

Tech focus: cloud

Sponsored by



ROBERT ROE LOOKS AT USE CASES FOR CLOUD TECHNOLOGY IN HPC

Cloud technologies are now reaching a level of maturity that is making them appealing to HPC users. Whether using public or hybrid cloud, these technologies offer unprecedented flexibility for users who can create or 'spin-up' nodes with specific architectural requirements, use cloud bursting to increase the capacity of their in-house infrastructure – or it can increase the agility of a company that shares data over multiple sites.

In previous years there have been some concerns around security or the cost of moving data to and from the cloud, but these reservations are slowly being eroded as more users see value in developing a cloud infrastructure as part of their HPC resource.

One aspect of designing and procuring HPC systems in the past was the need to create a balanced architecture. This means looking at the kind of applications that will be run on a particular cluster to try and match the requirements of applications with the technologies that are needed. For example, some workloads require large memory nodes, high-speed storage or interconnects or high-performance storage.

Advancing the technology

In a perfect world, all of these technologies could be

included in a single system but, in reality, this is not feasible for most HPC centres as the cost of such a system would increase drastically. Cloud HPC allows people setting up this infrastructure to make more efficient decisions, particularly if they are cloud bursting or developing a hybrid cloud strategy – as they can build their in-house resources to cater for 80 per cent of the user requirements while using the cloud to provide GPUs or specific node architectures that suit a small number of users.

This allows all applications to benefit from this balanced architectural approach, while still being able to cater to the specialised applications that have more niche requirements.

Matching resources to requirements

Univa started off work in HPC as a company that designed scheduling software that helps match jobs to resources to increase the utilisation of HPC resources. By coupling this with cloud technology, the company aims to make using cloud and in-house HPC resources as economical and efficient as possible.

Rob Lalonde, VP for cloud at Univa, said: 'We started out as a scheduler company and we were working for about eight or nine years in that scheduler/workload management/cluster optimisation part of the world. We acquired Grid Engine from Oracle, which originally came from Sun Microsystems, and we have spent a lot of time commercialising this software, moving it from its open-source roots to the enterprise.'

'That is where we have spent our engineering investments over the last eight years. We have



Chanchai Howharn/Shutterstock.com

"That is where we have spent our engineering investments over the last eight years. We have now built up a base of around 250 enterprise customers and we sit on a base of more than 3.3 million cores"

now built up a base of around 250 enterprise customers and we sit on a base of more than 3.3 million cores on-premise and various customers based in Chicago, Toronto and Munich,' Lalonde added.

This combination of scheduling expertise combined with the growing use of cloud technologies led the company to develop software that could help to drive both activities together creating a more efficient platform to manage workloads both on-premise and in the cloud.

'Over the last few years we have seen a migration to the cloud, whereas three years ago we did a survey and didn't see that cloud activity in the customer base. A little over a year and a half ago we did another survey and we found that 61 per cent of our base was migrating to the cloud,' said Lalonde.

'Of those customers, 70 per cent were going hybrid cloud – meaning they want to extend their cluster into the cloud – and 30 per cent wanted to run on dedicated cloud clusters for all or part of their computing requirements.'

Lalonde gave an example of a large pharma customer that works with Univa to develop its cloud strategy. He explained that they are based on the west coast of the US, with their dataset in the cloud, but the datacentre is based on the east coast. 'It doesn't necessarily make sense to be pulling their data down into the on-premise datacentre in the east coast, so why not build a dedicated cloud cluster for that application which reduces latency, reduces data movement and bring it closer to the users?' commented Lalonde.

This is just one example, but it describes the kind of challenge

that many large companies face. This solution provides a special-purpose dedicated cloud that is typical of the kind of user that wants to build a cloud resource while still maintaining their on-premise cluster.

'There are lots of our customers that want to take the peaks off those dedicated clusters and manage that more efficiently, not necessarily build their cluster to match peak usage but some normalised usage,' said Lalonde. 'Then, when they have additional workloads, put that in the cloud so you can optimise your financial models so you are spending the least amount in the cloud while still getting your work done and achieving your results.'

'That is where a lot of our customers come at the cloud, saying "how do we leverage our on-premise investment and still go the cloud to get more done when we need it or to access specialised resources such as GPUs because we do not have that on-premise to do a special project",' Lalonde added.

Building momentum

Andy Dean, HPC business development manager at OCF, noted that in OCF's higher education and research customer base the use of cloud is gathering momentum. While the technology is not being used in every installation, Dean stressed that the use of cloud is increasingly coming up in conversations.

'It is an interesting time, The HPC community is relatively new to adopt public cloud, at least in our customer base, but it is definitely picking up. A large proportion of our customer base is in higher education and research. In a lot of the conversations that we are having, they are discussing how they can adopt these technologies.

One thing that Dean noted was that it is not security or other concerns that dominate these discussions, but talk around the pricing of cloud technologies. This is because the cloud allows HPC users to spin up specific architectures for each application, which means that many examples will be using

different hardware with different costs. This can make comparing the price of hybrid and public cloud installations quite difficult as there is little like for like comparison.

'It is as much around how you charge your users for these technologies or services, as much as, the technical questions around getting their research computing in the cloud. Developing the cost models as much as how do you run a job in the cloud,' added Dean.

Dean also said that while costs were continually coming down that was where a lot of the conversations were taking place. Customers have different options available to them so it is important they understand the cost and pro and cons of a public cloud service or developing a hybrid cloud strategy. Each can suit the needs of their users but, depending on the scale and requirements of the operation, the costs could be quite different.

'What the public cloud gives you is the ability to spin up the right infrastructure for the right workload,' added Dean. 'That makes figuring out the cost side of things quite complicated, because you are not really comparing like for like.

This approach allows datacentre operators to limit capital expenditure, as the GPU users or those who require a high-speed storage appliance can have their workloads run in the cloud while the main research cluster benefits from cost savings from not having to provide this specialised technology.

'That means for some are probably running their jobs on infrastructure that is too expensive, or others could have done with some faster storage or a faster processor, for example. That is where I think it is really interesting going forward,' added Dean.

Developing a scheduler

With the rising interest in cloud migration and hybrid cloud deployments in HPC, Univa saw the opportunity to integrate its existing expertise in scheduling

Boston

FEATURED PRODUCT



Feat. AMD EPYC 7002 Series

Introducing the Boston Quattro 12256-T02 featuring four hot-swappable dual-processor nodes, utilising the latest technologies available on x86.

Each node supports two 2nd Generation AMD EPYC processors, up to 4TB DDR4-3200MHz memory in 16 DIMM slots, up to six SATA drives and multiple add-on card options. The 4th generation of PCI-E provides double the bandwidth over the previous generation, required for some high-performance network controllers and next-generation GPU's.

The Quattro 12256-T02 featuring AMD EPYC™ 7002 Series, is designed with power and cost-effectiveness in mind, leveraging Titanium efficiency power supplies to reduce power consumption.

The flexibility and rich feature set of the Quattro 12256-T02 makes it a great solution for high performance and high-density computing. Contact Boston to arrange testing across a number of solutions and AMD EPYC™ 7002 Series CPU SKUs.

**www.boston.co.uk
01727 876100**

with cloud migration software that could help to manage the flow of resources where they are needed.

'Certainly there is lots of activity for us now in the cloud, and that's where Navops Launch comes in,' commented Lalonde. 'To oversimplify, we set a dial so an organisation can turn that dial and determine what goes to the cloud and what does not. In our world, you ideally abstract the complexity of the cloud from the end-users, those researchers and scientists that are submitting work.

The system works much in the same way as using a job scheduler – in this case, Univa's Grid Engine. Users submit work that is then pushed to existing infrastructure or to cloud resources, depending on the policy and tags associated with

a particular job. They submit work in the same way that they always have to the Grid Engine master on-premise and then Grid Engine master determines in conjunction with [Navops] Launch and the policy where the workload is actually going to run,' said Lalonde.

An example would be a workload that may be tagged with GPU, for example. If the organisation does not have GPUs on-site then Navops could push that workload out to the cloud, while Grid Engine manages the on-premise workloads as normal.

Another example would be a job that is tagged 'secure' so it will never be run outside of an organisation.

'Through policy decisions, you can determine how that cloud cluster interacts with the on-premise cluster. How it grows →

■ Oracle FEATURED PRODUCT



Accelerate HPC Workloads in the Cloud

The time for HPC in the cloud is now. Oracle offers on-demand HPC infrastructure, suitable for any HPC workload — based on the most advanced compute, storage, networking, and software technologies — at a fraction of the cost of building it yourself.

Oracle Cloud Infrastructure provides on-premises HPC performance with cloud flexibility and pricing. HPC applications run on bare-metal servers, no virtualization layer, no jitter, full network line rates.

The latest GPU and CPU chips with RDMA enable performance designed to meet or exceed your current HPC environment. Running HPC in the cloud means that there is no waiting in queues for your job, you can run and burst when you need to. Spin up thousands of cores in minutes and pay for them only when you use them.

Our bare-metal HPC enables true 'lift and shift' capabilities. See why our customers choose to run their HPC, AI, and ML workloads on Oracle Cloud Infrastructure: <https://cloud.oracle.com/iaas/hpc>

■ Univa FEATURED PRODUCT

MANAGE HPC CLOUD SPEND



Navops Launch - Enhanced platform for HPC cloud migration, cost management, and cloud automation

Univa® Navops is a multi-cloud provisioning, monitoring, and management platform that helps enterprises integrate HPC environments with the cloud, control spending, rightsize cloud resources and automate decisions based on real-time cloud, application and workload-related metrics with composable automation applets.

Unlike some tools that provision cloud resources based on static policies or a cloud queue, Navops Launch is application, resource, and budget-aware, and can

adjust workload and resource deployments on-the-fly and takes into account, planned spending by project and department, on-premise capacity, actual cloud resource usage by application, and data access patterns.

With Navops Launch organizations can maximize and balance the utilization of on-premise and cloud resources, reduce total spending, and boost performance of end-user productivity by ensuring that the right resources are available for the right application.

Learn more about Navops Launch and Univa HPC hybrid cloud solutions at <http://www.univa.com/products/navops.php>.

"You ideally abstract the complexity of the cloud from the end-users, those researchers and scientists that are submitting work"

→ and shrinks, what data goes where, and also what data gets flagged with those workloads because we can facilitate that as well,' said Lalonde.

'Because we are the Grid Engine people, Navops Launch and Grid Engine work together through a common repository of metrics that enables Navops Launch to know what's going on in the on-premise cluster,' said Lalonde.

'There are lots of ways of lifting and shifting workloads to the cloud. Azure and Amazon have got their tools but you are building dedicated clusters. We are leveraging that on-premise

investment so that is really attractive to our customers. We are not asking them to change their world we are helping them to extend their existing infrastructure and that makes customers more comfortable,' concluded Lalonde.

Making the most of hybrid cloud

With the opportunity to push specialised workloads out into the cloud also comes a challenge in ensuring that the application fits the architecture. In order to maximise the efficiency of a public cloud set-up, it is imperative that 'you are

going to have to understand your applications a lot better to really benefit the cloud,' said Dean. 'You can spin up a LUSTRE environment and nodes with a high-performance interconnect but you really need to understand if that is what your applications needs.'

Dean acknowledged that, while some customers may have just one or two applications, that is not the case for the majority of research clusters. In a scenario like this, it is a simple case of working with the ISV or benchmarking the applications to determine which architecture suits them best.

However, most HPC centres have hundreds of users with nearly as many different applications. For these types of customers, Dean notes that it is

much more difficult to ascertain which hardware is best suited to the majority of users. 'Some customers have hundreds of applications, and then there has to be a little bit of picking the low-hanging fruit and understanding where the biggest benefits can be made by truly understanding the applications and figuring out what should be placed where,' said Dean.

'The great thing about public cloud for testing these technologies is that you can access them and only pay for what you use. You do not need to invest £100,000 on infrastructure just to see if your application scales and performs as you need it to. You can spin up 1,000 of these nodes if you want to and you are only paying for the time you use,' Dean concluded. ■