

# HPC

## *In the Cloud*

### Expanding enCORE with CFD

July 19, 2013 | Ian Armas Foster

---

Britain HPC systems integrator OCF has made noise with regard to high performance applications in the cloud, bringing online an 8,000-node HPC cluster called enCORE last year. This week, they announced an expansion of enCORE's capabilities. Specifically, they brought in the XFlow Computational Fluid Dynamics (CFD) software to expand their HPC capabilities.



The enCORE New Business Development Manager for OCF Jerry Dixon spoke with HPC in the Cloud on the challenges of bringing this CFD software to enCORE as well as the prospects of OCF and enCORE going forward as they look to expand their reach deeper into the manufacturing field, as well as other cloud-relevant HPC fields.

“We’re now starting to extend our horizons, to take on larger organizations with significantly bigger compute and storage requirements,” Dixon said.

According to Dixon, the project took off when he met Mark Hieatt, the commercial director of FlowHD, the company that sells XFlow in the UK and Ireland, at an event held at the Hartree Centre, where the enCORE server cluster lies. Further, the two paired with Dragon HPC, a system that aids with the visualization of large datasets, especially helpful in a manufacturing context. The collaboration has already seen movement in the automotive sector, according to Dixon.

“The engagement with Dragon HPC, pretty much a state of the art remote visualization company, together with the Hartree Centre, we operated to put this solution together, which is initially being used by a very significant automotive vendor,” Dixon said.

Usually, introducing a system like the XFlow CFD software into an 8,000-core cluster like enCORE requires a not insignificant time cost. However, Dixon noted that Hieatt’s experience working with HPC clusters made the transition relatively seamless. “Working with Matt and his team, we have pretty quickly got XFlow running successfully on the cluster,” Dixon said.

What made FlowHD interested in working with enCORE was the server cluster’s ability to properly scale, even at 1,000+ cores. Cloud environments, even ones designated for HPC workloads, do not always scale well, as servers tend to be farther apart the more numerous they are. However, enCORE’s servers are fairly tightly integrated, as noted in this previous [article on HPC in the Cloud](#) written by Tiffany Trader.

“Each of the 512 nodes sports Intel SandyBridge 8-core CPUs with either 36Gb or 128Gb RAM. This is true HPC-as-a-Service; the compute nodes are not virtualized. Service users can also access 48 GPU nodes outfitted with NVIDIA Tesla 2090 GPUs,” Trader wrote of enCORE in October of last year.

As a result, scaling happens much as it would on a typical HPC system (one that was not designated as a cloud, that is). “We have achieved very close to linear scalability, which is an important factor to [FlowHD],” Dixon said. “Matt and his team run tests likely to generate on about 1,000 cores so the performance scale-up is as close to linear as you’re likely to get. The performance from their point of view was more than acceptable.”

Those performance gains are happening on their preliminary workings with XFlow on enCORE, as Dixon

mentioned that currently four instances are provisioned currently. The plan is to scale that out to their planned extent using ScaleMP, with which tests have already taken place.

When one has successfully made available the amount of cores that OCF has in 8,000, as an HPC-as-a-Service, the challenges and problems shift from virtualization concerns, such that exist in AWS HPC instances, to things like ensuring security and what Dixon referred to as 'dynamic licensing.'

Interestingly, Dixon noted that, while security was of course still an issue, his clientele has started to accept the notion of computing in the cloud. "Security has always been an issue and always will be, but in the last few months talking to customers, I think it's becoming less of an issue. There's more of an acceptance of cloud technology, using remote infrastructure," he said.

Licensing, on the other hand, is another issue entirely. According to Dixon, the limitation comes in the form of costs rising disproportionately to the number of CPU cores used. This limitation exists generally with regard to commercially available software like XFlow, created for usage by manufacturing vendors and the like.

"With commercial software, the limitation is very often the fact that the license is based on a number of CPU cores that you run on," Dixon said. "If you want to run on that on multiple cores, you run into significant costs to operate a license. In many cases, it makes it economically [unsustainable]."

Licensing agreements add to the bevy of obstacles an institution like OCF must hurdle, including the volume of datasets that would plague most every HPC cloud provider. While some of those datasets are still transmitted by sending physical hard drives through the mail, Dixon noted that the partnership with visualization platform Dragon HPC eases the process of paring down ones data for submission to cloud systems such as enCORE.

The results so far on that front have been promising. "Typical engineering codes, they cannot produce large datasets, that's where the remote visualization comes in here with Dragon HPC," Dixon said. "The performance we've seen from that has been really good."

Beyond CFD, OCF and enCORE look to incorporate more engineering and manufacturing codes into their system, from both the commercial and open source markets.

---

Copyright © 1994-2013 **Tabor Communications**, Inc. All Rights Reserved.

HPC in the Cloud is a registered trademark of Tabor Communications, Inc. Use of this site is governed by our Terms of Use and Privacy Policy.

Reproduction in whole or in part in any form or medium without express written permission of Tabor Communications Inc. is prohibited.

Powered by **Xtenit**