



## How Integrators Ensure Top Performance for HPC Clusters

---

[insidehpc.com/2018/01/integrators-ensure-top-performance-hpc-clusters/](https://insidehpc.com/2018/01/integrators-ensure-top-performance-hpc-clusters/)

staff

January 16, 2018

*In this special guest feature from Scientific Computing World, Robert Roe looks at the role of integrators to deliver and support HPC systems.*

In today's HPC market, which provides a wide array of hardware options, the role of integrators is increasingly important as their expertise can help users to choose, set up and maintain HPC systems.

One example of the options available to HPC users can be seen in the processor market. Although Intel still dominates the field there are now credible alternatives from IBM's openPOWER, AMD and ARM. While there is no one correct choice for every user or application, integrators can help to select the technologies most suited to a client's budget and portfolio of applications.

Today a HPC system is much more than a collection of CPUs and, as the complexity of these systems increases, it is essential to balance computation with networking, storage and memory in order to deliver sustained application performance.

Integrators have become more important as HPC systems have become more complex. There is a requirement to balance the compute and the I/O to make the system efficient, and to get intelligent scheduling to ensure that time is not wasted – there isn't a tier-one provider that does all that," explained Julian Fielden, managing director of OCF.

### Selecting the right technology

It's the role of the integrator to bring together the best of breed technologies and ensure they are all balanced and work well together with ongoing support. This ongoing support is invaluable as our customers don't want to be dealing with numerous vendors on a regular basis," added Fielden.

Beyond the sale of components, and support and maintenance, Fielden also stressed that integrators provide expertise to help users get the most out of their computing system. 'It is not just the technology that is becoming more complex. Systems are now becoming

multi-purpose, and are required to be production machines serving internal customers – like an internal cloud,’ said Fielden.

Users increasingly want to do more with their systems; this could be reflected in cloud technology but also data portals or remote access to specific parts of the facility, or connections to other computing systems. This is particularly true in academic centres users but can also apply to enterprise users.

There’s a far greater requirement for things other than the ‘tin’ to make it work efficiently. That’s why there needs to be an integrator, as opposed to a reseller,” added Fielden. “The HPC system needs to be ‘designed’ not just sold for purpose. Without the services of a good integrator you risk getting a bag of bits, instead of a balanced system.”

It is not just computing hardware that integrators can help to set up when designing a new system, as demonstrated by a recent contract awarded to German based integrator Megware – which was selected to provide a hot water-cooled HPC system, ‘CoolMUC 3,’ at the Leibniz Supercomputing Center (Lrz) in Garching.

As the name CoolMUC 3 suggests, this is the third generation of of the CoolMUC computer cluster at Leibniz, which is part of the Bavarian Academy of Sciences and Humanities Garching research centre. This is the second time that Megware has supplied a CoolMUC system, which is designed to provide very high energy-efficiency.

The system was developed at the company’s technology development centre in Chemnitz and based on its own hardware and software. With the first generation of CoolMUC, Megware proved back in 2011 that it is possible to cool various processor technologies directly with hot water. In addition, an absorption chiller was used to efficiently reuse residual heat to generate process cooling, in order to cool existing servers in additional racks.

The new CoolMUC 3 system outperforms its predecessors in a number of respects, including its key feature of cooling in thermally insulated racks all compute and login nodes, power supply units, and Omni-Path switches directly with hot water, a combination the likes of which has never been seen before,” commented Axel Auweter, head of HPC development at Megware.

Hot water cooling relies on the idea that warm water does not need additional chilling equipment, so energy-efficiency can be increased as there is less of a requirement for cooling infrastructure. This can be further increased through the use of evaporative cooling or in the case of CoolMUC 3 an absorption chiller.

Even at a cooling water temperature of 40 degrees Celsius and a room temperature of 25 degrees Celsius, a maximum of just three per cent waste heat is produced in the ambient air. We’ve carried out a great deal of development work and are now more ready than ever to supply highly efficient, environmentally-friendly HPC technologies globally,” added Auweter.

## Disruptive technology

The CoolMUC 3 system supplied by Megware not only cools compute nodes but also features the latest generation of Intel Xeon Phi, which stands out for an integrated fabric and is directly interconnected via the Intel Omni-Path network.

With the huge amount of variation in HPC systems based on the kind of applications they might run, the number of users, hardware and the composition of the nodes system integrators must make sure they collaborate with their clients to ensure that the system is suited for the particular use case.

There is an increasing propensity for organizations to buy and grow systems, so we help our customers to think ahead about the future use of their system. It is important for customers to be mindful of where they might be in three years' time, where they are now and the path from one to the other. We discuss with them any future technology trends that may impact the way a system evolves and we have successful partnerships with World-leading vendors to help assist us in this exercise," stated Fielden.

"It is a collaborative approach working with customers to select the best technology for their needs. No two customers have exactly the same requirements; there are always individual nuances that may give them an extra five per cent efficiency on their particular HPC system," Fielden added.

The job of integrators is becoming increasingly complex as there are many new and disruptive technologies that can now be applied to new HPC systems. It is the job of an integrator to help select technologies that will help researchers accomplish their goals not just to sell them the latest kit. There is no reason to add technologies just because they are the flavour of the month – but, used in the right place, they can provide huge performance benefits.

One example of the use of disruptive technologies in HPC can be found in the recent contract from the Faculty of Mathematics at the University of Cambridge, which selected Hewlett Packard Enterprise to supply its latest HPC system.

The Cambridge Faculty of Mathematics will leverage this new system, in partnership with the Stephen Hawking's Centre for Theoretical Cosmology (COSMOS) to understand the origins and structure of the universe.

The new supercomputer combines HPE's Superdome Flex system with an HPE Apollo supercomputer and Intel Xeon Phi systems, to enable COSMOS to tackle cosmological theory with data from the known universe – and incorporate data from new sources, such as gravitational waves, the cosmic microwave background, and the distribution of stars and galaxies.

The in-memory computing capability of HPE Superdome Flex is uniquely suited to meet the needs of the COSMOS research group,” said Randy Meyer, vice president and general manager for synergy and mission critical servers at Hewlett Packard Enterprise. “The platform will enable the research team to analyse huge data sets and in real time. This means they will be able to find answers faster.”

“High performance computing has become the third pillar of research and we look forward to new developments across the mathematical sciences in areas as diverse as ocean modelling, medical imaging and the physics of soft matter,” said Professor Nigel Peake, head of the Cambridge department of Applied Mathematics and Theoretical Physics.

Beyond the use of Xeon Phi OCF’s, Fielden also noted several other potentially disruptive technologies that are now finding their way into mainstream use in HPC: ‘GPUs have already established themselves as being vital to obtaining a step change in performance, taking a lead in artificial intelligence and machine learning. One of the challenges HPC systems are going to face going forward is power consumption, and ARM is looking to secure a footprint in the HPC environment with its low-power processors which could prove disruptive.’

Fielden also noted two relatively new processing technologies that are also making potential waves in HPC over the next few years, as IBM and AMD both try to establish a foothold in the HPC processor market: ‘From the wilderness, AMD has come back to announce its EPYC processor, which could challenge Intel – which currently holds 90 per cent of the processor market. IBM is coming back with its Power9 processor, which was selected for two of the Coral research centre procurements in the US. Its OpenPower initiative has led to interesting collaborative arrangements with Nvidia and Mellanox.

Some people believe that the cloud will be disruptive, certainly AWS and Microsoft Azure are clearly taking market share. However, for more complex workloads and for people whom research is their business, it is unlikely that cloud will take the place of on premise because this would prove to be too costly. The hybrid cloud model will continue to evolve where organizations will have their own on premise systems and will burst out into the cloud when needed,” added Fielden.

## **Adding value**

While HPC integrators can certainly help users to select the right technology in the opinion of Fielden hardware is no longer the most important piece of a HPC system. ‘We’ve come so far and so fast in the past few years that services, support and consultancy have become equally, if not more important.’

The negotiation of the best deal with the various vendor partners is another vital part of an integrator’s added value,” Fielden added.

Much of the added value from integrators comes from support and maintenance which is

particularly attractive to academics or enterprise users that do not have large scale in-house expertise in this area. In these situations an integrator can help to manage and support the system to ensure that the system is kept operational and working as efficiently as possible.

Post-sales support is a key role in high-performance computing – whereby for every system supplied, customers have a support requirement and that will vary. Some customers need far more support than others, so integrator led project-development workshops add great value to work out what the best level of support for a customer is. It is, in the end, this ability to get close to, and, understand the customer, that is the secret of success for integrators,” said Fielden.

“The value of the integrator sits above the hardware and software – it’s the skills we have in designing the solution, pulling it together, and supporting it. We have a library of thousands of scripts written over the years that help the systems to run better,” Fielden concluded.

*This story appears here as part of a cross-publishing agreement with Scientific Computing World.*

*Sign up for our insideHPC Newsletter*