

Slide 1 – Real Time Monitor snapshot

As we have seen, the LHC accelerator creates extremely high energy particle collisions, which in turn create new particles that decay in very complex ways as they move through the detector. The detector registers the passage of these particles with a vast number of sensors and, finally, a digitised summary of this is recorded as what we call an “event”.

Physicists now face the challenge of unravelling all of this complexity in order to extract the science, and for this they need computers – lots of them.

In this talk, I hope to give an overview of the computing system we have put in place to support LHC data analysis: the challenges we faced, and the computational structures and human collaborations we built to overcome these challenges.

Slide 2 – LHC Computing Challenge

When we first started designing a computing system to analyze the vast amounts of data that we knew would be coming from the LHC, we were faced with three major challenges: (1) the scale and complexity of the data, (2) the size of the computing facility that would be needed to analyse it, and (3) the need for a service that can evolve rapidly as new discoveries change the way in which the data is accessed.

The volume of data produced at the LHC will be hundreds of times greater than in previous accelerators. This is because the detectors of the experiments have many more sensors, the “events” that they record at the LHC’s energy scale are much more complex, and the rate at which interesting collisions occur is much higher. Altogether, we expect that 15 PetaBytes of new data – 15 Million GigaBytes, or enough data to fill 20 million CDs – will be generated each year, and all of this must be carefully managed and made readily accessible to all of the users.

With 7,000 LHC physicists actively analysing the data, a lot of computing power is also needed. We are planning for about 100,000 processors and about 45 PetaBytes of disk storage just for the first full year of running.

Because this is a research programme – involving a new energy level, new detectors, and new physics – we do not know exactly how much capacity will be required, nor do we know how the data will be accessed.

What we did know when designing the computing system was that it would have to be very flexible – able to develop smoothly over a period of at least ten years, during which we would be continually adding capacity, improving performance, and absorbing new computing technologies, while striving always to maintain excellent access to the data.

Slide 3 – HEP Cluster Architecture

We had a good basic architecture to start from, that had been developed for experiments at CERN’s previous accelerator. This had proven itself well over the years in terms of scalability and the ability to migrate quickly to new, more cost effective, computing technologies. A similar architecture was being used in all of the high energy physics computing centres, and we were confident that this would scale up to LHC levels. But there was one other problem that we had to deal with. It quickly

became clear that the overall capacity required for the initial four experiments was far beyond the funding that would be available for computing at CERN. At the same time, we knew that most of the laboratories and universities that were collaborating in the experiments had access to national or regional computing facilities. And so, the obvious question was: “Could these facilities in some way be seamlessly integrated with CERN to provide a single computing service for the LHC?”

Slide 4 – Grid

Consequently, we decided to implement a distributed system as a *computational grid*, based on the ideas of two scientists working in the United States, Ian Foster and Carl Kesselman. Together, they had developed a concept which allowed computing centres to inter-connect in a very general way, integrating their separate resources to offer a single virtual computing service.

On top of this, basic service users could be grouped together into virtual organisations, which gave all members of the group the same rights to use the group’s resources at any of the grid sites.

Layers of software hide all of this complexity from the end user – so the physicists at one experiment only see a single service, enabling them to concentrate on their analysis without being troubled by the details. This meant that they did not have to be concerned about where the data is located, where the computational capacity is available, and how to authenticate themselves in order to use the resources at what has now grown to be one hundred and forty independently managed computer centres. Instead, physicists were free to think about . . . physics.

Slide 5 – Service Hierarchy

On top of this “grid” we then defined the LHC Computing Service architecture.

- The sites are organised as a hierarchy of layers, or *tiers*, depending on the services that they provide. CERN, as the “Tier-0”, is the top layer, performing initial processing of the data, and maintaining master copies of the raw datasets and other key information, pushing the data out rapidly to
- the next layer – the “Tier-1” centres. These are eleven large data-intensive centres, characterised by major investments in mass storage services, round-the-clock operation, and excellent network connectivity. They provide for long-term data preservation, hold synchronised copies of the master catalogues, are used for data-intensive analysis, and act as data servers for the smaller centres.
- At the level of the end-user analysis – the heart of the physics discovery process – are about 130 “Tier-2” centres. While these do not have to make the same level of commitment in terms of data and storage management, they must continually adapt their configurations to support the evolving demands of their client physics groups. The expectation is that the sheer diversity among the numerous Tier-2s will act as a stimulus to producing novel approaches to analysis.

Slide 6 – Collaboration

When we started to design the computing service for the LHC, the idea of a computational grid was still something new, and so in order to build a grid for LHC we had to collaborate closely with

projects, particularly in Europe and the USA, that were developing “middleware” – the basic software that makes the grid work. This was a very close collaboration – as many of the grid projects were initiated by and led by LHC collaborators.

We also had to learn how to manage and operate a service that interconnects more than a hundred independent computer centres, most of which also serve other communities – in physics and in other sciences. To integrate all of these centres into the grid we had to resolve many technical issues, navigating between competitive solutions, and work out flexible procedures addressing sensitive policy issues in areas including authentication, resource allocation and security.

Managing the distribution of the data across the grid was also a major challenge. Reliable, high-performance data-transfer between sites is only one of the issues. Each experiment has many millions of datasets to keep track of, and large disk storage resources spread across the grid to manage. The data management systems that have been developed to suit the specific needs of each experiment are complex but crucial components to managing a data grid.

Achieving agreement on so many details in an environment involving so many organisations and people with so many diverse backgrounds was itself a major challenge.

And of course – we would not have been able to contemplate building a grid without high-performance and ubiquitous networks.

Slide 7 – WLCG Site Map

With the first beams having circulated in the LHC machine, we can say that the primary goals for the worldwide distributed computing service for LHC data analysis have been met. Groups can pool their computing resources wherever they are located, large centres and small centres can all contribute, and users everywhere can get equal access to data and computation without having to spend all their time seeking out resources. At the start of LHC operation the LHC grid is operating across 140 computer centres in 33 countries, with more than 80% of computing resources located outside of CERN.

We have established a high degree of flexibility for growing the service and integrating new resources, wherever they are provided.

The Worldwide LHC Computing Grid has already begun to handle data from the first beams at LHC, and it is ready to handle the flood of data that will emerge over the coming months.

In summary, we have the beams, the detectors, and the computational resources all in place – the full suite of tools we need -- to embark upon some exciting physics. We now have the capacity to peer deeper into the nature of matter than anyone before, and – in the words of John Ellis, a theoretical physicist here at CERN -- find “weird and wonderful creatures that no one ever even thought of.”