The Red Oak Microsoft Azure HPC Collaboration Centre in partnership with Intel

From requirement to proof of concept, preparing for HPC and research computing in the Azure Cloud.

Dairsie Latimer

04 March 2021

Red Oak Consulting
Expert advice, exceptional delivery

Contents

Introduction

Red Oak Consulting's core business is high performance computing (HPC) consulting. With a reputation for expert knowledge combined with real-world experience, we've delivered strategy, given procurement advice and helped to install more than 30 on-premises HPC systems (including several in the Top 10 of the Top500). To date, we have also completed more than 20 cloud-based projects, bringing first-class technical skills to working with 'early adopter' customers in both higher education and commercial sectors.

Along with our partners from Microsoft and Intel, we want to share the principles of best practice and advice in a series of white papers on how to use the Azure Cloud for HPC and research computing. This work was funded by Microsoft in partnership with Intel.

This paper is principally aimed at higher education but if high performance computing is a key element of the support provided to scientific staff and researchers by your organisation then this paper will help in understanding the relevance of HPC in the Cloud.

The ability to scale resources in response to demand as well as allowing users (and the organisation) a painless and continual upgrade path to new technology while simultaneously outsourcing the purchasing, maintenance and housing of the equipment is highly attractive to many. If you see high-performance computing as a pivotal productivity tool for your researchers but have no intrinsic desire to manage your own physical equipment then Cloud will already be on your radar.

In this first white paper, of a series of three, we will consider the value derived from your HPC and research computing service, how the Azure Cloud is uniquely placed to help, and how to get the evaluation process and the project prerequisites right. We want to make the path to evaluating and adopting Azure Cloud for your HPC and research computing requirements as smooth and cost-effective as possible.

The next whitepaper will look at running a successful proof of concept evaluation, including standing up an Azure CycleCloud environment, building and benchmarking a range of popular codes using the Intel Software ecosystem and VMs under Azure.

How business requirements drive an HPC in the cloud strategy

Many organisations have already decided on a cloud first strategy when it comes to deploying new or replacing existing services. For enterprise applications and use cases this is now neither controversial nor unusual. However, for HPC, there are tangible reasons why this may not yet be quite such a straightforward choice, not least the perception of the cost of using cloud for HPC.

HPC is typically operated at a scale that means an organisation must also make significant outlays in capital costs for datacentres and associated facilities, alongside the cost of the HPC systems themselves. As many enterprise uses of datacentres have migrated to the Cloud, the need to maintain such large and capital-intensive infrastructure has decreased. This has made the proportion of cost of ownership that needs to be attributed to an on-premises HPC system climb in relative terms.

Even if you have previously had competitive on-premises HPC, as infrastructure ages and nears the end of its lifecycle, the merits of continuing to invest in one's own datacentre, versus taking advantage of the economies of scale, skills and capabilities of datacentre specialists, such as the hyperscalers, becomes an important element of understanding the whole life costs for HPC and research computing.

For many sectors, especially small and medium enterprises (SME) and the higher education sector, the capital costs necessary to maintain a modern and efficient datacentre capability can be significant compared to the actual amount of HPC capacity needed. The most common situation is that HPC capacity is constrained by one or more elements, including capital costs, space, poor understanding of return on investment for capital projects including HPC, and unpredictable future requirements.

This constraint based operating model, can lead to unworkably high levels of utilisation for HPC systems, with commensurate queue waiting times and occasionally the inability to effectively run certain workloads. This, as we will explore, can lead to lowered productivity, and is a constraint on at least the ambition if not the quality of the work undertaken.

In other circumstances, HPC capacity has a burst requirement, which means that the utilisation of the system is relatively low for periods of time, leading to a lower return on the investment. The reality is often a blend of both these extremes. Figure 1 shows a typical demand vs utilisation plot for an HPC system, running at lower than capacity for a period after installation and then running at a capped peak for much of the remainder of its lifetime. This leads to an effective utilisation over the lifetime of a service of less than 100% (which equates to increased costs) while not actually meeting growth in demand over time (lost value). Cloud offers a new way to deal with the feast or famine environment that a HPC service manager is often faced with over the lifetime of a HPC system.

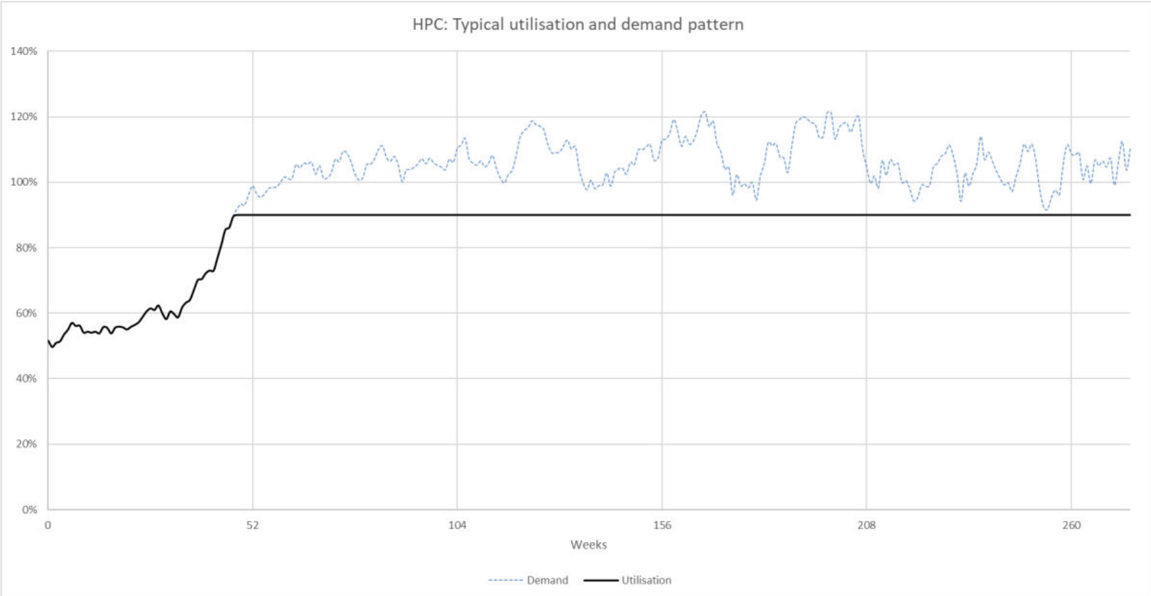

Figure 1: Representative demand vs utilisation plot for an HPC system over a five year period.

Regardless of which situation you find yourself in, the most common scenarios for cloud utilisation for HPC and research computing are to:

- cloud burst excess jobs (based on business rules);
- migrate partial workload(s) to the Cloud (selected for time to solution and cost to solution);
- migrate specific end-to-end workloads;
- migrate entirety of production workloads to the Cloud.

How to understand value of HPC in the Cloud

We have touched briefly on the need to understand the total cost of ownership but arguably more important is understanding what the value of HPC is to your organisation. Understanding how it drives value, is key to selecting the right strategy to provision it. Part of understanding value is determining which metrics to measure it by. With some effort it is usually possible to find a metric that encapsulates the value of HPC to an organisation.

Sometimes it is simply seen as a binary outcome; with HPC I can work, without it I cannot. But that misses an opportunity to look at a range of other metrics, such as productivity, return on investment and other value indicators, which are almost all affected by the capability and throughput of an HPC system.

For example, are there business opportunities left unaddressed because they cannot be accommodated by the current capabilities or capacity of the current HPC service? What is the opportunity cost? Does the organisation track the right metrics to calculate the return on an investment in HPC? Not all these questions must be fully answerable before you start making decisions and there are of course limits to the usefulness of data in making some decisions, especially longer-term strategic ones. We can become blinkered by the cost of something without fully exploring the value of something, especially if it enables a different way of working which adds value in a way that is not currently possible.

In our experience HPC has many of these inflection points and the flexibility of Cloud to rapidly adapt and react to these (at least in the context of a typical on-prem HPC system lifetime) can provide some of the strongest strategic reasons to consider Cloud as part of your long-term strategy.

Stakeholder engagement and management

Establishing a clearer understanding of the value of HPC to an organisationrepresents a significant effort during the early stages of a project. The best way to achieve this is to gather a representative set of stakeholders. This will include users, managers, IT directors, systems administrators and executive roles and functions. All these stakeholders will have valuable insight into the value of HPC and research computing to an organisation as well as governance requirements and advice.

This process is hard to do and even harder to get right, but if you can identify the key users and influencers early, making them part of the requirement capture, solution architecture, acceptance and signoff processes, you are more than halfway to getting your strategy right.

The total cost of ownership (TCO) model

Any attempt to compare between different methods of provisioning a research computing/HPC service needs to have a robust whole life total cost of ownership model. HPC services are usually costed in $'s per core hour delivered, especially if the organisation has a need to charge back costs to projects or users.

A word of warning, $'s per core hour is really only appropriate for comparing similar technologies and business models. It is a poor metric when comparing significantly different technologies (eg CPU vs GPU) and differing business models (fixed vs flexible resources). It also does not capture many broader value elements such as enhanced productivity and flexibility, or risk mitigation and business continuity. So as a result not all TCO models are created equal and any comparison is only as accurate as the model you use for costs and how you represent value from new ways of working.

In an academic setting, where the research computing and HPC service is free at the point of use, this cost is often not exposed to the user (though sometimes available in a 'full economic costing' figure used for funding bodies), but is still vital to understand. Also, if the only attempt to metricate value is based on the cost per core hour, then arguably the strategic value of HPC to your organisation is less well understood than it should be.

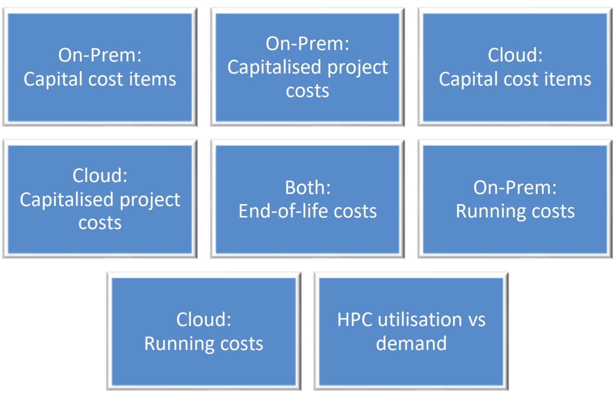

Figure 2: Main cost of ownership categories

In the past it was a reasonable question to ask if a cloud environment could run a particular workload, but today with modern HPC tuned VMs and infrastructure, it is no longer a question of if your jobs will run in the Cloud, but if they will run cost efficiently. To understand cost, it is necessary to have a cost of ownership model to compare against.

Total cost of ownership modelling is complex for any environment. Having a TCO, or better still a whole life cost model, allows important questions about how to manage the delicate balance between costs and investment in HPC against cashflow and investment in the whole organisation.

Prerequisites and readiness

While the list of prerequisites necessary to conducting a proof-of-concept evaluation isn't too daunting (essentially you need an Azure subscription, some credits), when looking at an extended pilot or production migration there are more things to consider. A non-exhaustive cloud readiness list includes:

Cloud migration strategy

- Has the organisation developed readiness and vision to drive a business change?
- Are all stakeholders aligned with the need for change strategy?
- What are the acceptance criteria for successful cloud migration?
- Is there a funding model or forward strategy supporting a cloud migration?

Connectivity and access

- What type of network connection to Azure or the internet is used?
- Are the internet and intranet network bandwidth adequate?
- Is there an existing Azure AD configuration and is there a user management interface?

Policies and procedures

- What governance structures are in place for Cloud?
- Are the necessary security requirements in place?
- Is there a current backup policy and how will it work with the Cloud?
- Is there a disaster recovery policy and can it be adapted with the Cloud?
- What storage lifecycle management policies exist and are required?
- How will Azure Cloud resources be paid for? Is there a charge back model?

Workloads and workflows

- Are all types of HPC workflows known or does an audit need to be undertaken?
- Are all characteristics of workflows known (eg compute, memory, fabric and storage)?
- Are there workload baseline benchmarks for comparison?
- Are there any specific license requirements for workloads in the Cloud?
- What data types (eg privately identifiable information) and storage are required?

Staffing and experience

- Can the staffing hierarchy for HPC support be outlined?
- Is there operations and support capacity for migration to the Cloud?
- Does current HPC support staff have any experience with Azure services?
- Is there a process for staff training and personal development?

Current on-premises estate

- What is the current hardware infrastructure estate and its condition?
- What are the power and cooling requirements, including current condition?
- What is the timeline for maintenance and support contract renewals?
- When is the planned date for hardware renewals and infrastructure refresh?
- What is the status of the data centre building and it end asset life?

Requirement capture process

Assuming that you have addressed most of the prerequisites and you understand where the value from HPC is derived from, you can start to gather and prioritise more detailed requirements for the HPC/research computing system and service. This is probably best done using a multi-phase approach, conducting a limited proof of concept phase to deliver performance baselines and rough costings which will qualify Cloud as a competitive or not, followed by a more comprehensive pilot study looking at the integration of further services and only then moving to a production migration planning stage when all relevant prerequisites have been addressed.

The success of any project is in the preparation and understanding what the goals and acceptance criteria are. Once you have the project governance in place you can start to collate requirements along the lines shown in Figure 3. As part of the process, consider creating user stories and then mapping current HPC workflows to the architecture.

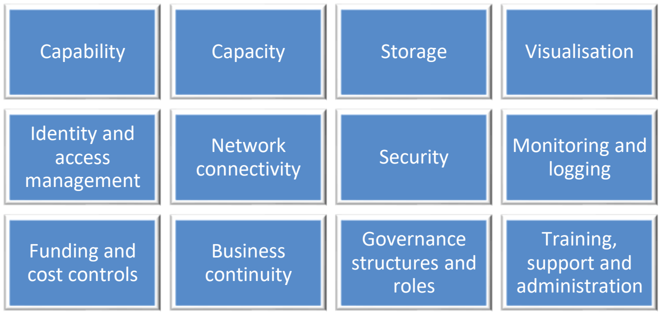

Figure 3: Major areas of consideration for a requirement capture process

When developing cloud systems, it's usually best to use a CI/CD methodology along with more traditional project management techniques for overall project control. Making sure you have the right stakeholders on your project board and ensure senior users and administrators are fully engaged, from the proof-of-concept stages all the way through to production is also important.

Examining the value of cloud HPC

Regardless of how efficient and cost effective your current HPC solution is, there are going to be occasions and use cases where a cloud capability can materially deliver value.

Three words to consider here are: *Agility*. *Flexibility*. *Scalability*.

All of these to a greater or lesser extent are impacted by the revolving door experience of conventional procurement cycles for HPC. It is not that conventional services cannot be these things, but it is by some measure harder to make them so.

In practice, if you can make a case for spending money on a project, then with a cloud capability you do not have to wait months for procurement to move through the gears. This is especially useful for time is of the essence opportunities. Yes, you could make this the highest priority job on your current system but in doing so it likely to impact other projects. What if you could simply burst certain jobs to the Cloud and not impact the rest of your user community? What value could that add?

Agility

Timescales for HPC procurements are still measured in months or even years for larger and national infrastructure scale services. However, once you have the right plumbing in place for the Cloud, it is not impossible to add new services or scale your use of existing ones in the Cloud in literally hours. This is in marked contrast to many experiences with procurement processes that are decidedly not agile.

Additionally, adding new services and testing them in a siloed environment is relatively straightforward using a CI/CD methodology and the ability to deploy multiple virtual clusters. In fact, it is not just virtual clusters, as containerisation, Software as a Service (SaaS) and Platform as a Service (PaaS) open up new ways of doing things and actively encourage experimentation.

Flexibility

Cloud offers the opportunity to be more flexible in many areas of HPC and research computing. This is especially true when it comes to capacity forecasts. No longer do you have to buy capacity on day one to last you three to five years (see Figure 1).

Additional flexibility and agility in terms of always being able to run on the most appropriate infrastructure and adopt new technology as it becomes available, seeing how it can be used in your workflows without having to buy it up front (try before you buy). Azure offers support for modern HPC infrastructure with low latency InfiniBand fabrics, SmartNICs, generous per core memory and support of several high-performance file systems including Lustre, BeeGFS and Azure NetApp Files.

No longer does one size have to fit all. You can provision different systems for different users and tailor the characteristics based on their need. In the higher education context, job mix is often very variable, with a tendency towards a high proportion of single threaded or sub-node jobs combined with a mix of small to medium sized parallel jobs and the occasional massively parallel workload. The Azure Cloud can cost effectively accommodate this heterogeneous workload mix without having to exclusively use HPC optimised VMs. Additionally, if you don't want to sink money into racks of GPUs and high memory nodes, which are necessary for a few users but are on average quite underutilised – Cloud can be a great fit.

Further, for many types of requirement, SaaS makes the job of administration easier by providing robust services that no longer must be directly under the care of the IT team, freeing them to spend more time with users and evaluating new technologies.

Scalability

Scalability in cloud computing refers to the ability to quickly increase or decrease the allocated resources as needed to meet a changing demand. In a traditional on-premises environment capacity is a fixed resource, meaning that depending on demand there are either queue waiting times or machines that are unused. The cost per core hour is also fixed for an on-premises system, but in a typical cloud environment there are multiple ways of charging for the available core hours based on the priority of the workloads. There are limits to the size of the scale sets, and are most optimal in the 100 VM range (but can scale to 1000 VMs), meaning a single job running on Azure's Intel Xeon based HC44 VMs can easily scale to over 4000 cores. Depending on capacity you can of course run many of these jobs at a time, enabling work at a scale perhaps not possible before.

Hybrid HPC – First steps towards HPC in the Cloud

Depending on how its procured and run, the cost of ownership for on-premises HPC can vary significantly. For large users of HPC who treat HPC as a strategic asset and maintain centralised, modern data centre facilities, then it is true to say that HPC in the Cloud is unlikely to be significantly cheaper per core hour than a best in class on-premises system. To match a competitive on-premises total cost of ownership, a cloud HPC system must utilise Spot and reserved instance (RI) VMs. This does require a change in mindset and sometimes how workflows are structured.

However, there will always be workloads that will be amenable to being run in the Cloud, either because they require capabilities not available on-premises (perhaps GPU or high memory nodes, perhaps even more modern network fabrics) or because the opportunity cost of not running them (due to lack of capacity on-premises) makes utilisation of cloud capacity more economically viable.

One must also look further into the future, where the cost of cloud cycles drops over time and take a more strategic view of Cloud for HPC and research computing. Should you simply wait another few years before you start evaluating Cloud for HPC or is it worth engaging now, such that you can utilise a hybrid cloud HPC model when appropriate?

A hybrid approach is still going to require effort to evaluate and deploy but it should be viewed as a forward-looking investment. It is a journey not a destination. Cloud HPC will enable innovation driven by agility, flexibility and scale.

Top 10 tips for engaging with cloud HPC

There are many best practice suggestions when considering an evaluation or migration to the Cloud but our top ten pieces of advice are:

1. Actively engage with stakeholders across the board, this is especially important with IT and the user community;
2. Train your system architects and administrative staff (and later users) as early as possible;
3. Conduct a cloud readiness review which sets out any organisational and IT changes necessary for a successful migration; [See Prerequisites and readiness section]
4. Determine what governance framework for Cloud is necessary for your organisation. This should at a minimum include finance, operational IT and support governance and roles;
5. Don't formulate your cloud adoption strategy too early, as it will more than likely be nuanced. Wait until proof of concept and pilot projects have delivered indicative time to solution and cost per core hour metrics per chosen workflows, and the business requirement phase has put those costs in context;
6. Ensure that you have a comprehensive cost of ownership model so that your cost to provision HPC can be compared on an equal basis and don't forget to also account for additional value around the flexibility cloud can deliver as well as possible risk reduction and mitigation;
7. Treat your chosen cloud provider and integrators as partners, they have a wealth of experience you can leverage for success;
8. Take the opportunity to reduce technical debt as part of any migration. Adopt a CI/CD mentality and constantly review the state of the art in SaaS, PaaS and IaaS offered by the provider. This is quite unlike a conventional HPC procurement where much of this effort is front loaded;
9. When developing your Infrastructure as Code (IaC) automate as much as possible including testing where possible;
10. Ensure that you monitor and manage spend, and that all necessary checks and balances are built into your deployed framework;

Next steps (proof of concept studies and beyond)

As we have already said, we believe it's no longer a case of if the Azure Cloud can run most HPC workloads, but can it do so cost effectively compared to your current on-premises infrastructure.

We have spent some time outlining the process and procedural steps an organization should take when engaging in a cloud HPC evaluation. The next whitepaper will look at running a successful proof of concept evaluation, including standing up an Azure CycleCloud environment, building and benchmarking a range of popular codes using the Intel Software ecosystem and VMs under Azure.

It is worth noting that for a proof of concept, the Azure CycleCloud deployment does not need to be fully integrated into all your other enterprise IT and cloud infrastructure. Such levels of integration would usually be addressed as part of an extended pilot project, which itself would feed into the requirements for, and the development of, a production quality deployment.

If after the proof of concept phase the decision is that cloud does not yet unlock specific value, consider when will it and for what workflows it will be best suited? Use time wisely and prepare and don't wait until you are starting your next procurement cycle.

By considering how you can engage early and extract value from doing so, you can make your organisation ready to adopt Cloud when it makes sense. In particular you can identify exception workloads, and how the elasticity and flexibility of the Azure Cloud can help you today.