

BP looks to ORNL, ADIOS to help rein in data

January 3 2020



The Adaptable IO System (ADIOS) provides a simple, flexible way for scientists to describe the data in their code that may need to be written, read, or processed outside of the running simulation. Credit: Oak Ridge National Laboratory

Researchers across the scientific spectrum crave data, as it is essential to understanding the natural world and, by extension, accelerating scientific progress. Lately, however, the tools of scientific endeavor have become so powerful that the amount of data obtained from experiments and observations is often unwieldy.

In other words, it is possible to have too much of a good thing.

Making sense of today's ballooning datasets has become a major scientific challenge in its own right, forcing researchers to not only tackle their domain science problems but also the problem of managing and processing their ever-growing datasets. Just ask researchers at BP, who are tasked with finding natural gas and oil in the ground and figuring out how best to extract it.

"New technologies in the field allow us to collect more data than we ever dreamed of," said BP HPC Computational Scientist Vladimir Bashkardin, referencing the properties of subsurface fluid and rocks obtained via energy responses to the company's probing. "We need to scale our ability to access large seismic datasets, which can measure half a petabyte at times."

To assist them in this monumental effort Bashkardin and his colleagues turned to the Department of Energy's Oak Ridge National Laboratory, home to Summit, the world's most powerful and "smartest" computer, and a wealth of expertise on how to manage and process today's large and complex scientific datasets.

Summit's debut marked the third time the laboratory has stood up the world's fastest supercomputer. These systems have been used to tackle some of the most pressing scientific challenges of our time including fusion energy, drug delivery, and the design of novel materials, efforts that have also made ORNL a world leader in the increasingly important arena of big data.

BP researchers turned to ORNL Scientific Data Group Leader Scott Klasky and ORNL Scientific Data Management Team Lead Norbert Podhorszki, principal investigators behind the Adaptable I/O System (ADIOS), an I/O middleware that has helped researchers achieve

scientific breakthroughs by providing a simple, flexible way to describe data in their code that may need to be written, read, or processed outside of the running simulation.

BP invited Klasky and Podhorszki to its Houston offices to give the company's high-performance computing team a tutorial of ADIOS and demonstrate how it could help them accelerate their science by helping tackle their large, unique seismic datasets.

"The workshop was awesome," said BP HPC Technology Analyst Bosen Du. "It was a great introduction to ADIOS, and we definitely saw plenty of possible opportunities to apply it to our specific challenges. Even better, Scott and Norbert asked specific questions to personalize the tutorial to BP."

Klasky shared Du's enthusiasm. "This was the one of the more enjoyable tutorials we have given due to the level of interest from everyone in the room," he said, adding that BP's interest led to what is likely the longest tutorial the team has ever given.

A natural partnership

Klasky and Podhorszki's trip was the result of a growing relationship between ORNL and BP.

BP's Director of HPC, Keith Gray, was already familiar with ORNL's Oak Ridge Leadership Computing Facility, the DOE Office of Science User Facility that is home to Summit, through the positive testimonials of colleagues who had participated in its Industrial Partnership Program ACCEL (Accelerating Competitiveness through Computational Excellence).

Gray even visited ORNL two years ago to give a guest lecture on how

BP's data center needs are smaller but similar to those of a center like the OLCF and on the importance of a reliable data center to support BP's commitment to being at the forefront of supercomputing technology.

That relationship, along with ADIOS's unique capabilities, made the choice an easy one. "We started doing research and ADIOS was always at the top of the list," said Gray, adding: "By collaborating, BP's world-class expertise in applying HPC to solve complex scientific problems could help the ADIOS team understand different workflows as they help us manage our data."

Managing that data is critical from a business perspective. In one recent project the BP team faced a 500-terabyte dataset. And that's before seismic processing, after which the dataset can grow ten-fold.

"Having something that can scale, do massively parallel I/O, and support compression would be a major advantage in helping us overcome our current data issues," said Bashkardin. MGARD, a technique developed jointly by ORNL and Brown University that is used for lossy compression of scientific data and which mathematically guarantees error bounds, seemed a particularly good fit for BP's compression issues, said Klasky.

He added that recent changes in ADIOS, made possible by the Exascale Computing Project, have helped the SPECFEM3D-Globe seismology code used by Princeton's Jeroen Tromp achieve a speed of more than 2 terabytes per second while writing data to Summit's general parallel file system. Such a speed could lead to further collaboration with Tromp's team, which utilizes ADIOS as the I/O backend, and help strengthen the data processing capability for a large part of the seismology community.

Overcoming issues such as I/O bottlenecks means a reduction in data

analysis turnaround time, which would allow the company to explore different ideas, identify and address bottlenecks, and achieve a better understanding of the subsurface. Taken together, these capabilities can create huge breakthroughs for BP's research program.

But a successful implementation of ADIOS into BP's current I/O code, dubbed the Data Dictionary System, would be beneficial in the short run as well. For instance, it would give their team valuable insight into whether they are pursuing the correct technologies and strategies to succeed.

"It may help us consider building additional file systems to deliver more bandwidth than our current clusters," said Gray, adding that "you don't need new file systems if your I/O is at peak, and we currently don't have all of the necessary I/O metrics." Researchers from the ORNL team have agreed to provide some support in helping BP to assess its data strategy.

Added Bashkardin: "We struggle with extracting I/O bandwidth out of our Lustre file system due to a number of factors. There's lots to be gained in these terms. Even doubling the performance with a single dataset would be an enormous improvement."

In theory, ADIOS could expedite some jobs from days to hours, fundamentally altering the workflows of BP's seismic researchers. And, according to BP HPC Computational Specialist Qingqing Liao, the middleware's built-in visualization capability is an excellent tool that pinpoints problematic areas of researchers' codes and models to help them best understand how to alter their algorithms. Klasky credits his colleagues Lipeng Wan and William Godoy for this capability, which allows users to instantly transition from file-based code coupling (e.g. asynchronously coupling a code to visualization) to in-memory coupling without changing their code.

But before ADIOS can be implemented, the BP team will need to specify what viable features they want to see on their I/O backend and create a new API layer with a specific set of API goals.

"Being able to leverage ORNL's ADIOS and working together to improve it will extend BP's expertise in using big [data](#) to solve critical energy problems," said Gray.

Provided by Oak Ridge National Laboratory

Citation: BP looks to ORNL, ADIOS to help rein in data (2020, January 3) retrieved 10 April 2024 from <https://techxplore.com/news/2020-01-bp-ornl-adios-rein.html>

<p>This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.</p>
--