HPC User Forum Update

Gaining a Competitive Edge with HPC Cloud Computing at DOW, HPC User Forum Dearborn Michigan

Bob Sorensen
April 2019

IN THIS UPDATE

The HPC User Forum was established in 1999 to promote the health of the global HPC industry and address issues of common concern to users. In September 2018, Hyperion Research hosted its 70th User Forum in Dearborn, Michigan. This update summarizes a presentation from that meeting entitled, *Gaining a Competitive Edge with HPC Cloud Computing at DOW,* given by William J. Edsall, R&D Information Research - High Performance Computing Team, DOW.

Edsall described how DOW is looking to develop and then leverage a hybrid HPC On Premise/Cloud Bursting environment by 2019. Edsall presented the details of their recent Proof of Concept for their cloud-bursting test capability with Rescale and Microsoft that consisted of a computational chemistry job using 14,000 molecules and a typical CFD day to day simulation. Edsall noted that both tests were successful and proved that DOW could benefit both technically and financially from wider implementation of a hybrid HPC capability.

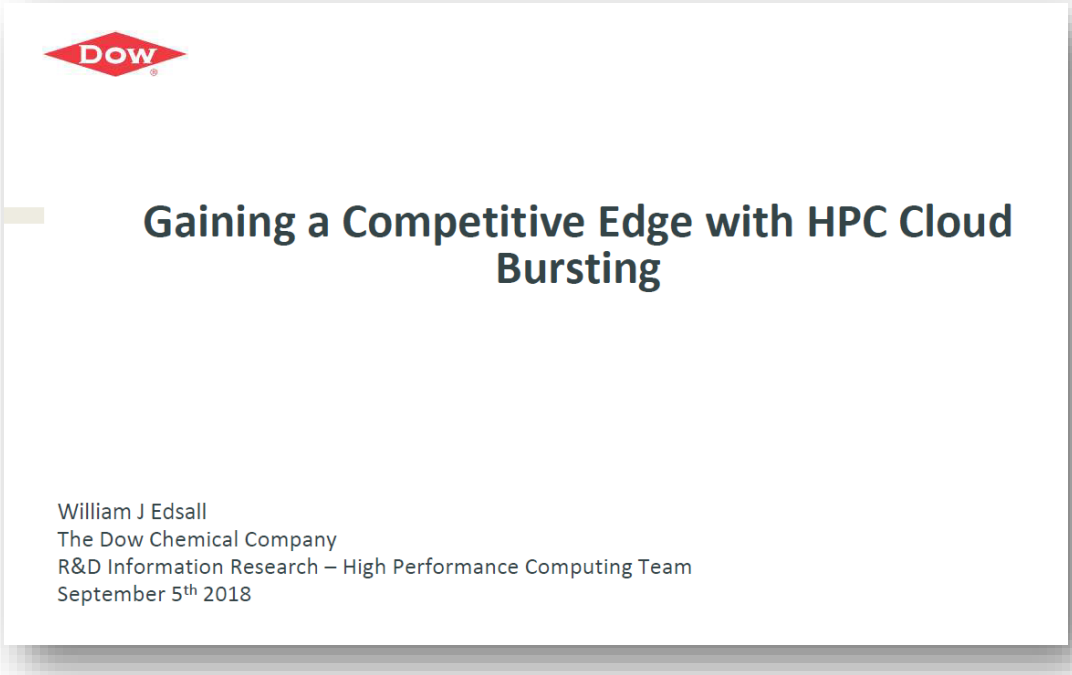

PRESENTATION: GAINING A COMPETITIVE EDGE WITH HPC CLOUD COMPUTING AT DOW, GIVEN BY WILLIAM J. EDSALL, DOW

Edsall shared his industry perspective on how DOW is exploring the option to transform to a cloud hybrid on-premise environment. DOW sees HPC as critical and thinks of it as a computational laboratory not much different from a physical lab with the added bonus of global accessibility.

Edsall noted that currently, DOW has high HPC utilization of their on-prem computers at 85%, that justifies their exploration in future growth opportunities for HPC system access. DOW uses general purpose clusters, and they don't build HPC systems for one specific application. Their systems must be able to run CFD, molecular-dynamics computational chemistry, and everything else that is needed.

- Although DOW uses InfiniBand interconnect for their CFD work and high-memory bandwidth for the computational chemistry, their systems are not custom tailored for any of these environments.
- They also have some deep-learning and machine learning going on with their NVIDIA servers. They procured an NVIDIA DGX1 that's considered a valuable addition to their overall HPC environment.
- Rounding out their computational collection, DOW uses high-end HPC desktops.

FIGURE 1

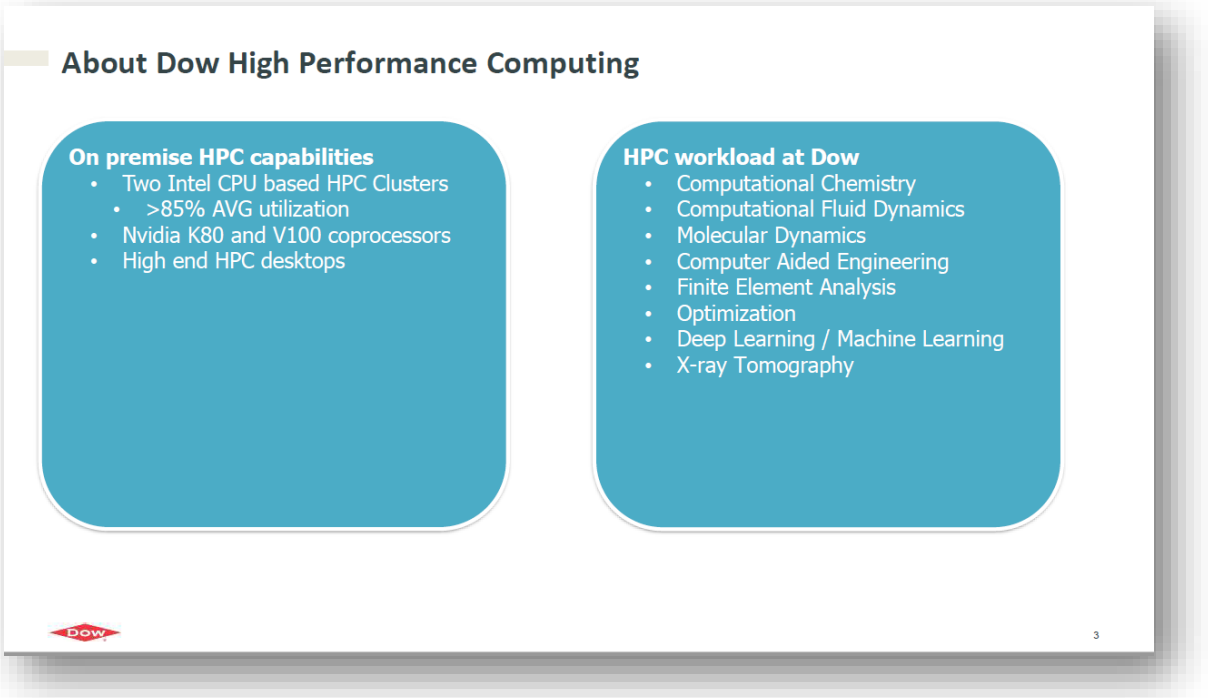

Source: DOW and Hyperion Research, 2019

Moving forward, DOW is looking to develop and then leverage a hybrid HPC on-premise/cloud bursting environment. To that end, DOW did a proof-of-concept with Rescale and Microsoft. Some of the important steps DOW had to go through are unique to its industry, such as the development of a specialized security model and a cost recovery model.

With respect to the Microsoft Azure implementation, DOW had several key requirements:

- DOW wanted the most secure environment, using a single sign-on and an environment that can use all the tools that their corporate IT is happy with.
- For ISV licensing, the system must understand what solvers are going to be used in the cloud and if their licenses are going to marry up with what they are using on premise; if one solver keeps DOW from running on the cloud, then that must be dealt with.

Rescale worked with the key ISV's to come up with a special pay-as-you-go or bring-your-own-license models, as well as addressing data transfer issues between DOW and Azure. Here, the design called for an express route, which is a dedicated network between the company and the cloud environment, which DOW maintains.

Driving the overall design for this proof-of-concept effort (see Figure 2 heat map) is the marriage of CPU's or CPU hours (x axis), from just a couple to several thousands to hundred-thousands, and the urgency of the job (y axis) with DOW. Edsall indicated that this is a good representation of the full range of computing that they could do at DOW and how it might look today.

FIGURE 2

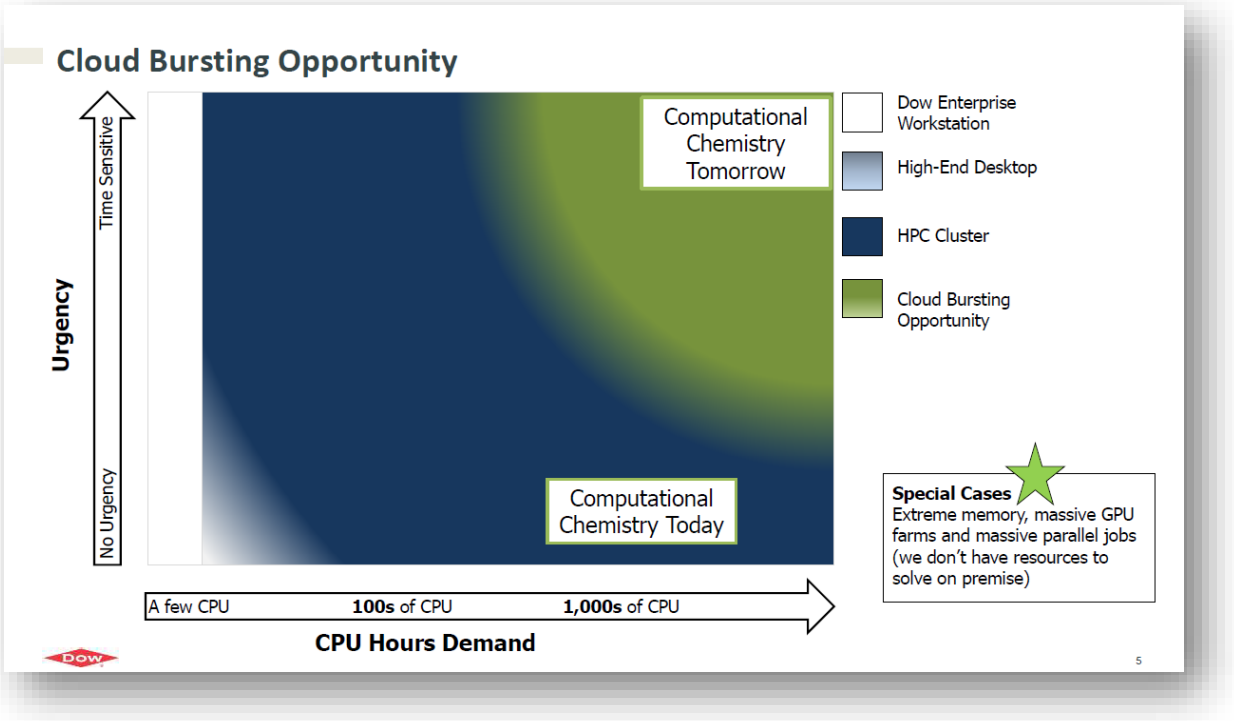

Source: DOW and Hyperion Research, 2019

- Small jobs, like a small computational chemistry job, could easily land on the few CPU range with their DOW Enterprise Workstation, which is what they call their laptops.
- Just beyond that, the high-end desktop range would require a couple more CPUs. Any amount of urgency can be addressed on those. Because they are small, they can be started immediately, and the computer can be returned to general availability.

Edsall noted that where it starts to get interesting is their cluster space, shown as the blue space in Figure 2. Here jobs can run in the hundreds to near-thousands of CPU range at any time with high-urgency but without much work. They might have to juggle some work around or make some special cases and, in an hour or two, the space to run is available.

Where it really starts to get challenging, according to Edsall, is when they get into the thousands to tens of hundreds or even millions of CPU hours. The computational chemistry work at DOW is one of their highest workloads, and the green quadrant is seen as DOW's opportunity statement for cloud bursting and where DOW would want to have a strong hybrid capability.

Edsall presented the details of their recent Proof of Concept for their cloud-bursting test capability with Rescale and Microsoft. They used two important use cases:

1. One was the computational chemistry POC seen on the left of Figure 3 that used an input deck of ~14,000 computational chemistry molecules (under 50 atoms per molecule). The exercise was not to benchmark speed or efficiency, but instead to see if they could burst to something more massive than they could get done on premise.
2. Their second objective was to see if they could run a day-to-day CFD simulation in the cloud similar to a small simulation they would do on premise.

For computational chemistry jobs, each one was run on a seven CPU Azure instance (H16r). At one point, testers spun up close to 10,000 cores. It took less than 30 minutes just to stand the whole thing up. At the time, they had 500 CPUs not busy on premise. It was 13 or 14 times faster to do this in the cloud, and the overall test goal was met.

The second case was the CFD test case. Researchers scaled up 2,000 CPUs. The input deck, in this case, did not justify 2,000 CPUs, but researchers wanted to see if it was possible. The test didn't run completely through its paces because it wasn't necessary. Researchers weren't benchmarking speed, but feasibility. The goal was met in 15 minutes. To carve that amount of resources on an inhouse cluster would have been, at minimum, a half-day effort simply to figure out where they were going to take the sacrifice from and how to justify that.

FIGURE 3

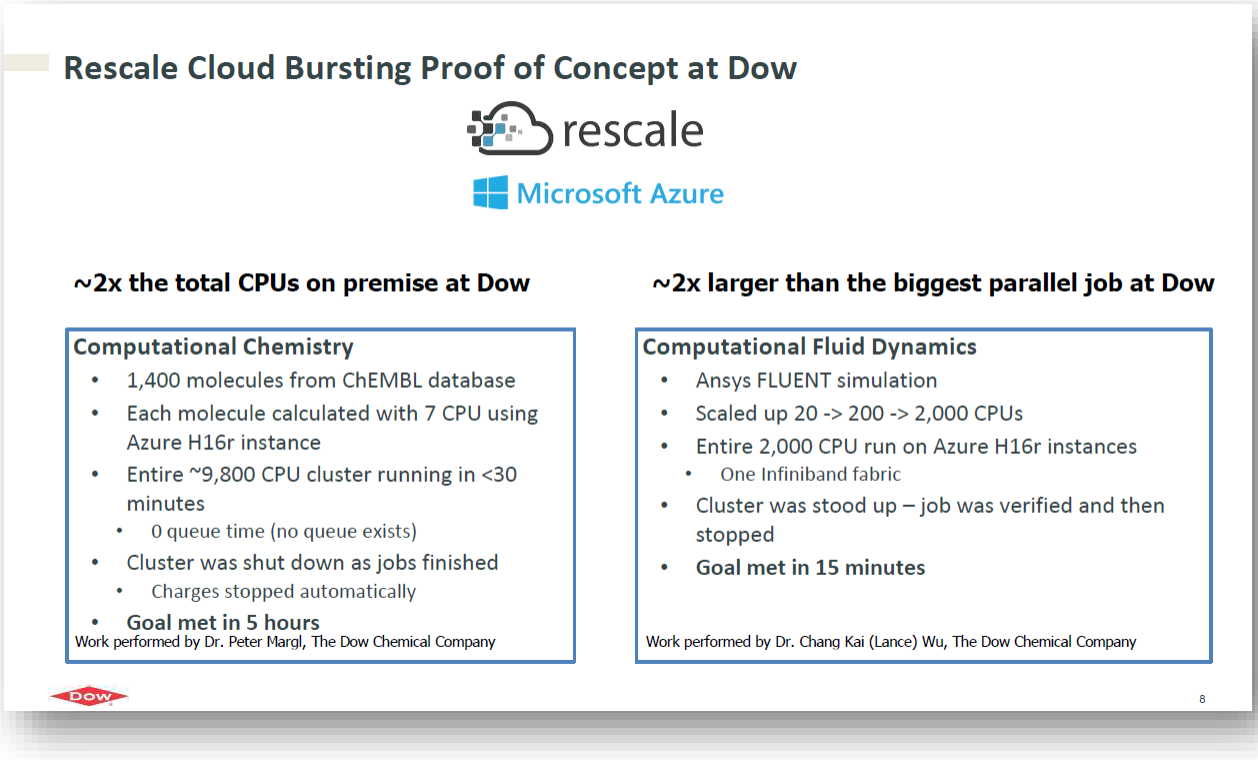

Source: DOW and Hyperion Research, 2019

In this POC, the street cost of those hours was $2,500. Edsall highlighted that based on the speed-up and knowing how much capacity DOW has internally, a job like this could take a one-year simulation down to a week or even less.

Edsall concluded by saying that 10,000 cores may not be a lot to most of the people in the HPC sector, but it still is a significant amount of hardware. He noted that Univa published an article about their deployment of 1,000,000 CPUs in AWS. The whole job was ready in two and half hours. Although 10,000 isn't a big deal, 1,000,000 certainly is.

- This announcement was a comfort to Edsall, coming out a week after DOW did their POC. For DOW, their realistic burst case would be, perhaps 50,000 CPUs. This shows them what is possible in the near future.

For more information or to view this and other presentations given at HPC User Forums dating back to 2008, visit www.hpcuserforum.com.

About Hyperion Research, LLC

Hyperion Research provides data-driven research, analysis and recommendations for technologies, applications, and markets in high performance computing and emerging technology areas to help organizations worldwide make effective decisions and seize growth opportunities. Research includes market sizing and forecasting, share tracking, segmentation, technology and related trend analysis, and both user & vendor analysis for multi-user technical server technology used for HPC and HPDA (high performance data analysis). We provide thought leadership and practical guidance for users, vendors and other members of the HPC community by focusing on key market and technology trends across government, industry, commerce, and academia.

Headquarters

365 Summit Avenue
St. Paul, MN 55102
USA
612.812.5798
www.HyperionResearch.com and www.hpcuserforum.com
