

BY DAVID BAUM



Wolfgang von Rüden, Head of openlab at the European Organization for Nuclear Research (CERN)

When Worlds Collide

THE STRONG BOND BETWEEN CERN AND ORACLE PUTS A POWERFUL SPIN ON PHYSICS.

Ever since Albert Einstein unveiled his general theory of relativity in 1915, physicists have attempted to devise a universal theory that can explain how all the particles and forces in the universe interact. The so-called Standard Model reveals the relationship between three of the four fundamental forces—electromagnetism, the strong nuclear force, and the weak

nuclear force. But it doesn't explain why things have mass, nor does it fully incorporate the physics of general relativity, such as gravitation and dark energy.

To plug the holes in the Standard Model, scientists have proposed the existence of the Higgs boson, a hypothetical particle that could reveal how other subatomic particles are thought to

acquire mass. Until now this elusive particle has been demonstrated mathematically but has escaped the observation of scientific instruments. But thanks to the work of the Large Hadron Collider (LHC) at the European Organization for Nuclear Research (CERN), where many Oracle databases are deployed in support of a massive data grid to store scientific and operational data, researchers are getting closer to confirming the existence of the Higgs boson.

Situated more than 100 meters below the gently rolling countryside around Geneva, Switzerland, CERN built the LHC, the world's largest scientific instrument, to study the properties of subatomic particles. Within the LHC's subterranean tunnels, powerful magnets are chilled to 1.9 degrees above absolute zero to guide beams of superenergetic protons or lead nuclei around a 27-kilometer ring. The particles are accelerated to speeds approaching the speed of light and collide head-on, re-creating similar conditions at laboratory scale as during the massive bursts of energy that occurred a fraction of a second after the Big Bang.

Scientists hope that the particle activity observed at the LHC will shed some light on many fundamental questions in physics. CERN relies on Oracle technology and applications to keep the LHC operational—and to accurately share its findings with researchers all over the world.

A MASSIVE DATA ANALYSIS CHALLENGE

Hadrons are a class of subatomic particles that includes protons and neutrons. When they collide, their energy is transformed into dozens of other infinitesimal particles, such as the theoretical Higgs boson. Being able to create mathematical models that describe these particles is one thing—observable proof is quite another. To verify the existence of the Higgs boson, LHC researchers have to process a staggering amount of information.

Four very large detectors, comparable to 100-megapixel 3-D cameras, collect “images” up to 40 million times per second, yielding raw event data on the order of 1 million GB/sec. None of today's computing systems are capable of

CERN openlab: Forging Connections Between Science and IT

The Large Hadron Collider (LHC) at the European Organization for Nuclear Research (CERN) is a multi-disciplinary project that has involved thousands of scientists, engineers, operations personnel, and administrators. To push the use of new technologies in support of the state-of-the-art IT infrastructure and information systems for the global scientific community working with the LHC, CERN devised openlab. Through close collaboration with Oracle and other leading IT vendors, CERN acquires early access to technology that is still years from reaching the general computing market. In return, CERN offers expertise and a highly demanding computing environment for pushing new technologies to their limits and provides a neutral ground for carrying out advanced R&D with various partners.

“Our slogan is, ‘You make it, we break it,’” quips Wolfgang von Rüden, head of CERN openlab. “The participants obtain hardware, software, and networking products in the very early phase—often beta releases or pre-beta versions. We install them in our production environment at CERN and put them through rigorous testing. For the last seven years, Oracle has

been helping us push database technology and related solutions forward, and in turn we are helping Oracle enhance its products.”

For example, CERN has been an active contributor to Oracle Database 11g Release 2. During the beta testing period and in the months following, the CERN team performed extensive work using Oracle Streams, Oracle Automatic Storage Management, Oracle Active Data Guard, Oracle Real Application Clusters, and Oracle Advanced Compression. This collaborative effort yielded new techniques for replicating large databases, scaling compute capacity, and managing enterprise software environments. The openlab team also pioneered new virtualization techniques through close collaboration with the Oracle JRockit Virtual Edition team.

According to Juan F. Rada, a senior vice president at Oracle with worldwide responsibility for public sector, health-care, and education industry solutions, openlab is an ideal test bed for evaluating and integrating new technologies and services. He says the openlab collaboration has had a decisive impact not only on the Worldwide LHC Computing Grid (WLCG), distributing

information and results to a worldwide community of researchers, but also on the deployment of cutting-edge large-scale solutions that will become mainstream in the market.

For example, CERN uses Oracle Active Data Guard standby databases to replace existing servers and storage with new, more powerful models, enabling complete technology refresh with minimal downtime. CERN simply creates a physical standby database on the new technology platform and then executes an Oracle Data Guard switchover to move production to the new platform. Oracle customers can apply these principles to minimize downtime during OS upgrades, cluster resizing, migration from nonclustered to clustered systems, and other data-center moves.

“What makes CERN really useful to Oracle is the fact that we are able to test configurations and create proofs of concept in a very challenging production environment,” Rada concludes. “Oracle has been able to develop better products as a result of the openlab collaboration. Our goal is to support the needs of the CERN scientific community while improving the technology for all Oracle customers.”

recording such data rates. Fortunately, the vast majority of these events contain known physics, and only in very rare cases does an event hide new physics. Nevertheless, every single event has to be analyzed in real time to avoid missing a breakthrough. A multilevel pipelined trigger system, based on dedicated electronics and a large PC farm, applies sophisticated criteria to keep the significant data of each experiment below 1 GB/sec. All told, CERN's four big detectors—ALICE, ATLAS, CMS, and LHCb—produce more than 15 million GB of data per year.

In addition to the physics data, each experiment yields information about the positions of all of the detectors, the gas mixtures, the temperatures, and so forth—the conditions about each experiment—along with summary data that identifies the types of particles and their associated energy levels. The scientists rely on these measurements and instrumentation records in order to calibrate the “camera” before proceeding to the physics analysis.

“In addition to collecting the physics event data, the experiments measure and collect conditions data from about 1 million channels in each detector,” explains Wolfgang von Rüden, head of CERN openlab, a collaboration between CERN, Oracle, and other technology vendors. “The results are stored in an Oracle database as environmental parameters—temperatures, pressures, voltages, currents—that reveal the state of the machinery.” CERN needs precision on the order of a few microns to follow the tracks of the particles. The data is extremely important for making precision alignments.

In a similar way, the LHC accelerator physicists manage many hundreds of thousands of machine settings and record terabytes (TB) of operational data in databases as well, to guarantee that the conditions of a particular accelerator setting can be reproduced with high accuracy for optimal running conditions. Conversely, when things go wrong, scientists can determine what changes they need to make based on deviations from the standard setup.

Scientists depend on this data to verify exactly what an operator has done, when, and what the effects were, so if they restart the accelerator with similar conditions they are able to repeat the experiment. “It’s complicated machinery that you have to watch all the time,” adds von Rüden. “We even detect earth tides created by the gravitational attraction of the moon and the sun on the rotating earth, and we have to compensate for these variations.”

This operational data is also closely

“The experiments measure and collect conditions data from about 1 million channels in each detector. . . . It’s complicated machinery.”

—Wolfgang von Rüden,
Head of openlab at CERN

the LHC still gaining speed; in the summer of 2010, the collider had only attained half its ultimate energy. The critical parameters are the very high intensity and density of the beams. Given the required precision of the operations and the high power stored in the beam and in the accelerator equipment, keeping track of system anomalies is a critical undertaking.

“Oracle databases are used to manage the settings and controls configuration necessary to drive all accelerator installations, and the real-time interaction with the database is mandatory for running the accelerator. Short-term measurements and long-term logging are also stored in databases,” explains Eric Grancher, section leader of database services in CERN’s IT department. “The logging information amounts to more than 3.5 terabytes per month and has to be kept for the next 20 years in order to understand, tune, and boost the performance of the LHC accelerator. That’s why these data are so important.”


FUELING THE GRID

Once valuable collisions are observed and recorded, the data from each experiment is shared with physicists and researchers around the world to analyze and to look for the Higgs boson and other discoveries. To meet that goal, scientists have connected tens of thousands of computers into a distributed computing network called the Worldwide LHC Computing Grid (WLCG). Supported by replicated Oracle databases, the WLCG permits scientists everywhere to actively participate in analyzing the LHC data.

The WLCG supports the offline computing needs of the

LHC experiments via 130-plus computer centers in 34 countries. When experiments are conducted, the project generates about 27 TB of raw data per day, plus 10 TB of event summary data. The data produced by the LHC on all of its distributed computing grids is expected to add up to 10 to 15 petabytes (PB) per year. CERN is the central hub of this massive computing grid, the Tier-0 center. An initial copy of all data from the LHC experiments is held here. The central hub is connected to 11 Tier-1 computing centers using dedicated 10 Gb/sec optical networks.

The Tier-1 sites make data available to Tier-2 centers, each consisting of one or several collaborating computing



This computer model of activity in the ALICE detector shows particle tracks spinning off after high-speed collisions in the LHC.

facilities, which can store sufficient data and provide adequate computing power for specific analysis tasks. These Tier-2 sites, grouped in federations, cover most of the globe. Individual scientists access these facilities through local computing resources. Previously scientists had to wait weeks to get access to the processed data sets of their experiments—now they get these data sets the same day.

“We’re permitting a global research community to work on the data as it’s produced. The Oracle database is a major enabler for this effort,” von Ruden says, adding that Oracle Streams, a feature of Oracle Database 11g Release 2, provides a unified solution for information sharing. “We tuned Oracle Streams to reach the performance we needed to make sure those databases can stay in sync. It’s much more efficient than sending tapes around.”

A COMPLETE STACK OF TECHNOLOGY

CERN uses Oracle databases not only to control the accelerators and to store data about the measurements taken by the experiments, but also for many business functions.

“Oracle databases underpin not just physics results but also accelerator operations and certain administrative activities,” says Tony Cass, who heads up the database services group at CERN. “Therefore, as the capabilities of Oracle Database improve, so does the level of monitoring and control that we have.”

Cass manages a team charged with making sure that the database management systems are effectively stable and available 24/7. He says universities worldwide collaborate with the institutions working on the accelerator, and they need to be able to order equipment and supplies for the labs, check on budgets, and perform other administrative functions. “The Oracle databases and tools support those administrative efforts, just as they support technical design studies and engineering change requests in the engineering database management system,” he adds.

In addition to Oracle Database 10g and Oracle Database 11g, CERN uses Oracle Enterprise Manager for monitoring, Oracle VM, Oracle Application Server, Oracle WebLogic Server, Oracle Application Express, and Oracle Data Guard to create synchronized replicas of the databases to ensure business continuity. Approximately 40 PB of data is stored on four large StorageTek tape libraries from Oracle.

CERN also uses Oracle E-Business Suite, commercially and home-developed applications running on Oracle Application Server, and Oracle Database to manage procurement, human resources, and other business functions. It’s not just scientists running experiments who need to procure equipment and

supplies, but operations personnel throughout the organization, who need everything from PCs to paper clips to do their jobs. As Cass says, being able to handle small details quickly is important. “We’ve moved to interfaces where we have e-business contacts with other suppliers, so we don’t maintain stock onsite, but we do maintain the ability to have essential items delivered rapidly.”

CERN developers adopted Java as the strategic platform for developing their in-house software applications. CERN received the Duke’s Choice Award at the 2008 JavaOne Conference for the entire collection of Java applications that CERN has developed for the installation and operation of the LHC. These custom interfaces simplify access to administrative information services encompassing all of the procurements, workflows, staff management, budget management, project management, and some executive dashboards.

“The use of Java means we have a common Web-accessible interface that anyone can use through whatever computing environment they like,” adds Cass. “We don’t have to tell them which operating system or browser to use.”

>> SNAPSHOT

European Organization for Nuclear Research (CERN)
cern.ch
Founded: 1954
Industry: Scientific research
Employees: 2,500, plus 10,000 visiting scientists
Oracle products: Oracle Database 11g, Oracle E-Business Suite, Oracle WebLogic Server, Oracle JRockit Virtual Edition, Oracle Streams, Oracle Enterprise Manager, Oracle Automatic Storage Management, Oracle Active Data Guard, Oracle Real Application Clusters, Oracle Advanced Compression, StorageTek tape libraries, Oracle Fusion Middleware, Oracle Application Express, Oracle VM

SOLVING QUANTUM RIDDLES

As the activities in CERN pick up speed, scientists the world over hope to dramatically improve our understanding of quantum physics. The experiments taking place at the LHC and the huge amount of data they collect helps scientists understand important details about the birth of the universe, dark matter, and hopefully, the elusive Higgs boson, which will round out the Standard Model by lending mass to matter.

“As an organization that supports a worldwide, distributed community, CERN must be able to rely on the databases that underpin the operation,” Cass sums up. “Thanks to the data grid, physicists everywhere can be involved in the experiments. This worldwide level of participation is bound to have a major impact on the field.” <>

DAVID BAUM is a frequent contributor to *Proff*.

>> FOR MORE INFORMATION

Oracle Database 11g Release 2
oracle.com/us/products/database

Oracle E-Business Suite
oracle.com/us/products/applications/ebusiness

Oracle WebLogic Suite 11g
oracle.com/us/products/middleware/application-server/weblogic-suite