annual report 01

to sponsors

CERN

openlab for DataGrid applications

Developing Solutions for the Data-Intensive Science of the Large Hadron Collider

CERN
openlab for DataGrid applications

CERN openlab for DataGrid Applications

http://www.cern.ch/openlab

Annual Report 01 to Sponsors

Executive summary 3

CERN: A world laboratory for studying Nature and a technological powerhouse 5

LHC - the Large Hadron Collider 7

LCG - building the LHC computing environment 9

CERN openlab for DataGrid Applications: the industrial dimension 11

Where do sponsor contributions fit in? 13

Achievements with individual corporate sponsors 17

Looking ahead 23

This document is the Annual Report to the Sponsors of the CERN openlab for DataGrid Applications for the year 2001.

Note added in press:
This report summarizes the status and plans discussed with sponsors in March 2002.
Subsequently, KPNWEST has filed bankrupcy.
CERN is working to minimize the impact of this event on planned activities and continues enthusiastically its work with the remaining sponsors.
CERN also continues to seek additional sponsors to enhance the CERN openlab for DataGrid Applications.

Prof. Manuel Delfino
June 2002

executive summary

CERN is an intergovernmental organization, funded by twenty European states, with the mission to foster international collaboration in the field of fundamental research in particle physics. The CERN research programme involves over 6000 researchers from some 500 research centres and universities all over the world, making it the world's leading particle physics research laboratory.

CERN is in the process of building the *Large Hadron Collider (LHC)*, a large particle accelerator operating at an energy level much higher than with today's machines. Four detectors are being built, which will generate a huge amount of data to be analysed by physicists all around the world.

To construct the LHC and the infrastructure to exploit it, CERN is facing tremendous technological challenges. One of the biggest is the building of a global computing environment, the *LHC Computing Grid (LCG)*, to meet performance requirements far beyond any in the past: 11 petabytes of data to be recorded per year, a total processing power equivalent to that of 200,000 of today's fastest processors.

The development, prototyping and deployment of the LCG are organized as an inter-institutional project coordinated by CERN. CERN organizes its technical participation in the LCG through a set of activities entitled *CERN openlab for DataGrid Applications*, of which interested parties are invited to become sponsors. The CERN openlab for DataGrid Applications has many distinguishable features: it is building a concrete system with defined performance and strict timelines, based on open standards and leading to a production system, the complexity and scale of which are unique.

The LHC Computing Grid system is formed of technological components that constitute the building blocks of the LCG architecture:

1. The Computing Systems, assemblies of highly performing CPUs, storage systems, local and wide area networks, data management systems

2. The Grid Technologies, mainly formed of a series of cooperating software systems called Grid middleware.

3. The Application software, the collection of common frameworks, client software, user interfaces, which are directly used by physicists.

A first step for the CERN openlab for DataGrid Applications was the appropriate positioning of the sponsor contributions in the LCG architectures, in order to ensure the overall project consistency. These contributions, all belonging to the Computing Systems building block, concern three key areas of technology and services: very-high-speed networking equipment, wide area network services and processor technology.

• **Enterasys Networks** is providing a *Ten-Gigabit Ethernet infrastructure* to form the LAN intra-fabric network at CERN.

• **Intel Corporation** is providing Itanium® *2 family processors* for the servers of the LCG fabric.

• **KPNQwest** is providing *gigabit-per-second long-haul network* services to connect CERN systems to one or several regional centres.

The second step has been the exploration with the individual sponsors of the details of their contribution (and the design where appropriate of possible phased scenarios), leading to a work plan and timelines. The third step has been, when possible, the implementation of the initial phase of the plan.

• Enterasys Networks is providing a Ten-Gigabit Ethernet infrastructure to serve as the LAN of the LCG fabric at CERN. The network being constructed is formed of concentric layers of increasing switching and bit rate power: *Ten-Gigabit Ethernet routers* in the core, *E1 Gigabit switches* in the intermediate layer, and *Gigabit Ethernet* switches in the outer aggregation layer.
A first implementation has been completed, with a Ten-Gigabit Ethernet link between two E1 switches, interconnecting 120 computers. The plan includes the extension of the set-up to 240 computers by May 2002 and more than 500 by the end of the year.

• Intel Corporation is providing Itanium® *2* family processors for the servers of the LCG fabric at CERN. The exploration of the 64-bit processor technology started at CERN before the CERN openlab for DataGrid Application activities. Several lessons were learned from these tests, including the importance of compiler technologies to take maximum advantage of the parallel architecture of these processors.
Intel's compiler experts and CERN will work together to ensure that LCG applications become available on successive versions of the Itanium® *2* family processors. The CERN openlab for DataGrid Applications is in the process of identifying an industrial assembler to integrate them into LCG fabric servers.

• KPNQwest is providing high-speed long-haul services between CERN and regional centres. Three scenarios have been devised with various levels of technology innovation and differing end points.
Scenario 1: One or more 2.5 (or 10) Gbps links to Tier1 centres.
Scenario 2: Lambda connections to two European regional centres.
Scenario 3: Dark fibre to the closest regional centre -IN2P3 in Lyon.
These scenarios are not mutually exclusive, and indeed may serve as a basis for constructing a compound scenario. The decision on the chosen scenario should be made in April 2002 and the first connection could be delivered in September 2002.

In summary, the collaborative work with sponsors has started and solid results have already been achieved. Though the various aspects of the work are progressing at differing rates due to the nature of the contributions and their time-critical positioning in the overall work plan, the achievements so far are extremely promising.

The analysis of the project in terms of technological building blocks has demonstrated the key position of the technologies brought by the three founding sponsors. It also shows additional opportunities in a number of technological areas for corporate partners to contribute to the CERN openlab for DataGrid Applications, and join the world's largest data-intensive challenge.

CERN: A world laboratory for studying Nature and a technological powerhouse

CERN, the European Organization for Nuclear Research, is one of the world's largest and most respected international centres for scientific research. Its core business is fundamental physics - finding out what the Universe is made of and how it works.

CERN is unique amongst physics research laboratories in that it is an intergovernmental organization with 20 Member States . CERN's highly skilled staff of about 2500 people work side by side with almost 6000 scientists and engineers from collaborating universities and research laboratories towards the common goal of building the world's largest particle accelerators and detectors. These complex scientific machines are used to study Nature's ultimate building blocks - the fundamental particles. By studying what happens when these minute fragments of matter collide, physicists unravel the basic laws of Nature. In order to achieve these scientific marvels, CERN and its partners often have to push existing technologies to their limit or participate in exciting new developments.

When CERN was founded in 1954, it was one of Europe's first joint ventures, with specialists working together to build scientific research installations that were beyond the means of individual nations. It has gone on to become a shining example of international collaboration - a laboratory for the world. From the original 12 signatories of the convention establishing CERN, membership has grown to 20 European Member States[1], and nations from all over the globe contribute to and participate in the research programme.

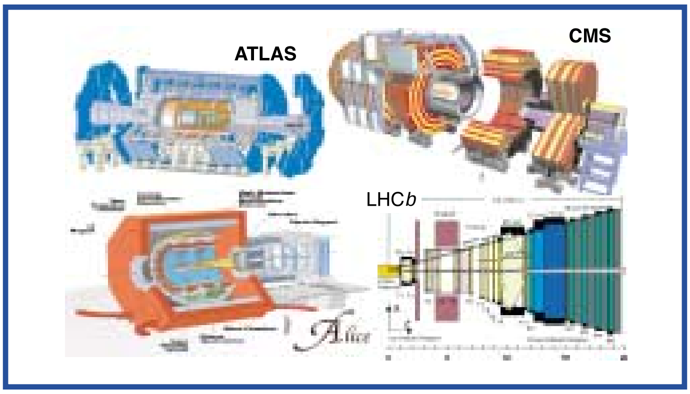

The Large Hadron Collider project: the four detectors

1 At present, CERN Member States are Austria, Belgium, Bulgaria, the Czech Republic, Denmark, Finland, France, Germany, Greece, Hungary, Italy, Netherlands, Norway, Poland, Portugal, Slovakia, Spain, Sweden, Switzerland and the United Kingdom. Israel, Japan, the Russian Federation, the United States of America, Turkey, the European Commission and Unesco have observer status.

CERN builds accelerators

CERN's machines are particle accelerators and detectors, and CERN's users are physicists. The Laboratory's oldest still-functioning accelerator is the Proton Synchrotron (PS), commissioned in 1959. The Super Proton Synchrotron (SPS) was commissioned in 1976, when it was the world's largest particle accelerator. In the early 1980s it was the scene of research that brought CERN its first Nobel Prize (1984). The LEP electron-positron collider, built in a 27-kilometre underground tunnel, was the Laboratory's flagship research machine from 1989 to 2000.

CERN's successive generations of machines attain higher energies, each machine feeding the next. The LEP ring has now been dismantled, and in its tunnel CERN is building the Large Hadron Collider (LHC) to enable the world's particle physicists to probe even deeper into the heart of matter.

Physicists build detectors

While CERN provides these giant machines, physicists in universities and national research centres build the giant detectors which record what happens when the particles collide or interact. As physics has advanced, experiments have become progressively larger, and a major collaboration can now include more than a thousand scientists. Altogether, the CERN research programme involves over 6000 researchers from some 500 research centres and universities all over the world.

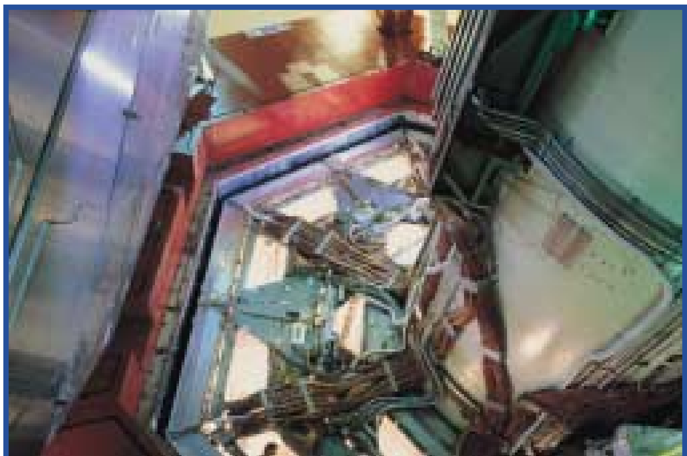

The technology transfer dimension

Fundamental research is CERN's reason for being, but the Laboratory also plays a vital role in developing the technologies of tomorrow and bringing important benefits to society. Techniques developed for particle detectors have been successfully adapted for medical imaging and for security scanning. The best known of CERN's spin-offs is the World Wide Web, invented in 1989 in response to the need for new ways to easily share information amongst the members of CERN's community.

But CERN has made numerous other contributions that have helped to advance technology to the benefit of society, in particular in the field of Information Technology where CERN has been a consistent innovator. CERN played a key pioneering role in the development of the Internet Infrastructure in Europe, and became in the early 1990s the biggest Internet hub in Europe, aggregating up to 80% of the total international Internet transmission capacity. In 2002, the CERN Internet hub that interconnects dozens of institutional and industrial network operators is a vital component of the European Internet landscape. CERN also played a pioneering role in High Throughput Computing for over a decade. In the early 1990s, CERN and physicists from many of CERN's Member States developed a revolutionary computing architecture called 'SHIFT', which allowed multiple storage and processors to interact over high-performance networks. For this achievement, CERN received in 2001 the prestigious Twenty-First Century Achievement Award from the Computerworld Honours Programme.

From material sciences to computing, particle physics demands the ultimate in performance, making CERN a unique test-bed for industry.

LHC, the Large

The Large Hadron Collider is a particle accelerator which will probe deeper into matter than ever before. It will ultimately collide beams of protons at an energy of 14 TeV, 300 times higher than at CERN's present day experiments.

Why the LHC?

The LHC is the next step in a voyage of discovery which began a century ago. Back then, scientists had just discovered all kinds of mysterious rays: X-rays, cathode rays. Where did they come from? Were they all made of the same thing?
These questions have now been answered, and have changed our daily life, giving us television, transistors, medical imaging devices and computers.
At the dawn of the twenty-first century, we face new questions: Why do elementary particles have a mass, and why are their masses different? Are forces in Nature really just manifestations of a single force? The LHC is designed to address these questions. Who can tell what new development the answers may bring.

The LHC accelerator technology challenges

To keep the LHC's beam on track needs stronger magnetic fields than have ever been used in a CERN accelerator. Superconductivity – the ability of certain materials to conduct electricity without resistance or energy loss – makes such fields possible, but a superconducting installation as large as the LHC has never before been built. The LHC will operate at about 300 degrees below room temperature - even colder than outer space. With its 27-km circumference, the accelerator will be the largest superconducting installation in the world.

The LHC detector technology challenges

To study what happens when the LHC's beams collide, four giant detectors are being built: ALICE, ATLAS, CMS and LHCb. These detectors will handle as much information as the entire European telecommunications network does today. Installed in caverns, they will stand up to 20 metres high, as tall as office blocks, and bristle with electronics.

At the centre of the detectors, protons will collide some 800 million times a second, with a combined data volume of more than 40 million megabytes. To reduce the flow of data to manageable levels, specialized multi-level computing systems, or ‘triggers’, will be used to select events with promising characteristics.

Typically, in the case of ATLAS, the level-1 trigger is a massively parallel system of specialized electronics. Of 40 million bunch crossings per second, fewer than 100,000 pass level-1. The level-2 trigger is a system of programmable processors from which fewer than 1000 events per second are sent to the third stage, the event filter. The event filter is a large farm of CPUs, which perform a detailed analysis of the full event data. Less than 100 events per second remain and are sent to the data storage system to be recorded for offline analysis.

The LHC computing challenges

The LHC data storage and analysis presents tremendous challenges in computing power, data storage and software technology.

For all four experiments, an aggregate of five to eight petabytes (five to eight million gigabytes) of data will be recorded and accumulated each year, to be shared and analysed by physicists around the world. To this, must be added the requirements for reconstructed and simulated data, and all the calibration information, leading to a total mass storage in excess of 11 petabytes per year for all experiments. This corresponds to the storage volume of 16 million of today's CD-ROMs.

The analysis process will require 10 petabytes of data to be recorded on disks, as well as a total processing power equivalent to that provided by 200,000 of today's fastest processors. Using the latest generation of processors may allow to reduce this number to about 20,000, which will need to be interconnected by a local area network operating at ten gigabits per second - one order of magnitude faster than today's local area networks.

Between CERN and regional processing centres, long-haul international connections will be required with an aggregate bandwidth exceeding ten gigabits per second, that is a capacity equivalent to that needed for transmitting 5,000 digitized movies of TV-quality simultaneously.

All these computing and networking requirements are far beyond the demand of any project ever undertaken.

LCG, building the

The computing organizational model

The development, prototyping and deployment of the computing environment necessary for the LHC is being carried out as a project that includes many scientific institutes and industrial partners. The project is coordinated by CERN and nicknamed LCG (after LHC Computing Grid).

The computing environment has been modelled from a multi-tier hierarchical architecture where centres geographically distributed all over the world cooperate to provide the necessary storage and analysis power.

In this model, for each experiment, raw data storage and reconstruction would be carried out at a Tier0 centre. Analysis, storage and generation of simulated data as well as data distribution, would be the task of several (four to five) national or supra-national Tier1 centres. The Tier1 centres would be followed by a number (about 15) of national or infra-national Tier2 centres, which in turn are followed by institutional Tier3 centres (about 50). Finally, the physicists operate from their end-user workstation (Tier4).

CERN will host the Tier0 as well as one Tier1 centre, which together will represent about one third of the total computing capacity. Though they form two distinct logical entities, the CERN Tier0 and Tier1 will share a single physical, partitionable facility.

This multi-tier model, when combined with appropriate security, authentication and data management mechanisms, results in a distributed integrated system which easily maps to the innovative technology known as Grid.

The organizational model of the LHC computing

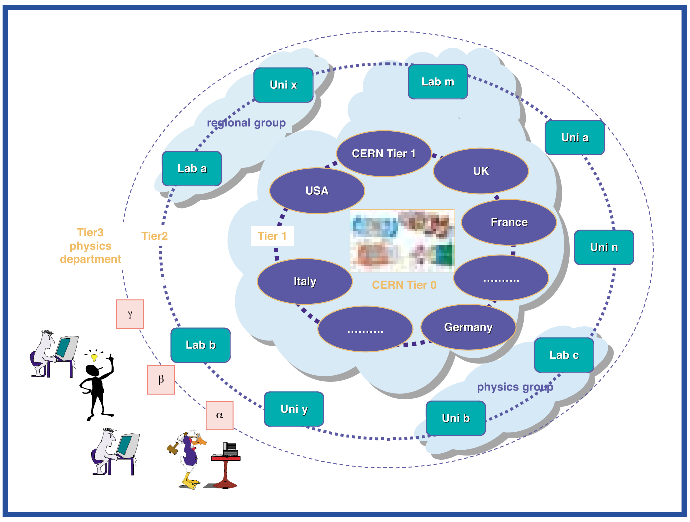

Grid: the technology for the LHC computing environment

The overall worldwide computing system for LHC that the LCG project is developing will be implemented as a global computational Grid, with the goal of integrating large geographically distributed computing fabrics into a virtual computing environment. The Grid approach may be summarized as distributed computing on a global scale between numerous, heterogeneous organizations.

The LCG project is integrated with several European national computational Grid activities and it collaborates closely with other projects involved in advanced grid technology and high - performance wide area networking, such as DataGrid and DataTAG, two projects funded by the European Union and coordinated by CERN.

There are numerous challenging problems to be tackled in many areas to build the global computational Grid of the LHC, which will involve tens of thousands of computers. They include distributed scientific applications, computational Grid middleware, automated computer system management, high-performance networking, object database management, security and global Grid operation.

Fostering the next generation of computer engineers

The LCG project will be carried out by cooperating teams all over the world. At CERN alone, the required human resources for the LCG development activity amount in 2002 to nearly 100 people for the pure development activities of the project. To this, must be added another 50 people for equipment acquisition, installation and service operation of the supporting services. The development activities cover the most advanced fields of Information Technology and require an ambitious recruitment programme: every year, young computer scientists and engineers will be hired, the aggregated number of required newcomers reaching 50 in 2006. By involving those young people in the world's largest data intensive challenge, to work in teams gathering some of Europe's best researchers, the LCG will foster a new generation of computer engineers in Europe and the world.
This will be made possible through special agreements with Member and Observer States.

Origins

The LCG activities at CERN are taking place in the context of the CERN openlab for DataGrid Applications - the collaboration between CERN and industrial partners to develop data-intensive Grid technologies to be used by the worldwide community of scientists working at the LHC.

The CERN openlab for DataGrid Applications stems from the initial observation that though our complex global world has real needs for systems where many technologies are integrated to provide high functionality (e.g. air traffic control, integrated hospital care management, electronic commerce), it is often easier to start creative work on advanced integrated systems by focusing collaborative work on a challenging focal point in the academic and research environment.

Thus, the idea to create a *collaborative forum* with both the *public sector and industry* focused on solving a *well-defined problem* through *open integration of technologies.*

CERN's environment periodically provides well-defined complex problems where it is difficult or impossible to achieve appropriate performance and an affordable price through existing solutions. As a result, CERN is uniquely positioned to provide such a focal place for open integration in several areas of technology. The concrete problem of solving the LHC Computing Challenge provides an exciting focal point for advanced development of integrated distributed computing systems, and of data-intensive Grid computing in particular.

Key features

CERN invites parties interested in testing and integrating Grid-related systems, products or services to become sponsors of the CERN openlab for DataGrid Applications and carry out large-scale, high-performance evaluation of their solutions in an advanced integrated environment.

The CERN openlab for DataGrid Applications has several key characteristics which distinguish it from many other initiatives.

- The objective of building a concrete system with defined performance requirements, respecting strict timelines, as the LCG project is one key component of a wider undertaking: the LHC project.
- The nature of the outcome of the LCG project, which is a production system, not a demonstrator to be dismantled at the end of the exercise.
- The performance requirements, beyond those of any other project.
- The complexity: the wide spectrum of technologies covered.
- The scale and the global dimension of the project.
- The preference given to open standards.

for DataGrid Applications:
industrial dimension

Where do sponsor

A major aspect of the work carried out during the first phase of the CERN openlab for DataGrid Applications was the appropriate positioning of the sponsor contributions within the overall LCG architecture.

Technological components of the LCG

The LHC Computing Grid system is formed of technological components that constitute the building blocks of the LCG architecture.
• **Computing systems** - assemblies of high-performance CPUs, storage systems, local and wide area networks, database and cluster management systems.
• **Grid technologies** - mainly formed of a series of cooperating software systems called Grid middleware.
•**Application software** - the collection of common frameworks, client software, graphical user interfaces, and application programming interfaces which run on end-user workstations or are directly used by physicists.

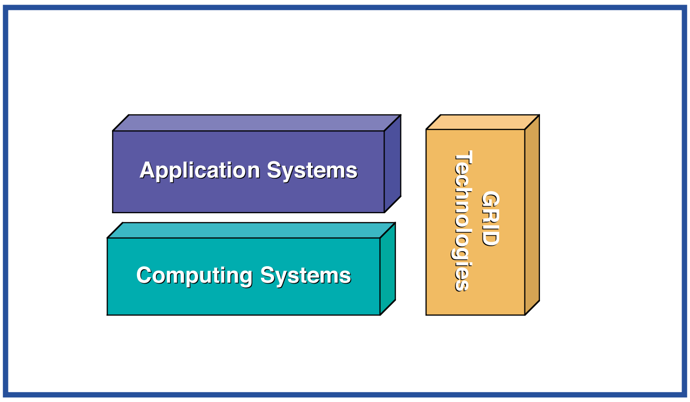

The key technological components of the LCG architecture

Computing systems and LCG fabrics

The computing systems form the core of the LCG environment. They include *LCG fabrics* interconnected with wide area network connections.

LCG fabrics are giant clusters of *CPUs, mass storage systems* and *disks*, glued together by an ultra-high bit rate *intra-fabric network*. These fabrics, of varying degrees of complexity and computing power, are distributed at the nodes of the LHC computational Grid, that is Tier0, 1, 2, and 3.

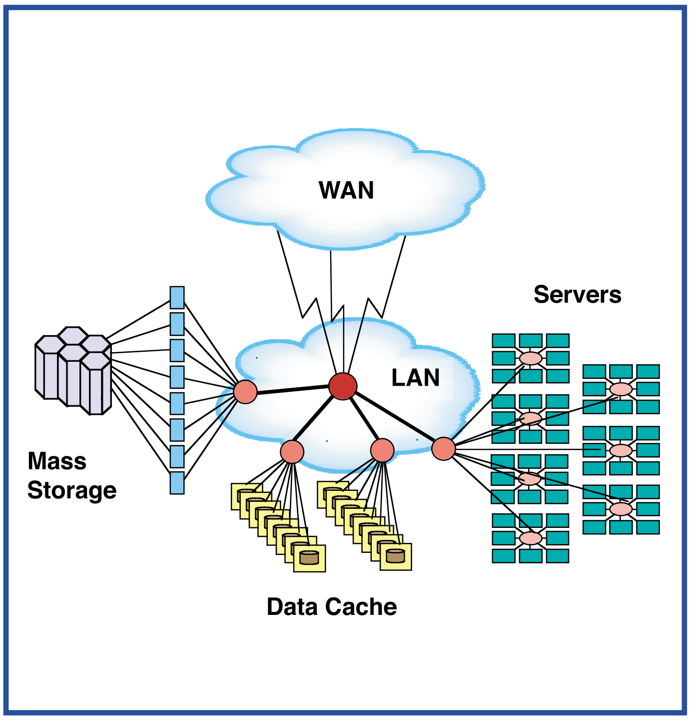

Anatomy of an LCG fabric

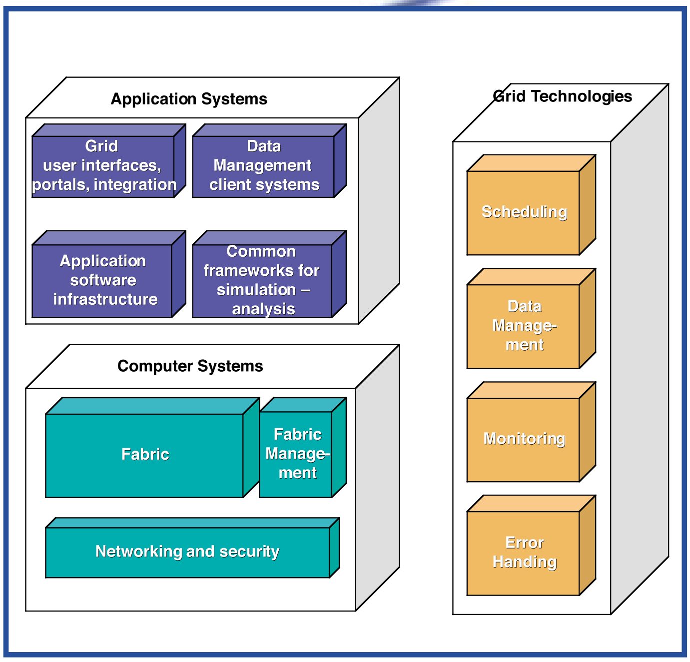

Details of technology building blocks of the LCG architecture

Grid technologies

The Grid technologies are essentially formed of the Grid middleware: a series of software systems that collectively provide a seamless computational Grid based on geographically distributed fabrics. These software modules are called middleware because they are not used directly by end-users, nor do they lie at the base level of operating system or networking. They provide essential – though largely end-user transparent – services including scheduling, data management, monitoring, and error detection and recovery.

Application software

All software of the LHC will be built using modern programming techniques, namely object and component approaches. Within this context, a basic environment for physics *software development* must be built, including scientific libraries, software development tools and monitoring tools.
In addition, a *common framework* and *toolkits* for simulation and analysis must be developed, as well as client systems for *data management*. Finally, end-system *Grid tools* and *portals* must be developed and physics applications must be adapted to the Grid environment.

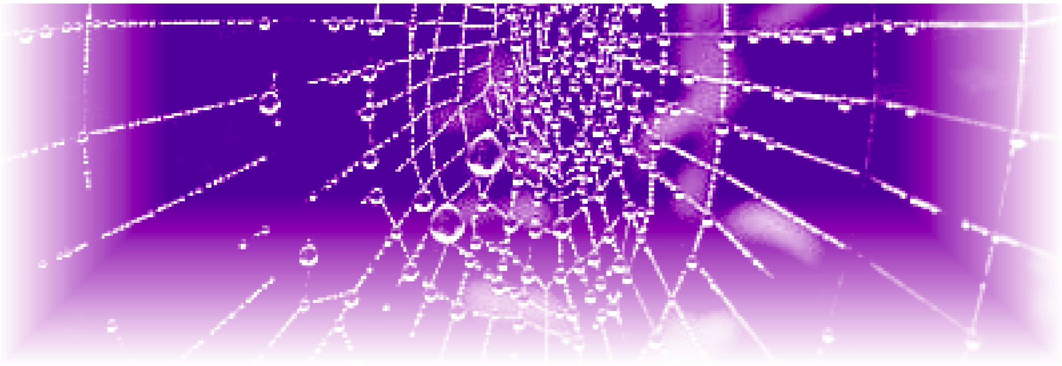

Structuring the project into areas of work

The LCG project has been structured around four areas:

• Three areas are mapped to the three key technological building blocks. The work carried out consists in developing, assembling, testing and validating the required technologies.

• A fourth area, of a different nature, will be devoted to the actual progressive, worldwide deployment of the LHC Grid, leading to a global production computational environment.

The computing system area will develop the physical data storage system, select the appropriate data management system, install fabrics and develop their management system, install and evaluate high-performance intra-fabric and inter-fabric networking, select and implement appropriate security and authentication systems.

The applications area will provide an infrastructure for the development of physics software including compiler expertise, development tools and standard libraries. It will also develop a common framework for simulation and analysis and provide support for physics applications. It will develop Grid interfaces for end-users including portal systems, carry out the necessary adaptation of physics applications to the Grid environment, and, finally, provide end-system software for physics data management.

The Grid technologies area will liaise with Grid technology development projects, in particular the European Union DataGrid project, as well as other high-energy physics projects. It will monitor these developments, select as appropriate the components produced as an outcome of these projects, and test their mutual compatibility in order to validate the concept of a coherent global computing environment.

The Grid deployment area will ensure the gradual deployment of the environment. It will plan the wide area networking and the integration of the Grid and physics environments, will coordinate the setting up of the regional centres, and will organize the overall Grid operation including the security and access mechanisms. The area will also plan the large-scale tests – called data challenges – conducted by the LHC experiments.

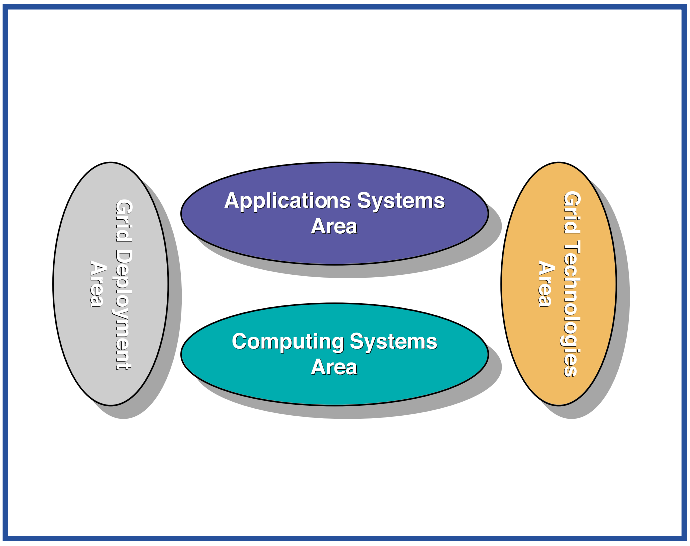

LCG project areas of work

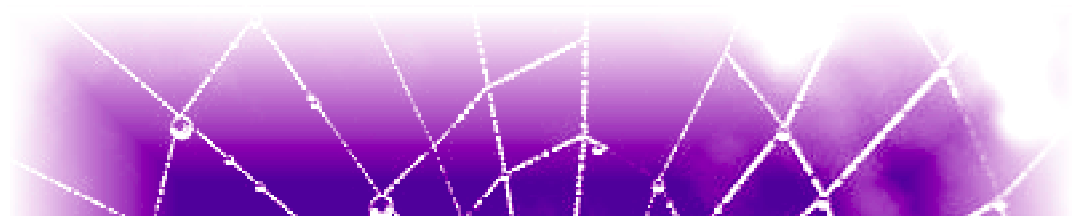

Positioning sponsor contributions in the LCG architecture

At the time of writing this report, three corporate sponsors have joined the CERN openlab for DataGrid Applications: Enterasys Networks, KPNQwest, and Intel Corporation,

These founding sponsors represent three areas of technology and services which will be key to the success of the LCG: very-high-speed networking equipment, wide area network services and processor technology, respectively. Specifically for 2002, they contribute as follows:

• Enterasys Networks provides a Ten-Gigabit Ethernet infrastructure to form the LAN intra-fabric network at CERN.

• Intel Corporation provides Itanium® 2 family processors for the servers of the LCG fabric.

• KPNQwest provides very-high-throughput long-haul network services for technology evaluation between CERN and one or several Tier1 regional centres.

The status of these contributions, together with the plans for the coming year, is detailed in the next chapter. The following figure provides a synthetic view of the positioning of the sponsor contributions on the overall LCG technology map.

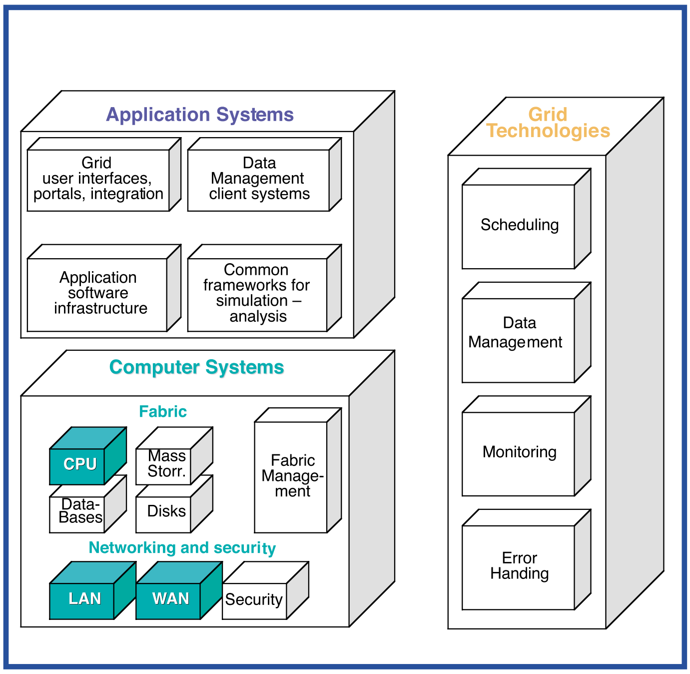

Positioning of sponsor contributions on the LCG technology map

Enterasys Networks

The needs for Ten-Gigabit networking: facts and figures

The first few months of the collaboration have been used to design a gigabit testing and deployment plan, and to start implementing it.

CERN pioneered the usage of Gigabit Ethernet technologies as early as 1999 with Enterasys technology, and has now set up a production gigabit-based network infrastructure, mainly deployed in the computer centre, for cluster interconnection as well as tape and disk server access.

However, Gigabit Ethernet will be unable to face the high demand of the LHC computer fabrics. Ten-Gigabit Ethernet technology is the natural – as well as mandatory – next step.

According to the LCG computing model, the CERN Tier0 and Tier1 will share a common physical, partitionable facility, which will eventually require of the order of 20,000 systems to beconnected to the intra-fabric LAN. Indeed, all processors of the CPU farm need to be connected individually - that is several thousands of processors, each requiring a dedicated one gigabit per second access to the LAN. In addition, hundreds of disk servers and mass storage devices will need to be connected at gigabit speed.

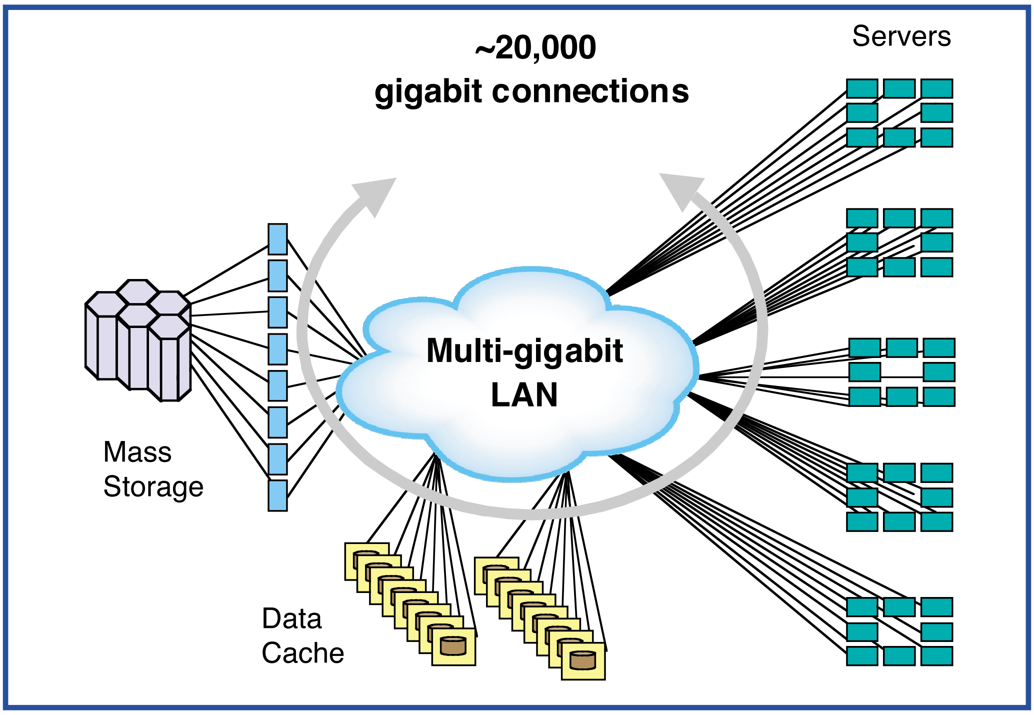

Eventual LAN connections for the CERN LCG fabric

The Enterasys Networks contribution and the network infrastructure model

By providing a Ten-Gigabit Ethernet infrastructure, Enterasys Networks will demonstrate its ability to meet the technical challenge of the CERN LCG fabric. The Enterasys products will be used to implement an end-to-end Enterasys solution providing Gigabit Ethernet connectivity to several thousands of computers and devices.

The infrastructure that will be demonstrated includes:
• an aggregation layer, composed of *Gigabit Ethernet switches* from Enterasys, aggregating the traffic from ten computers or devices onto a two gigabit per second trunk;
• a distribution layer, composed of *E1 Gigabit switches* with Ten-Gigabit uplinks, aggregating six trunks of two gigabit per second trunks into a Ten-Gigabit pipe;
• a routing/switching core composed of *Ten-Gigabit Ethernet routers.*

Achievements and work plan

As a first step, the E1 product family will be integrated in the current test-bed set-up, demonstrating a Ten-Gigabit Ethernet link between two E1 switches, interconnecting 120 computers.

This first prototype network is now in place and is ready for testing.

From May 2002, this set-up will be extended to a few more switches interconnected with Ten-Gigabit Ethernet links to a central Ten-Gigabit router, interconnecting a total of 240 systems. At the end of 2002, the set-up will be extended to at least two high-capacity central routers to accommodate the connectivity of more than 500 systems.

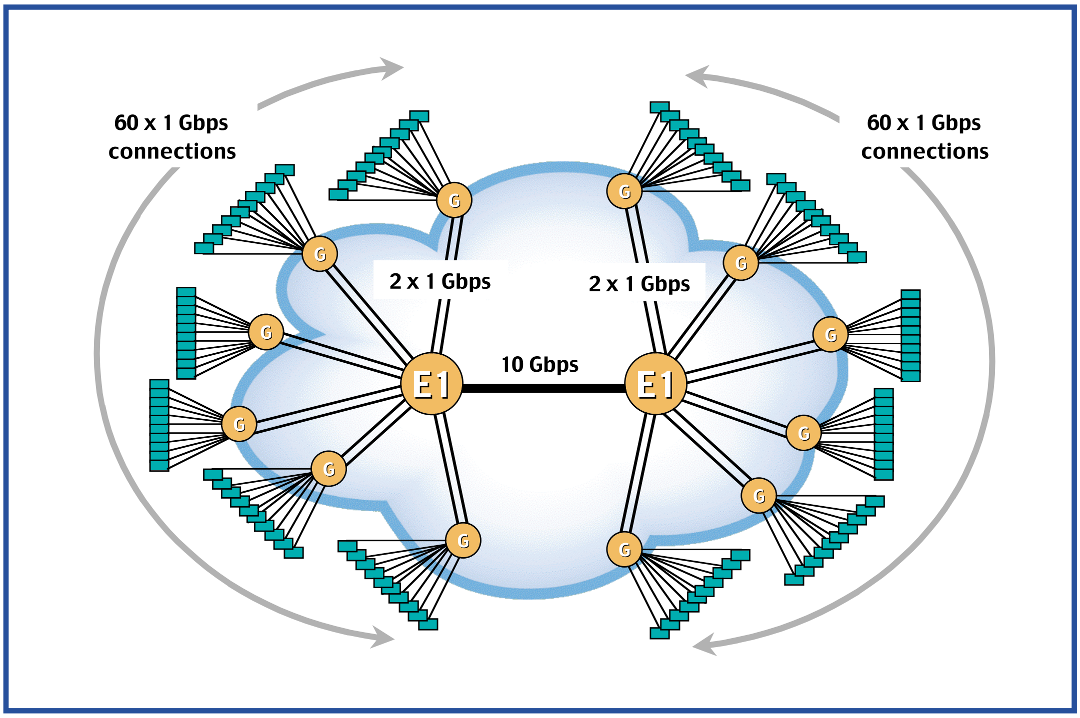

First prototype for the CERN LCG intra-fabric LAN

Intel Corporation

Intel Corporation

The LCG will bring exponential growth in processing power

Intel Corporation is providing Itanium®2 family processors for the servers of the LCG fabric.

To meet the demand of the LHC, the computing power available at CERN will have to be increased exponentially over the coming years. Between 2002 and 2006, the CPU capacity will have to be multiplied by a factor of 100. In 2002, the CERN computer centre cluster comprises 2200 processors. With processors of the same power, the number of computers required at CERN for the LHC would exceed 200,000.

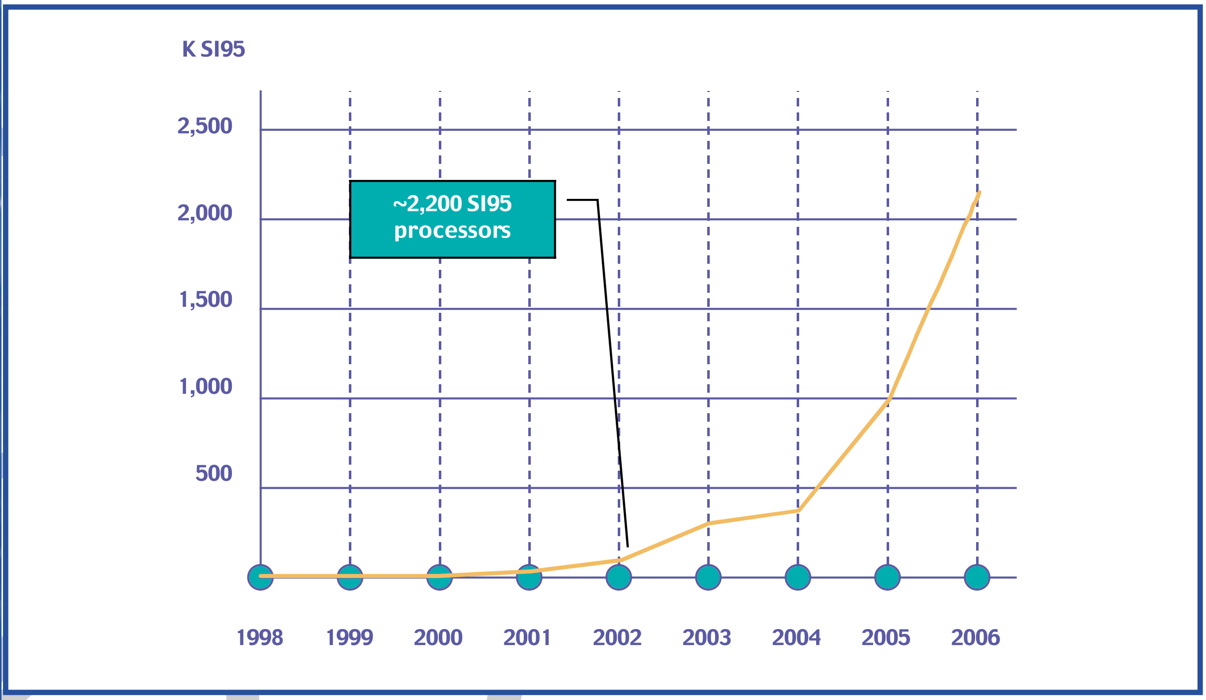

Estimated CPU capacity required at CERN

Gaining experience with 64-bit processor technologies

To cope with the explosive growth in processing power, it is essential for CERN to explore the use of the most advanced processors. In this context, CERN has been conducting since 1999 exploratory projects to assess the capabilities of 64-bit processors, and in particular the Itanium® family processors. These efforts included porting the Linux operating system to the Intel® Itanium® architecture and benchmarking farms of Itanium®-based processors from Fujitsu Siemens Computers. These tests demonstrated the high potential of the 64-bit technology as well as the crucial importance of the compiler technology to take full advantage of the parallel architecture of these processors.

The Itanium® test-bed cluster initially comprised 16 Itanium® processors at 800 MHz. The combination of Itanium® processors, Linux and LSF has proven to be highly effective for speeding up accelerator simulations, which are so challenging that the Sixtrack applications actually form part of the SPEC CPU2000 benchmark suite.

Plan for integration of Itanium® family processors

Intel is providing Itanium® 2 family processors in the context of the CERN openlab for DataGrid Applications.

To boost the computing power by a factor of 100, not only will very powerful processors be required, but also the level of replication of these processors must be dramatically increased.

It is therefore the product of the CPU power increase by the replication factor increase that may provide the necessary overall computing power increase. Assuming that the processing power increases by a factor of 10, the CERN CPU fabric will then have to host, by 2006, 10 times more processors than in 2002, that is of the order of 20,000 units.

CERN has already acquired solid experience in processor replication, in particular when evolving RISC-based processor farms (tens) to Pentium-based processor farms (thousands). Yet, building the 20,000-processor LCG fabric represents a formidable challenge. The CERN openlab for DataGrid Applications is in the process of identifying an industrial assembler to integrate Itanium® 2 family processors into LCG fabric servers.

As the initial tests have shown the importance of compiler technologies to take maximum advantage of the parallel architecture of the Itanium® 2 processors, Intel's compiler experts and CERN will work together to ensure that LCG applications become available on successive versions of the Itanium® family processors.

KPNQwest

KPNQwest is providing gigabit-speed long-haul network services to the LCG project. The first months of the collaboration were devoted to designing possible scenarios for setting up gigabit connections between CERN and other sites, considering different options for the technology to be used and the end-points of the connections.

The context of the LCG wide area networking

Several initiatives have already been launched by CERN and some of its partners to evaluate gigabit wide-area connections and progressively build the required LCG long-haul network.

CERN is connected to GEANT, the federative, European Union-sponsored network, which interconnects the European national research and academic networks.

In addition, CERN and three other partners – PPARC (United Kingdom), INFN (Italy) and University of Amsterdam (Netherlands) – are collaborating in the European Union project DataTAG. The objective of the DataTAG project is to allow a large-scale intercontinental Grid test-bed involving the European Commission DataGrid project, several national projects in Europe, and related Grid projects in the USA such as GriPhyN, PPDG and iVDGL. In this context, a 2.5 gigabit per second (Gbps) transatlantic connection will be made available in July 2002 between CERN and STARLIGHT in Chicago, the major North American interconnecting point. This 2.5 Gbps circuit will be used as a dedicated 'bridge' between EU and US Grid projects.
The partners will communicate with CERN and access the transatlantic bridge via the GEANT network.

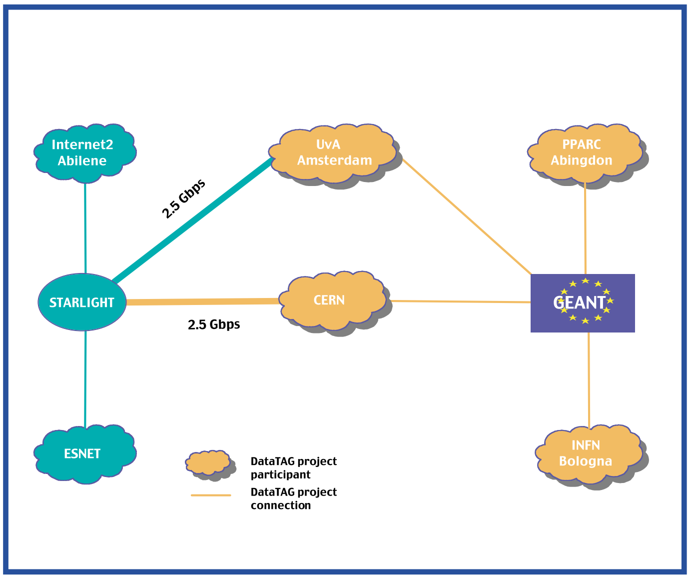

A simplified view of the DataTAG participants and their connections

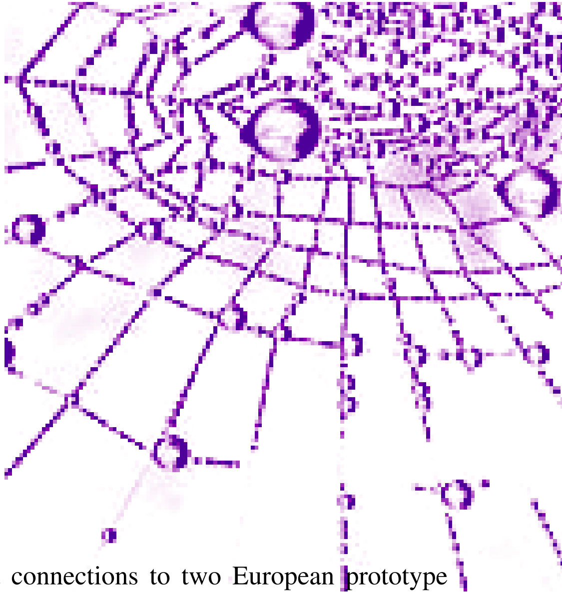

Scenarios for KPNQwest Gigabit connections

Several scenarios have been devised during the first phase of the project, with varying levels of technology innovation and differing end-points. These scenarios are not mutually exclusive, and indeed may serve as a basis for constructing a compound scenario.

Scenario 1: One or more 2.5 (or 10) Gbps links to Tier1 centres

In this scenario, KPNQwest would provide, for evaluation and testing, connections at 2.5 Gbps or, preferably, 10 Gbps to one or more of the following European prototype Tier1 regional centres:

- Forschungszentrum Karlsruhe, Karlsruhe, Germany.
- INFN Computing Centre (CNAF), Bologna, Italy .
- IN2P3 Computing Centre (CCPN), Lyon, France.
- Rutherford-Appleton Laboratory (RAL), Abingdon, United Kingdom.

The aim of this scenario would be to gain experience and evaluate the performance of applications operating at 2.5 Gbps+ speed.

Scenario 2: Lambda connections to two European prototype Tier1 centres

In this scenario, CERN connects first to the KPNQwest point of presence in Geneva with a lambda connection. From there, two lambda connections are set up to two European prototype Tier1 centres.

Signalling-enabled applications would be used to activate one of the two possible paths dynamically. This would allow testing and evaluation of advanced technologies such as Generalized MPLS (G-MPLS).

This scenario has a significant R&D component, in particular the exploration of the generalized MPLS technology.

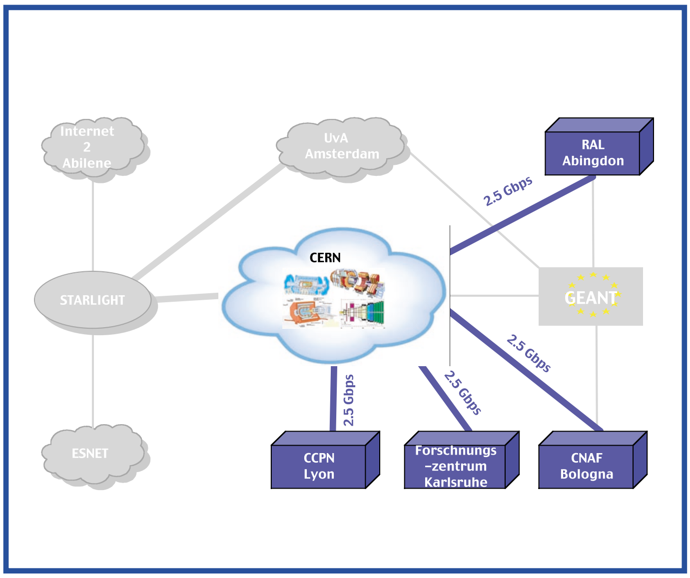

Scenario 1:Connections at 2.5 (or 10) Gbps to one or more prototype Tier1 centres

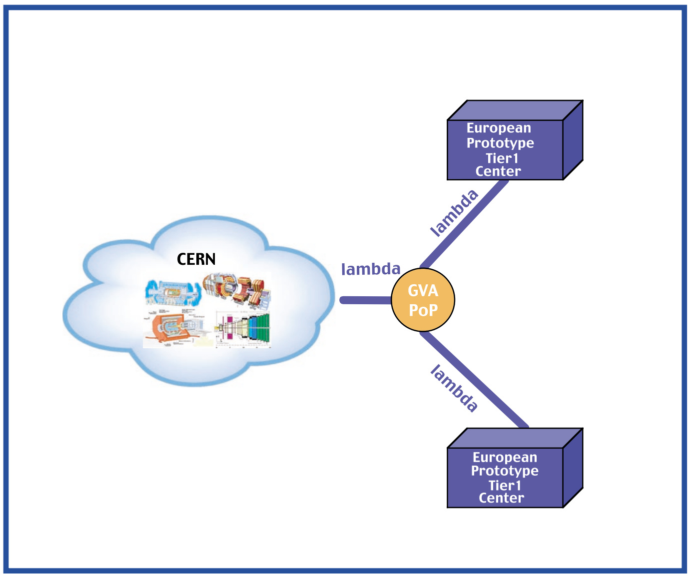

Scenario 2: Lambda connections to two European Prototype Tier1 centres

Scenario 3: Dark fibre between CERN and the closest prototype Tier1 centre

In this scenario, KPNQwest provides a dark fibre between CERN and the CNRS/IN2P3 Computing Centre in Lyon (CCPN), which is the closest prototype Tier1 centre to CERN. Indeed, dark fibres require signal regenerators at regular intervals, and for the project to keep to a reasonable size, the length of the connection should preferably not exceed about 150 kilometres.

This scenario has a strong R&D component, as the connection would be used to test advanced technology and evaluate new services. It would however require extra equipment and partners in particular for providing at least one regenerator system between Geneva and Lyon. The implementation of the local loop on the Lyon side would also need further study.

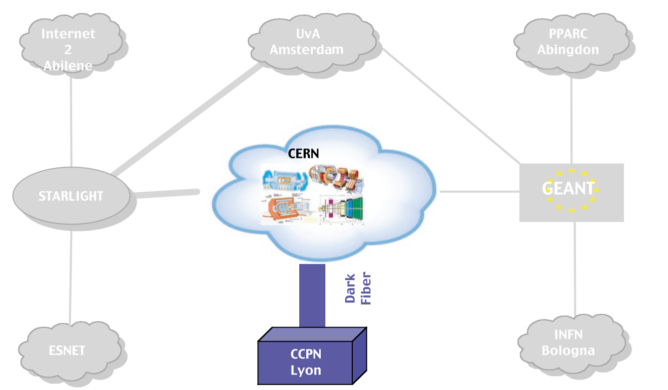

Scenario 3: Dark fibre between CERN and the closest prototype Tier1 center

Timelines

The three scenarios above have to be further examined and a feasibility study needs be conducted for each of them.

The work plan for the coming months is to carry out these studies and to evaluate the scenarios and their possible combination, in order to make a decision by the end of April 2002. The delivery of components of the circuits and equipment could then start in September 2002.

Looking ahead

The LCG project comprises two phases, the first covering development and prototyping in the years 2002-2005 and the second aimed at installation and operation of the full worldwide initial production Grid in 2006-2008. At the end of the first phase, a production prototype will have been deployed. It will represent a 50% model of the final Grid (that is, 50% of the complexity of one of the LHC experiments), including the availability of the committed Tier1 regional centres. To attain this objective, three annual prototypes of increasing complexity will be released during the years 2002-2004.

Plans for Prototype I

The objective is to build elements of a computer fabric at CERN (CPU clusters, disk storage, mass storage system, system installation, intra-fabric network, system monitoring) and to carry out performance and scalability tests using physics applications.

Prototype I will be developed during 2002. It will build on the existing CERN computing centre systems, with additional components from the CERN openlab for DataGrid Applications sponsors.

High-throughput tests will be possible using up to 500 nodes in test fabric at CERN. Part of this fabric is already up and running, including commissioning of the Ten-Gigabit Ethernet network.

Tests will also be carried out with regional centres using emerging Grid technologies.

Timelines for LCG and sponsor contributions in 2002-2003

As many as 400 additional computers will be added to the prototype LCG fabric in the course of the year. New disk units and mass storage systems will also be added and a Data Challenge will be run in collaboration with users from the LHC experiments.

In parallel, the contributions of the CERN openlab for DataGrid Applications sponsors will allow the following achievements:

Enterasys Networks: Extension of the network for interconnecting a total of 240 systems by May 2002, and further extension to accommodate more than 500 systems by the end of the year.
Intel Corporation: Identification by CERN of the industrial partner for assembling processing units based on Itanium®2 family processors around September 2002.
KPNQwest: Choice of the scenario (or combination of scenarios) for very-high-speed long-haul connections in April 2002, and first delivery of the connections by September 2002.

This is summarized in the following figure.

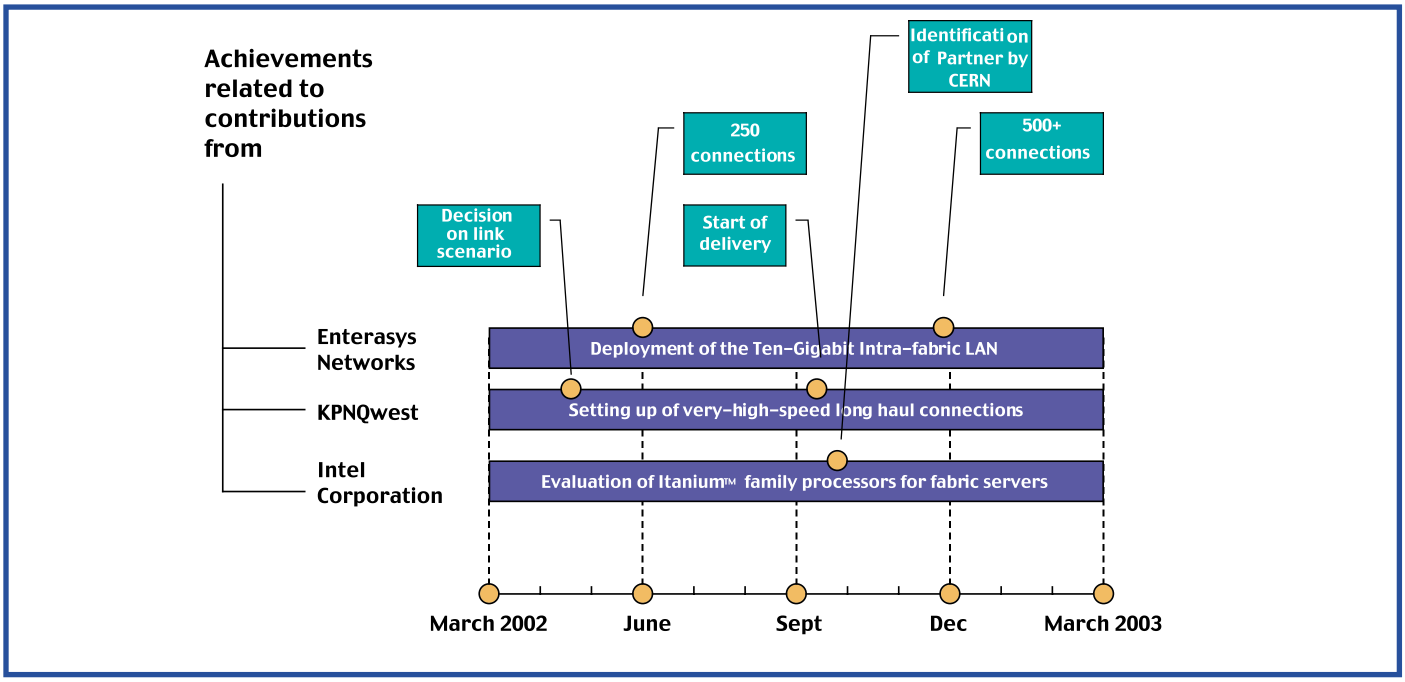

Timelines for 2002 sponsor contributions

Extending the openlab

The chapter of this report on the positioning of the sponsor contributions in the LCG architecture bears witness to the key role played by the technologies provided by the three founding sponsors.

It also demonstrates the scale and complexity of the undertaking, and shows that a number of technological areas present opportunities for corporate partners to join the CERN openlab for DataGrid Applications, participate in the precursor of the upcoming international Grid infrastructure, and benefit from taking part in the world's largest data-intensive challenge.

For more information, visit our website: http://www.cern.ch/openlab

Public relations:
Renilde Vanden Broeck,
CERN Press Office,
Renilde.Vanden.Broeck@cern.ch
Tel. +41 22 767 41 01

Executive Contact:
Manuel Delfino,
Head of the CERN openlab for DataGrid Applications, CERN Information Technology Division,
Manuel.Delfino@cern.ch
Tel. +41 22 767 45 85

Technical Contact:
Les Robertson,
Project Leader of the LHC Computing Grid, CERN Information Technology Division,
Les.Robertson@cern.ch
Tel. +41 22 767 49 16

ENTERASYS NETWORKS

intel

kpnQwest

CERN

openlab for DataGrid applications

Developing Solutions for the Data-Intensive Science of the Large Hadron Collider