ON CLOUD NINE

HPC in the cloud reduces runtime for a complex multiphase CFD model with realistic particle loading from five days to two days.

By Sam Zakrzewski, Fluid Dynamics Specialist, FLSmidth A/S, Copenhagen, Denmark
and **Wim Slagter**, Lead Product Manager, ANSYS, Inc.

Large industrial plants and their need for higher accuracy are increasing the computational intensity of simulation in process industries. For example, at FLSmidth, a multiphase flow model of a flash dryer with realistic particle loading in the ANSYS CFX solver takes about five days to run on its local infrastructure. FLSmidth was interested in reducing the solution time and, if possible, increasing mesh size to improve the accuracy of its simulation results without investing in a computing cluster that would be utilized only occasionally. So, the decision was made to participate in the Uber-Cloud Experiment, which was initiated by a consortium of 160 organizations and individuals for the purpose of overcoming roadblocks involved in remotely accessing technical computing resources in high-performance computing (HPC) centers and in the cloud.

FLASH DRYER APPLICATION

FLSmidth is the leading supplier of complete plants, equipment and services to the global minerals and cement industry. The company recently designed a flash dryer for a phosphate processing plant in Morocco. The dryer takes a wet filter cake and produces a dry product suitable for transport to markets around the world. The flash dryer procedure was designed by FLSmidth's process department; the structural geometry was created by its mechanical department based on engineering calculations and previous experience.

Before investing large amounts of money and time to build the dryer, it is essential to verify that the proposed design will deliver the required performance and evaluate if alternatives can be found that cost less to build and operate. This dryer is the largest ever built, so FLSmidth used ANSYS computational

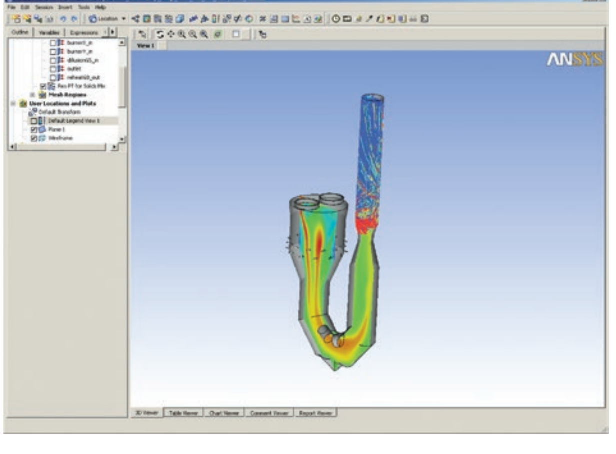

Flash dryer model viewed with ANSYS post-processing tool

Uber-Cloud Experiment

The aim of the first round of the Uber-Cloud Experiment was to explore the end-to-end process of accessing remote HPC resources and to study and overcome the potential roadblocks. The project brought together four categories of participants: industry end-users, computing and storage resource providers, software providers such as ANSYS, and HPC experts. ANSYS was engaged with multiple teams, and this article is a summary report of one of the teams.

fluid dynamics (CFD) tools in addition to extensive pilot testing to evaluate the operation of the dryer prior to installation. This approach helps identify designs that reduce construction and operating costs and reduces the risk of having to modify the dryer during the installation phase.

COMPLEX FLOW SIMULATION PROBLEM

ANSYS CFX is considered by FLSmidth to be the tool of choice for flash dryer design because it provides the breadth and depth of physical models and interaction among these models needed for the demanding multiphysics requirements of designing dryers and other process equipment.

Accurately simulating the performance of a flash dryer requires modeling the flow of gas through the dryer, tracking the position of particles as they move through the dryer and calculating the moisture loss from the solid particles. For example, the multiphase flow model of the dryer used in the phosphate plant employs Lagrangian particle tracking to trace five different species in time steps of 1 millisecond for a total time of 2 seconds.

FROM DESKTOP TO CLOUD

As mentioned, the model originally took five days to solve on a local machine with a Xeon® processor running at 3.06 GHz with 24 GB RAM. FLSmidth recently ran this same model on a cloud solution provided by Bull extreme factory (XF) with technical assistance from science + computing. The model ran on 128 Intel® E5-2680 cores in the cloud in about 46.5 hours. The cloud solution demonstrated the potential to run models faster, increasing the speed of sensitivity analysis while reducing the amount of internal resources that need to be devoted to IT and infrastructure issues. The XF resources used in this project are hosted in the Bull data center at the company's headquarters outside Paris; they feature Bull B510 blades with Intel Xeon E5-2680 sockets with eight cores, 64 GB RAM and 500 GB hard disks connected with InfiniBand® quad data rate (QDR) serial links.

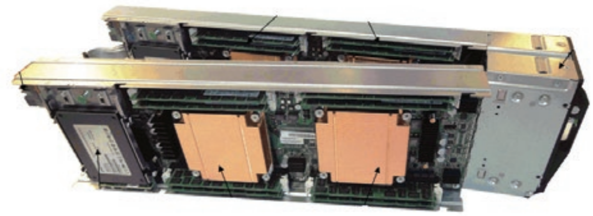

Per Node

- 2 Intel E5-2680 – 2.70 GHZ sockets
- 8 cores per socket – 8.0 GT – 20 MB – 130W
- 64 GB RAM = 4 GB/core DDR3 1600 MHZ
- 2 x 500 GB SATA hard disks
- 1 Infiniband QDR port (40 Gb/s)
- 2 x 1g Ethernet ports

▲ Compute nodes deployed by extreme factory

This project demonstrates the feasibility of migrating powerful computer-aided engineering applications to the cloud.

– Marc Levrier, HPC Cloud Solution Manager, XF

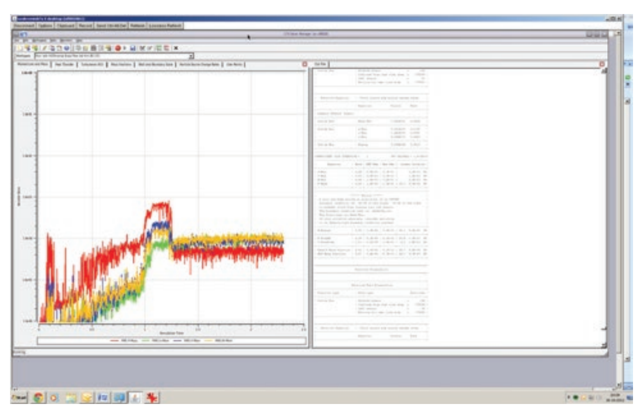

▲ ANSYS CFX solve manager

RUNNING SIMULATION IN THE CLOUD

The XF team integrated ANSYS CFX into its web user interface, which made it easy to transfer data and run the application. The XF team spent around three man-days to set up, configure and execute the ANSYS CFD experiment. science + computing provided technical assistance in setting up the problem to run in the cloud. FLSmidth engineers spent around two man-days to understand, set up and utilize the XF interface methodology. Users were able to easily manage ANSYS licenses through the interface. "This project demonstrates the feasibility of migrating powerful computer-aided engineering applications to the cloud," said Marc Levrier, HPC Cloud Solution manager for XF.

SUBSTANTIAL TIME SAVINGS

The ANSYS CFX solver was designed from the ground up for maximum parallel efficiency. All numerically intensive tasks are performed in parallel, and all physical models work in parallel. Administrative tasks, such as simulation control and user interaction, as well as the input/output phases of a parallel run are performed in sequential mode by the master process. This approach guarantees good parallel performance and scalability of the parallel code. It also ensures that the input/output files are the same as those produced by a sequential run.

A few early runs failed due to hardware or software glitches, but soon the model was successfully solved on 128 cores. The runtime of the successful job was about 46.5 hours — so FLSmidth's primary goal of running the job in one to two days was met. The XF team installed ANSYS CFD-Post visualization software and made it available from the portal in a remote 3-D visualization session; it was used to view the simulation results. The users monitored the runs with the solver manager user interface, avoiding the need to download large output log files.

TAKE-AWAY POINTS

This experiment demonstrated that the initial deployment of HPC applications in a cloud requires a lot of experience and planning, plus a team to both deploy and tune applications and support software users. For example, some hardware provisioning delays were experienced because pressure from production made it difficult to find free resources and tune them to get good results.

Within the first round of the Uber-Cloud Experiment, there was not enough time available to perform scalability tests with the CFD solver. These would have been helpful in determining how additional resources could have been deployed to further reduce the runtime of the job. Due to the size of output data and transfer speed limitations, the users concluded that the remote visualization solution was necessary. Remote visualization in a Windows®-based environment means that a dedicated visualization server is required to render the results in a timely and usable manner. Users also would like to have access to more cluster metrics than are currently readily available, such as CPU, memory and I/O usage. In the end, this was a very promising experiment that demonstrates the potential of the cloud to reduce the time required to solve large CFD models while avoiding the need for investing in and administering HPC resources.

Initial deployment of HPC in a cloud requires a lot of experience and planning, plus a team to both deploy and tune applications and support software users.

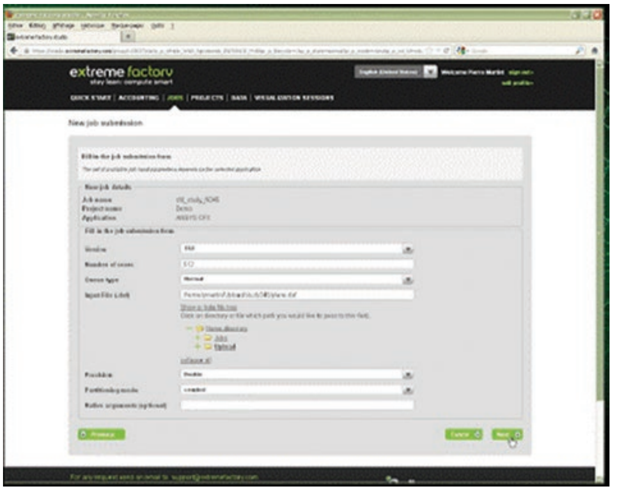

ANSYS CFX job submission web form

This experiment demonstrated the potential of the cloud to reduce time required to solve large CFD models while avoiding the need for investing in and administering HPC resources.