer machines typically require recent operating systems and / or versions. Even without continuous middleware bug fixing and enhancements, this alone would necessitate regular release updates. Some of these changes are much more disruptive than others – the move to 64bit that accompanies Scientific Linux versions 4 and 5 has taken a very long time to follow in terms of middleware releases. A list of changes already foreseen on the 2009 timeframe is given below and more are likely!
- Scientific Linux 5 support;
- Oracle 11g;
- Move to the CREAM Compute Element;
- Use of pilot jobs and new authorization framework;
- Data Management fixes and enhancements motivated by production experience;
- Possible – or even probable – changes in the experiments' computing models, based on experience from first data taking;
- 2009 hardware resources;
- Continued commissioning, 7 + 7 TeV operation;
- The EGEE III transition to a more distributed scenario (in preparation for an eventual move to EGI.)

6. Petascale Computing on the Grid

WLCG originally foresaw the need for 100,000 PCs to perform the offline processing of the data collected at the LHC. This has since been revised to 100,000 of today's cores. Whilst floating point performance is of less importance than one might expect for our workloads – they are more typified by various SPECint benchmarks – the processing power, data volumes and transfer rates (sustained rates of multiple GB per second from Tier0 to Tier1s as well as inter–Tier1s and Tierx–Tiery) surely qualify for the "petascale" label. Thus, we can justifiably claim a production demonstration of petascale computing on the grid – something that is more usually associated with state–of–the–art supercomputer installations. This is a success not only of grids in general but also opens the door from individual institutes that make up the grid to petascale resources.

7. Grids versus Clouds versus Supercomputers

One of the important characteristics of today's grids is that they permit resources at an institute collaborating in a large project to be used both for the benefit for local users as well as for the overall collaboration. As such, sites participating in WLCG – be they Tier0, Tier1 or Tier2 – all have a crucial role to play. The value that they provide can be seen from the fact that all reprocessing is done at Tier1 sites – and all physics results can be expected to come from reprocessed data, particularly in the early years - whereas Tier2 sites are responsible for generating simulated data and (typically) also providing resources for end-user analysis. Whilst there will be Tier1 and Tier2–like activities also at CERN, it is today impossible to envisage a scenario where the external institutes do not play a key role. In fact, the "one-thirds" (at CERN) "two-thirds" (outside) rule for computing resources has been in place since well prior to the startup of the LHC's predecessor, LEP. (In WLCG, the sum of resources at each level is approximately constant, thus the sum of Tier2 + Tier1 resources follows the above-mentioned rule).

The fact that funding bodies can invest locally and that researchers from all participating institutes can play a key role from their home institute adds value to these institutes and stimulates the local communities. It is clearly much more attractive than "CERNtralizing" all such activities, which in turn would have a negative impact on external institutes and would make fair returns to the various funding bodies much more of a challenge (but not impossible, as this is followed for all CERN purchases and contracts as a fair returns policy).

At least with today's commercial cloud offerings, it is hard to see how such benefits could be provided to local communities. Score 1 for the grid.

From a more technical point of view - but not yet proven - it is hard to understand how data intensive applications, such as those that exploit today's production grid infrastructures, could achieve adequate performance through the very high-level interfaces that are exposed in clouds. In both the case of data management and databases deep knowledge and direct control of the implementation is required to get adequate performance. Indeed, data management been one of the most turbulent areas in WLCG for at least the past 5 years for precisely this reason (with the deployment of important new functionality at the data management layer only proven to be production-ready in CCRC'08) and this can be expected to continue in the future.

On the other hand, should the higher level interfaces actually work for applications as demanding as those of the LHC, this would definitely be a big plus for cloud computing. It would, however, beg the question as to why such interfaces were not also possible in grids.

Another sanity check can be made with supercomputers. At least today – as reported at CCP 2008 – expert help is required to fully exploit the capabilities of these systems. This is true also in the case of grids, so "equality". However, the lower cost of acquistion of grid systems, coupled with the advantages listed above should allow grids to gain back the lead. Indeed, many communities feel - and most probably are - excluded from areas of research that require supercomputing resources. It has been proven that this is not the case with grids.

So the challenge for grids remains: to demonstrate low cost of entry (helping to attract new application communities), as well as low cost of ownership (to retain their competitive edge over supercomputers and other challengers).

8. The Requirement for Continuity and Stability

With the LHC entering operation in 2008, discussions have started concerning the LHC luminosity upgrade programme, which will in effect be a refreshement (read replacement) of much of the accelerator infrastructure at CERN. This will eventually lead to an increase in luminosity of around 100 fold around the year 2018. One can therefore predict that the LHC (or "super LHC") will continue to take data well into the following decade. Given the accelerator cycle described above, there is clearly a very long term need for computing services, today provided via grids. The LHC, however, is but one application community that relies on this infrastructure. Although the largest single community (in terms of resource usage) other communities summed together reach at least the same level of usage at peak periods. If grids indeed pay off on their now–proven promise(s), this can only be expected to grow. In the not–too–distant future, High Energy Physics might well be a minority.

However, user communities will only be attracted, and will only remain, if the cost of entry and usage (or ownership, from a service provider point of view) are sufficiently low with respect to alternative solutions.

In the preparation phase of the LHC, CERN migrated several hundred TB of data from one data management solution to another – admitedly at significant cost – but will not be afraid to move to a new solution if it proves sufficiently attractive, or is required for other reasons (an existing solution becoming de–supported, for example).

9. The European Grid Initiative Design Study

The European Grid Initiative (EGI) Design Study (EGI_DS) is a 27–month project whose aim is to establish a long–term sustainable e–infrastructure that follows on from the current projects that are typically funded on a 2 year cycle.

According to the "blueprint" document currently being produced, the goal of such an infrastructure is "to realize a large-scale, production Grid infrastructure for sharing of IT resources and data - built on National Grids that interoperate seamlessly at many levels, offering reliable services to a wide range of applications, ranging from mission critical to prototyping and research." The timescales involved are rather aggressive – the future EGI needs to be established early enough so that it can be fully operational when the current projects end. EGEE III will terminate in April 2010 and a minimum overlap of some 6 months will be required. Whilst it is strongly hoped that this project will not only be established on time but will also provide the functionality required by WLCG, it is clear that to exploit the LHC machine in 2010 and beyond a fully functional grid infrastructure must be available and that the project must take steps to ensure that this is the case. It is important to note that WLCG and VO (ATLAS, CMS, ALICE, LHCb) operations represent a significant layer on top of those provided by the general infrastructure and that a number of the key middleware components are maintained by CERN or other High Energy Physics laboratories / national infrastructures (DESY, FNAL, INFN). Thus, to extend the WLCG operations to cover those currently provided by EGEE III infrastructures would be a relatively minor, albeit undesirable, step.

10. WLCG Operations

Whilst WLCG operations relies heavily – by construction – on the operations, application support, middleware development, certification, testing, release and other functions of underlying grids such as EGEE, these are insufficient. As mentioned above WLCG holds daily operations conference calls that act as a focal point for all short term issues and provide a mechanism for exchange of information between sites, service providers and the experiments. These calls typically last around 10 to 15 minutes and notes are produced and circulated the same business day. These notes are used to provide a weekly service report to the WLCG Management Board and as input to the quarterly reports used for higher-level management. At least once per year a "Collaboration Workshop" is held, with attendance typically of around 200 people. These are complemented by topical workshops sometimes held at CERN but whereever possible in conjunction with another event (e.g. Open Grid Forum, EGEE annual conference, EGEE Users' Forum, Computing in High Energy Physics conference).

In addition, the experiments hold their own workshops, which can also attract well over 100 people!

It is simply unrealistic to believe that an infrastructure project – as EGEE has been and as EGI is hoped to be – can supplant application community specific needs. This can clearly be seen when comparing against other infrastructure providers, such as telecoms, networking or even rail track!

11. WLCG S.W.O.T. Analysis

As part of the final review of the outcome of CCRC'08, a "Strengths, Weaknesses, Opportunies and Threats" analysis was carried out. The majority of this analysis has already been reported on above. What remains is primarily the threats. The largest of these is felt to be the risk of falling back into “fire-fighting mode” at the first sign of serious problems. It is hoped that this can be avoided by the continuous communications established not only through the daily operations meetings but also by the extensive applications support that has been successfully used over the past years. Problems are a certainty, but we also know that some sort of solution, possibly not perfect from a computer science point of view, but adequate from that of our priority – extracting as much science from this world–class machine as efficiently and rapidly as possible – will surely be found.

12. Summary

We have shown how several heterogenous grids have been used to satisfy the production needs of a number of High Energy Physics experiments – well known for being particularly demanding in terms of computational systems. We have aruged how such a solution provides additional and significant benefits to scientific communities at all institutes participating in these projects and how this adds value not only to these insistutes but also well beyond. The application and operations support requirements have been briefly compared with alternatives, showing that grid solutions can at least be competitive in these respects. We have suggested how application communities with widely differing needs can co–exist to their mutual benefit.

A question raised at CCP 2006 – as to whether grids would cross the chasm from niche usage to ubiquity – remains unanswered.

The challenge that remains is for a sustainable grid “e–infrastructure” to build on these not insignificant successes and hopefully provide a solution that not only lives as long as user communities, such as those from the LHC, but broadens out into other areas of science and hopefully beyond.

Lightening never strikes twice?

Acknowledgements

The work of the many people involved in the LCG projects and the underlying Grids (OSG, EGEE and NorduGrid) is gratefully acknowledged.

In addition to the Tier0 situated at CERN, WLCG Tier1 sites are ASGC, Taiwan; GridKA, Karlsruhe, Germany; CNAF, Bologna, Italy; IN2P3, Lyon, France; NIKHEF/SARA, Amsterdam, Netherlands; NDGF, Nordic Countries; RAL, Didcot, UK; PIC, Barcelona, Spain; BNL, Long Island, USA; FNAL, Batavia, USA; TRIUMF, Vancouver, Canada.

The WLCG service coordination role is currently by members of the Grid Support in CERN’s IT department, in close collaboration other “physics groups“.

References

[1] The Worldwide LHC Computing Grid (WLCG) – in the proceedings of the Conference on Computational Physics 2006 (CCP 2006), volume 177 (2007) 219 - 223.

[2] The European Grid Initiative Design Study – website at http://web.eu-egi.eu/.

[3] CERN webpage with information on LHC startup – http://lhc-first-beam.web.cern.ch/lhc-first-beam/Welcome.html.

[4] The Enabling Grids for E-sciencE (EGEE) project – http://public.eu–egee.org/

[5] The Open Science Grid – http://www.opensciencegrid.org/

[6] The Grid middleware development and testbed deployment project in the Nordic countries – http://www.nordugrid.org/

[7] Robust and Resilient Services - How to design, build and operate them, in the Proceedings of the Third Conference of the EELA Project, Jamie Shiers, Gavin McCance, Patricia Mendez Lorenzo.