The team at Verne Global and I are really looking forward to ISC18 – the European high performance computing (HPC) conference in Frankfurt, where our industry meets. It's an interesting and engaging show that annually highlights the tremendous developments taking place across many sectors and only made possible by greater adoption of HPC. We're especially looking forward to publicly introducing hpcDIRECT - our superb bare metal HPC platform - to the international HPC community.

Once again it's a packed program and personally I have penciled in the keynote from **CERN openlab's** CTO **Dr. Maria Girone** and the talks from **Mark Parsons** (EPCC University of Edinburgh) and **Dirk Pleiter** (Jülich Supercomputing Centre) as ones to watch. It is also good to see *"Industrial day"* and *"Machine Learning day"* coming into the agenda – more on the latter later in this blog.

While I've been busy making my preparations for ISC18 it gave me a chance to review the status of the HPC market, which has been experiencing strong, steady and diverse growth. According to Hyperion Research *(if you can get up early, their HPC Breakfast Update is worth dropping into on the Tuesday),* the market for HPC related services will continue to grow at a compound annual growth rate (CAGR) of 7.8% from 2017 to 2021, and is projected to reach over $14 billion dollars in the next five years.

Interestingly, the HPC industry is playing an increasingly larger role in the world economy, as reported by HPCWire: [HPC-reliant US economic sectors contribute almost 55 percent of the total US economy, encompassing $9.8 trillion and accounting for over 15.2 million jobs](). Another thing that's interesting about the growth is it isn't confined to the traditional HPC customer base of scientific research, oil and gas exploration or computer-aided engineering, but can be attributed to a number of different changes in the way that HPC is being perceived, sold, and applied.

One factor responsible for this growth is the increased accessibility of HPC technology over the last two decades. Historically, the HPC infrastructure would need to be purchased by the user and a skilled team of engineers, techs and end users recruited to set it up and get it working. Commercially available HPC solutions (such as our hpcDIRECT platform) have streamlined the deployment of HPC systems, lowering the barrier of entry for companies to begin using HPC, and helping to change the image that HPC is an abstruse, inaccessible science. This has been compounded by the advent of HPC computing in the public, hyperscale cloud. Today, companies have a fast, flexible means of accessing HPC resources that they didn't have even five years ago. This democratisation has also helped to slash the hourly per-core price of HPC resources from about $1.50 in 2007, to less than 10% of that cost today, clearing the way for a more democratic HPC adoption.

On the subject of HPC in the public cloud – intriguing debates are evolving, and I read an [excellent article recently in Computer Weekly]() by **Francis Lam**, director of HPC product management at Huawei. He made the case for dedicated HPC infrastructure, despite the growth of public cloud services and it's a view I would agree with. Public cloud HPC has its place and for certain applications and workloads it provides a convenient and elastic service. However, for more attuned, intensive HPC workloads a more 'genuine' or _True_ HPC solution utilising bare metal, fast interconnects for low latency and all clustered in one location (as opposed to virtualised servers) is preferable. Anyway – more on this debate at ISC18, and if you're interested in this angle do come and listen to my colleague **Tate Cantrell**'s [presentation in the Exhibitor Forum]() on Tuesday 26th at 4.20pm CET.

Another key factor that's spurring HPC growth is the diverse industries that have started to ramp up their HPC usage. Take for example the highly-competitive consumer goods manufacturing industry. Although not traditionally a major HPC customer, these manufacturers are turning to HPC in greater numbers to facilitate the development of prototypes for innovative products and make optimal use of new materials. Procter & Gamble,

for example, has used HPC clusters to help develop better products for over a decade, but the size and scale of its HPC cluster has grown dramatically in recent years, from around 128 cores

at the dawn of the millennium, to what is today many thousands of cores. Although P&G keeps the technical details of its systems under wraps, the company has been open about how they use their HPC capacity to test a variety of products, including their electric shavers, laundry detergents, and even nappies!

Many consumers may not realise that their favorite products have been carefully optimised using HPC-driven computational analysis like CFD and finite element analysis (FEA). Even the humble potato chip has been subjected to the rigours of HPC! P&G are not alone in this endeavor. Other consumer goods companies like Unilever have also committed to HPC to expedite the product discovery process, recently establishing a base at the UK's Hartree Centre ... ...boratory to do so.

HPC's steady growth is also to do with the proliferation of machine learning, deep learning, and other artificial intelligence (AI) technologies, which are delivering a new level of insight and flexibility to industries that haven't used HPC before. Take for example the cyber security industry, where a wave of new companies have started using machine learning to derive value from the security log data generated by company networks. Before the proliferation of HPC-enabled machine learning, analysing the security logs in medium- and large-sized enterprises was functionally impossible to do with any degree of efficiency. That isn't quite the case anymore, as companies like Anodot, which recently tripled its revenue, and Dynatrace, a leader in the multi-billion dollar Application Performance Management (APM) industry, have appeared to meet that need.

From where I sit, there is still quite a gap between the HPC infrastructure purists – all of whom will be walking the corridors and aisles of ISC18, and the AI and machine learning fraternity who appear to be tackling their compute challenges in a different way, driven by the application first and foremost. This is why *"Machine Learning day"* at ISC18 is so crucial and elements like this will help to converge these two technologies further.

Blurring the borders that once separated the HPC from other industries is inspiring both creativity, entrepreneurship, and helping to keep growth within the HPC industry strong. Our team at Verne Global is enthusiastic about the prospect of the HPC industry further developing and we're confident our hpcDIRECT platform using latest spec **Intel** hardware will

help spur wider adoption of HPC in a more diverse range of industries, creating a virtuous HPC ecosystem that benefits everyone.

If you are attending ISC18, please come along to our stand (F-912) and I'd be delighted to meet you. You can contact me at spencer.lamb@verneglobal.com.

Share:

Sign up for the Verne Global newsletter

Opinion, thought leadership and news delivered directly to your inbox once a month.

Enter your e-mail

SIGN ME UP

MEDIA KIT   **Privacy Policy** | Copyright © Verne Global 2018. All rights reserved.