



## HPC Innovators – Adrian Wander

Adrian Wander is the Chief Executive Officer and Director of **A Wander Consulting**. With over 30 years' experience of working in the high-performance computing and IT sectors, Adrian is a well known and highly respected figure within the international HPC community. He has worked across all areas of IT including HPC, AI, big data, and data centres. Adrian has been intricately involved with a full spectrum of activities ranging from the development of long-term IT strategies through to the management of highly technical procurements in the UK university, government, research council and international sectors.

**His vast experience made for an interesting and wide ranging conversation with Kao Data about the state of the HPC industry, the need for colocation services and the impact of moving HPC to the cloud.**

### 1. How did you get started in your career?

I'm a research scientist and I've got a PhD in theoretical solid state physics. I began using HPC to run calculations at the **University of London Computer Centre (ULCC)** when I was a graduate student. I did some post-graduate work in the US at the **Lawrence Berkeley National Laboratory** at the University of California, Berkeley campus. Upon my return to the UK, I was a chemistry lecturer at **Cambridge University**, working under **Professor Sir David King**, who would later go on to become the chief scientific advisor to Tony Blair's government.

The next part of my career was with the **Science and Technology Facilities Council** where I went on to be Director responsible for managing the Scientific Computing Department that totaled over 190 staff split between the Daresbury and Rutherford Appleton

Laboratories and operated on a budget close to £20M per annum. I was also heavily involved in setting up the **Hartree Centre** which is home to some of the most advanced computing, data and AI technologies in the UK. I then moved to the **European Centre for Medium-Range Weather Forecasts (ECMWF)** as it provided an opportunity to expand my experience beyond research and into operations. My team delivered computational services both internally to ECMWF and externally to users from member states. These were highly critical, time sensitive computations that relied heavily on HPC modeling for speed and accuracy.

### 2. What are the most critical issues you are seeing your clients face today?

I think one of the biggest problems facing traditional HPC research users is in regards to hosting HPC systems. At ECMWF, we had sufficient power to run these systems, but insufficient cooling, so we ended up moving the data centre as we could no longer support it on-site. Organisations are running out of power, and they're running out of space for cooling, and particularly for some London-based organisations, the floor space is expensive.

The post Covid-19 business environment is going to impact this as well as big model offices are likely to disappear. These big monolithic offices come with big monolithic machine rooms. As we see people moving to smaller offices, or hot desking as they go in for meetings while working from home, what will happen to all of the IT equipment and data centre infrastructure on-site? I think there will be a big move in the commercial sector to look at colocation solutions. Time will tell, but Covid-19 will continue to present some interesting challenges and opportunities. Many of the **HPC facilities across Europe, and across the world, have been giving away cycles for Covid-19 research**. There's been a huge upsurge in the life sciences via simulation and modeling of protein folding and the structures of proteins and those kinds of areas. And, of course, a lot of AI is looking for new drugs by trying to automatically sample huge numbers of drugs on the target treatment. I don't see this slowing down in the life sciences sector any time soon.

### 3. How do you see HPC moving to the cloud?

One of the big pushes over the next few years will be the democratisation of HPC. Historically, HPC has been seen as the domain of specialists. People who actually understood HPC already had HPC systems, and other people just thought it was too complex. Ease of use, particularly for SMEs, is going to be imperative. If you look at some of the genome sequencing now, the Oxford nanopore system takes you down into quite small organisations that are doing this kind of work, but they need the compute to do the genome assemblies. They want to focus on the research, not the IT infrastructure. A move to the cloud eases some of those constraints for those types of organisations, but there are other challenges still to be addressed.

One of which will be in the movement of data. If you look at weather forecasting, it's a huge output file. It's a 9km grid over the surface of the earth with 137 vertical levels on each grid point. Do you really want to be moving millions of observations into the cloud running a forecast and then moving huge data sets back to your home? Probably not.

There is also the fact that while many of the HPC research centres are running mixed workloads, some centres do not. Ninety percent of the workload that was run on the ECMWF system was an integrated forecast system based on a single code. There is a huge benefit to be gained from optimising that code to a very specific hardware and software stack which, when using servers in the cloud, is not always easy to do.

Ultimately, we will end up with some kind of cloud provision, but I don't think we are there yet. More and more people are looking at doing HPC in the cloud, but as for a cloud first strategy, they generally run up against a cost barrier and security concerns. So the question will be how do we get over that? How do we make HPC in the cloud cost effective? How do we not only secure the cloud, but ensure data transferring to and from the cloud is secure and cheap as well?

### 4. What does the future of HPC look like five years from now?

The simple answer is there will be a rise in AI, a rise in big data, and increasing convergence of HPC, AI and data workloads. People are trying to analyse more and more data and as the size of the training data gets bigger, the impact on the time taken to train deep learning systems will increase as well. We are just at the beginning of that phenomenon today. On the hardware side, we've got a large range of chips between Intel, AMD, Marvel and Fujitsu in the market, not to mention NVIDIA dominating the AI and GPGPU segments. They're all going to have applications that have users interested in the different performance characteristics. Whether that's sustainable over five years, I don't know. Currently, we're heading towards heterogeneous systems with accelerators, typically GPGPUs, but to fully exploit such systems will require us to do a lot of work on the software. So there's going to be the software engineering piece that needs to be looked at and this is before we fully understand the impact of new paradigms such as quantum computing as well.

### 5. What innovations are you seeing in the industry that you are most excited about?

HPC goes through cycles of technology pull and application push. We've been in a period where it's been a technology pull, but that's not necessarily what is really important. What's important is the output. It's solving problems. It's doing science. It's improving products. Within the industry itself one of the problems we need to solve is the energy consumption around HPC. The importance of having a genuine green agenda will be critical moving forward. As will demonstrating the solutions that are being created are environmentally friendly as well.

**For more details on Adrian's work and to contact his consultancy please go to: [awanderconsulting.com](http://awanderconsulting.com)**