CERN accelerates the future

Intel's 10 Gigabit Ethernet tackles the challenge of data-intensive particle physics research

Case Highlights

Profiled Organisation	CERN: The European Organisation for Nuclear Research.
Challenge	To construct an environment to support the transfer and analysis of Petabytes of data collected from the next-generation Large Hadron Collider (LHC).
Milestone	To achieve aggregate speeds of 10 Gbps across the LAN, supported by Intel® 10 GB Ethernet technology.
Result	A pilot infrastructure has delivered high performance and high availability; it's a promising result for the future development of the global LHC Computing Grid.

Summary

Founded in 1954, CERN is the European Organisation for Nuclear Research, the world's largest particle physics centre. Here, approximately 6500 visiting scientists, half of the world's particle physicists, come to explore what matter is made of and what forces hold it together. CERN exists primarily to provide these physicists with the necessary tools: accelerators, which accelerate particles to almost the speed of light, and detectors, which make the particles visible. CERN employs nearly 3000 people.

Construction is already underway for what will be the world's largest particle collider, the Large Hadron Collider (LHC), which is anticipated to start operation in 2007. CERN's Information Technology (IT) Department is anticipating unprecedented amounts of data from the four giant detectors on this ring, which measure the results of collisions between protons travelling in opposite directions around the ring. Buried in this data physicists expect to find elusive traces of exotic fundamental particles, such as the Higgs boson, which have been predicted by theories of the origins of our Universe, but not yet detected experimentally.

Every year, up to 15 Petabytes of data from the LHC will be distributed on the LHC Computing Grid (LCG), a worldwide grid linking hundreds of computer centres in universities and major research labs. The data must be collected and transported across the LAN/WAN infrastructure at CERN. In order to plan for the future, CERN's IT Department set up the CERN openlab in 2003, in close collaboration with leading IT companies. The first major project of CERN openlab was to develop a pilot environment called the opencluster, based on 10 Gbps networking technology, in order to verify that the infrastructure would deliver the required performance.

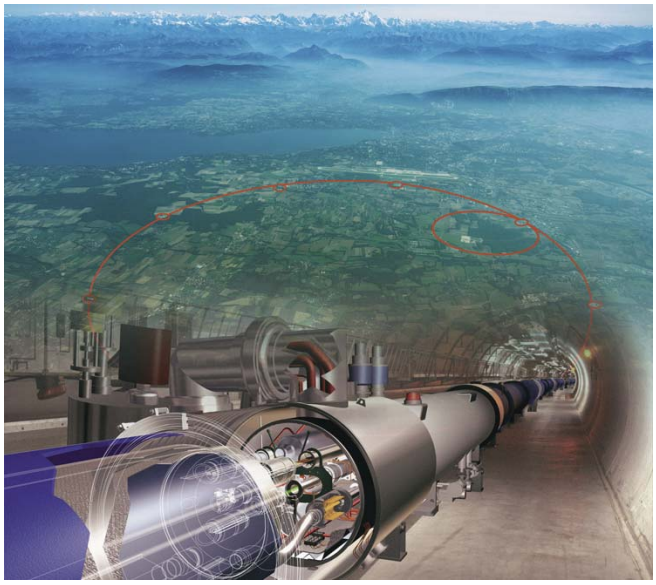

Figure 1: Artist's impression of the Large Hadron Collider

Challenge: Achieving leading-edge performance and availability

The CTO of CERN openlab, Sverre Jarp, heads a dedicated team of engineers, responsible for researching innovative technology that can be incorporated into CERN's overarching IT architecture. This R&D team works as an integral part of the CERN IT Department.

Having joined CERN in 1974, Jarp has witnessed the continual evolution of technology, and his team is recognised as an 'early adopter'. The opencluster project aims in particular to test cutting edge technology in the uniquely demanding environment of the LHC Computing Grid (LCG) a grid of computing centres involving CERN and some 80 partners worldwide. As Jarp comments: "This project has plenty of challenges. But for us this is business as usual." High energy physics has traditionally been an innovator in networking technology, and it is no coincidence that the World Wide Web was born at CERN.

"We work with industrial partners, like Intel, to verify cutting-edge technology," Jarp continues. "For our current project, we have setup a pilot environment to simulate the unprecedented amounts of data that we are expecting from the LHC experiments, and to ensure that the infrastructure can handle the challenge, from data collection through to the huge storage silos."

Peak performance, of course, is not the only measure of success. The experiment requires fast, granular and continuous collection and transfer of information, meaning that the IT infrastructure must be highly available to ensure that nothing is missed.

Process: Opting for a logical progression

The process of choosing to work with Intel as one of the collaborating partners was a logical one for Jarp. "We have a long tradition of working with Intel," he says, "and we are very satisfied with the Intel products we are using today."

The current IT infrastructure at CERN is supported by an environment that is based on Intel® Xeon™ dual processors and Intel's 1 Gbps Ethernet solutions. Jarp explains: "Despite the fact that this sort of environment is only now becoming commonplace in commercial environments, we've been running ours successfully for several years and are familiar with high-performance Ethernet technology."

"Of course, we are always evaluating alternatives and new solutions—we need an environment that will move us forward by at least one order of magnitude, after all—but one of the great things about 10 Gbps Ethernet is that it is a logical progression from the environment we have today. This represents a strategic investment, both for us and our scientific partner institutions worldwide."

CERN has a huge Ethernet base comprising 25,000 outlets and 10,000 devices within its own organisation. Furthermore, collaborating partner institutions will have similarly significant installations.

The LHC is anticipated to run for well over a decade, and there will be upgrades that will result in even higher data production rates. Given these long-term demands, established industrial partners with defined roadmaps are valued. "We know we are at the limits of what the technology can handle today. Because we are an early adopter of technology, Intel's intention to push Ethernet to greater speeds is of real interest to us. Additionally, we value their flexibility, such as providing us with Open-Source versions of the I/O drivers, which makes it easier for us to control the pilot environment."

Solution: Driving 10 Gbps performance

The pilot environment is about achieving 10 Gbps performance, end-to-end. “The most critical aspect for us is aggregated throughput,” highlights Jarp. “It’s not just the interface cards we are testing; we need to know that the server and switching equipment can handle the 10-gig rates.”

In the spring of 2003, CERN deployed equipment from several vendors to achieve the ground-breaking performance they need. The Openlab environment consists of 100 Dual RX2600 servers with Itanium® -2 processors, most of which are running at 1.5GHz. 16 of the servers have an Intel PRO/10 GBE SR Ethernet PCI-X Server Adapter.

In the production environment that will be deployed by 2007, data collected from the LHC will be passed, across the 10 Gbps LAN backbone, to the processing servers. These engines reconstruct the raw data, adding context, before sharing it with the wider community of so-called Tier-1 centres, which store complete data sets from the experiments. There will be several of these centres at high-energy physics institutes in Europe, including Rutherford Laboratories in the UK and IN2P3 in France. Other organisations are based around the world and include Fermi Laboratories in Chicago.

“So far, we are very pleased with the performance of the pilot environment. Even with a single stream of data coming form one application, we’ve been able to use more than 50% of the cards’ available bandwidth—that’s about 5Gpbs out onto the network,” smiles Jarp. “We found it quick and easy to get the Intel PRO/10 GBE Server Adapter cards working and they’ve delivered solid and consistent performance ever since, allowing us to meet our aggregation targets.”

Figure 2: An illustration of the LHC Computing Grid (linking over 80 computer centres worldwide for the processing of LHC experimental data)

Future: The need for speed

The LHC will be the most powerful instrument ever built to investigate the fundamental properties of matter. However, CERN scientists already have their eyes on tomorrow and even higher energy experiments. The need for speed will continue for the foreseeable future.

“40 Gbps is on the horizon,” says Jarp, “and that means we are investigating, processors, memory, LAN and WAN technology that can deliver this performance. Through the CERN openlab framework, Intel is a key partner in this process, helping us to continuously push the technological envelope, not just for the network interface cards, but also from processors through to entire systems.”

Find out more about a business solution that is right for your company by visiting the Intel website at:
www.intel.co.uk/network

Solution provided by: