National Ignition Facility & Photon Science Drives Research with Robust Data Infrastructure

National Ignition Facility & Photon Science – Lawrence Livermore National Laboratory
Livermore, CA
https://lasers.llnl.gov

Industry:
High Technology

Employees:
800

Oracle Products & Services:

Oracle Database
Oracle Real Application Clusters
Oracle Partitioning
Oracle Active Data Guard
Oracle Total Recall
Oracle BPEL Process Manager

Key Benefits:

- Gathered and stored large volumes of experimental data in less than 30 minutes
- Gained ability to support a database that is expected to grow by approximately a petabyte each year
- Provided a flexible and scalable solution
- Ensured security of sensitive data
- Enabled accurate budgeting and planning

"Oracle infrastructure software is integral to our mission to study fusion energy as a source for clean power." – Tim Frazier, Senior Software Architect, National Ignition Facility & Photon Science - Lawrence Livermore National Laboratory

As the movement for clean energy technologies continues to grow, the National Ignition Facility (NIF) & Photon Science Directorate at the Lawrence Livermore National Laboratory is conducting ground-breaking research that could help dramatically reduce the carbon footprint by generating power from nuclear fusion.

"Fusion energy is the holy grail of clean energy and the future of meeting the energy needs of our civilization. People have been working on fusion energy for 50 years, and, if successful, the National Ignition Facility is going to be the crowning achievement in that search," said Ed Moses, principal associate director for the National Ignition Facility & Photon Science Directorate.

NIF's scientists and engineers use highly powerful lasers to produce controlled thermonuclear fusion through a process called inertial confinement fusion. Housed in a 10-story building the size of three football fields, the facility's laser system shoots 192 beams that focus intense energy onto a small target in the center of a target chamber, creating conditions similar to those that exist only in the cores of stars and planets or inside a nuclear weapon. The resulting fusion reaction is expected to release many times more energy than the laser energy required to initiate the reaction. This energy gain will be a key milestone in the pursuit of fusion energy as a source of electricity.

The laser system uses as many as 90 instruments—from cameras to mirrors, to oscilloscopes—to set up and measure the results of each experiment. The scientists and engineers need to precisely manage the placement of the equipment involved in the experiments and then gather and quickly analyze all resulting data.

Published September 2009

"Oracle Real Application Clusters gives us the flexibility to add nodes if we find we're under capacity. Since implementing Oracle Database 11*g* and Oracle Real Application Clusters, we have seen dramatic improvements in performance."

Tim Frazier
Senior Software Architect.
National Ignition Facility & Photon Science

To manage all of this data, as well as the business of constructing and budgeting for the facility, NIF selected Oracle Database 11*g* with Real Application Clusters.

"Oracle is integral to our research in many ways. We plan the construction of the facility and manage the budget with Oracle. We also use Oracle Secure Files for two databases, one to control the laser itself and one to collect and analyze all of our scientific data," said Tim Frazier, senior software architect at NIF.

Achieving Extremely High Performance

NIF's instruments have a very small time frame, approximately 30 minutes, to gather results in the database.

"If we don't have a high-performance method for saving images into the database, the instruments may time out, and we'll lose data," Frazier said.

The data produced by each experiment is generated in 10 megabyte to 100 megabyte hierarchical data format (HDF) files that may contain anywhere from 1 to 50 images. These files are written into the database and then the data is analyzed and inspected, which, in turn, increases data volumes by roughly 300%. At this point data volumes are approximately 350 gigabytes, which is ingested and then made available to scientists in less than 30 minutes. NIF will store its data in a database that makes use of Oracle Automatic Storage Management.

When operating at full capacity, the facility is expected to generate approximately a petabyte of data per year. This made the compression feature of Oracle secure files critical.

"We need to economically manage data storage, while ensuring the data is accessible to scientists," Frazier said.

Further, the data needs to be accessible for the lifetime of the facility, which is targeted for 30 years. With the potential to reach three to four petabytes in the next three to four years, this is a major undertaking.

Published September 2009

Ensuring Flexibility and Scalability

Prior to implementing Oracle Database 11*g* with Real Application Clusters, NIF was evaluating file system architectures to handle the amount of data that would be generated by each experiment. By deploying the database that gathers scientific data on Oracle Real Application Clusters, the organization can quickly add processors if it is falling behind in data collection.

“Oracle Real Application Clusters gives us the flexibility to add nodes if we find we’re under capacity. Since implementing Oracle Database 11*g* with Real Application Clusters, we have seen dramatic improvements in performance” Frazier said.

Oracle Real Application Clusters also enables NIF to provide different performance levels to different users based on need. For example, gathering the data from experiments requires the highest level of performance. However, scientists doing ad hoc queries likely do not require as much. Oracle Real Application Clusters gives NIF the flexibility to balance these loads.

NIF is also exploring Oracle's dynamic grid capabilities, built on Oracle Database 11g with Real Application Clusters, to deliver even more flexibility. This enables the dynamic sharing of resources within a cluster across widely different applications using Oracle Real Application Cluster's services and provisioning features.

“We are very enthusiastic about the concept of dynamic grid and look forward to working with Oracle as it further develops this strategy,” Frazier said.

Enabling Streamlined Analysis

Once the NIF laser system generates data and it is stored in an Oracle Database, it must be analyzed. NIF uses the Oracle BPEL Process Manager workflows to trigger analysis and ensure that appropriate analyses are performed and in the proper order.

“Our use of Oracle BPEL workflows has proven to be a good decision because now we can write analysis routines and know that they will be activated in response to the right data being generated without writing code for each one,” Frazier said.

Published September 2009

Enabling Faster Querying

NIF is also leveraging Oracle Partitioning to enable faster querying. For example, each day, the facility monitors environmental factors—temperature, humidity, pressure, and more—generating nearly 30 million rows of data per day. This large volume of data needs to be available for querying to track trends on a daily, weekly, monthly, and yearly basis.

Oracle Partitioning also helps NIF manage information lifecycles to ensure optimal performance and efficient storage utilization. NIF uses a tiered storage environment that incorporates high performance disk, serial advanced technology attachment (SATA), and tape.

Ensuring Security and Backup

The scientific data NIF manages must be protected without fail. NIF uses Oracle Data Guard to segment the database, allowing users to view only the data they are authorized to see. It also uses this feature to create a real-time backup of the database, protecting the facility against data loss due to a major disaster.

NIF is also beginning to use Oracle Total Recall to enable scientists to re-analyze old data if necessary due to instrument calibration updates.

Improving Budgeting and Planning

In addition to the scientific aspect of NIF's Oracle implementation, the organization also leverages Oracle Database 11*g* for budgeting and planning processes, such as workforce planning. Oracle helps NIF track workforce factors such as completion of its very rigorous safety program, to ensure it is able to meet workforce needs now and in the future.

Why Oracle?

An Oracle Database 10*g* customer, NIF stayed with Oracle because it found that Oracle Database 11*g*, with its unique features, was capable of scaling to meet the organization's needs for high volume and high performance.

Published September 2009

"Being true to the form of a research lab we don't leave any stones unturned. We created a series of blind tests that would help us convince ourselves that a migration to Oracle Database 11*g* would help us attain the high performance metrics we needed to move all the data," Frazier said. "We were able to stand up two databases and run our ingestion routines through both. We found that Oracle Database 11*g* with Oracle secure files was an order of magnitude faster than Oracle Database 10*g*."

Implementation Process

NIF upgraded to Oracle Database 11*g* over a four-month period. It also deployed Oracle Real Application Clusters and built its system using a three-node cluster to meet the needs of its various databases and optimize hardware usage.

Throughout the construction and testing process, the organization has been performing experiments. In spring 2009, the laser system is slated to be fully operational. Once it is, NIF plans to explore additional areas, including data mining.

"The work we've done so far makes it possible to manage data and to complete the construction of the facility. Once we have experimental data coming, in we'll look for more advanced ways to understand how they relate to each other. We feel this will be a key area for collaboration with Oracle," Frazier said.

Advice from NIF

- Leverage the expertise of your partners whenever possible. This is a core value at NIF. We are not trying to build a better database. Oracle has already done that; we leverage technologies from companies like Oracle whenever possible.

The National Ignition Facility & Photon Science (NIF) is the largest single research project in the United States, with the largest laser in the world. NIF's system of 192 giant lasers, housed in a 10-story building the size of three football fields, is expected to make significant contributions to national and global security and could lead to practical fusion energy. The project is a national collaboration among government, industry, and academia, as well as many industrial partners throughout the world.

Published September 2009