

This briefing paper is aimed at researchers who are considering how cloud computing might help their research. It is supported by a detailed report which examines the issues in detail.

'Cloud' as a term is being applied to a broad range of computing activities, some of which are novel, and some (eg webmail) which are already well-established.

Cloud Computing variously combines considerations of hardware, software, services, programming environments, resource management, billing services, and legal, regulatory and ethical governance. Various definitions and characteristics of Cloud Computing have been put forward; the NIST defines Cloud

Computing as: "a model for enabling convenient, on-demand network access to a shared pool of configurable computing resources (eg, networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction..."

This briefing paper does not contain a technical analysis of cloud systems (see the JISC Technical Review of Cloud Computing), but rather considers the overall picture of cloud computing for research, focusing on infrastructure clouds which provide compute and storage, rather than 'software as a service' or other approaches.

There is currently an extremely rapid pace of change and innovation across the range of activities considered to be Cloud Computing. This covers not just the details of costs of services from specific providers, but also the range of services offered and the available tools for configuring and controlling cloud resources. Entirely new ideas, technologies, business models, providers, strategies and risks continue to appear.

There is a wide range of cloud vendors, although three make up the bulk of the market: Amazon Web Services (AWS), Google App Engine, and Microsoft Windows Azure.

There is a broadening range of academic research into the use of Cloud Computing for research, although as a proportion of research computing overall it is insignificant. At present, the vast majority of research computing using the cloud is undertaken on AWS, funded by Amazon.

The place of cloud systems in research

We believe that in the short to medium term, cloud computing will not significantly replace the procurement, provision and use of departmental, institutional and national HPC resources. The current cloud infrastructure technologies and business models do not provide the capability which current users of these facilities demand.

Some institutional systems will likely become virtualised using cloud technologies, but we do not believe that small-scale private clouds will provide the key benefits that large, commercial clouds do. Cloud Computing is made efficient and effective by its very large scale. This scale permits the flexible and elastic nature of the services that are provided.

However, there are a range of cases where cloud provision may provide important new capabilities, including:

Short duration requirements: specific research needs to be undertaken that has a short-duration requirement for resources. For example, it would take longer to purchase and deploy infrastructure internally, negotiate for and learn to use an internally or externally provided cluster, or wait for a queue on a specific shared cluster, than it would to undertake the research.

Infrequent use: users may only require occasional use to servers or cluster systems, potentially wanting to package up what is running on their own machines and run it elsewhere. Alternatively, the overheads of maintaining infrastructure with a relatively low level of utilisation are too high.

Cloudbursting: additional compute capacity may be required on demand at specific times, to cope with unpredictable peaks of research computing, than is available within the system which the researcher is able to access.

Transfer to commercial use: a system may be required to be accessible, usable and manageable by, or directly transferrable to, a commercial partner, for example if a the researcher is building and testing applications in the Cloud that will also be usable and modifiable by a third party. Neither organisation needs to develop or maintain the infrastructure.

Flexibility with system: users may want to be able to install and run their own applications without negotiating with system administrators; alternatively, users want to offset the potential up-stack impacts of system upgrades and patches, downtime, and related system factors.

Data hosting: research organisations may wish to make use of the levels of redundancy provided by Cloud storage, and/or retain off-site backups. Alternatively, individual researchers, research groups, or wider research communities wish to have data available "where the computing power is".

Cloud-based publications: to support the repeatability of science, researchers may publish papers that contain a reference to virtual machines and storage that comprise the data, applications, and other essential elements of the published experiment.

Ad hoc activities in support of research: researchers occasionally need to make use of web-based software in order to undertake efforts that either cannot be supported readily within the organisation due to policy limitations or manpower availability. An example is a limit being placed on

the size of files that can be emailed to project partners, or wanting to host websites outside the confines and templates of an institutional content management system.

Paying for cloud resources

One of the defining features of cloud services is that access is on-demand, and charging is correspondingly pay-as-you-go. The costs of using the resource are allocated in a range of ways, most commonly charging per machine-hour and for data transfer.

At first glance, the costs of cloud computing frequently appear high to potential users from academia. But the core issue is that the costs of research computing are often hidden from the users of the service. Institutions typically provide some or all of the facility as an overhead, paid for from institutional or departmental budgets. These overheads include buildings, power, cooling, staff and administration. This creates an environment where very few researchers have a clear idea of the true costs of the resources they are consuming.

Cloud Computing provides an environment where many more of the costs are exposed to the users of the service, and so the service is more expensive *to the users*. Unless an organisation has fully identified the costs of local provision, and understood how those costs would be changed by adoption of cloud systems, it is not possible to judge which approach is less expensive *overall*.

Security and information assurance

There are a wide range of concerns about the information assurance aspects of cloud provision. Many of the issues that are raised are not well addressed within current local or institutional provision, but there is often a sense (rightly or wrongly) that local services are more trustworthy than external ones.

A common question is “who can access my data on the cloud?” This question has little meaning – “the cloud” is not an entity as such, and it is up to the user to decide which services to use, and each service has a distinct approach to data. Generally, services bought from cloud vendors are private services – your data (and other resources) belong to you and will not be accessible to others unless you make it so. From a security (confidentiality) aspect, many potential users expect absolute guarantees that their data cannot be accessed without their authorisation, but it is never possible to give these guarantees. New vulnerabilities are constantly discovered in all elements of the internet, and until they are disclosed, they will be exploitable.

Regarding the availability and integrity of data (commonly referred to as “backup”, although the issues are broader), cloud provision provides just one more option in the range of services available to institutions or their researchers. Approaches must be considered on a case-by-case basis.

Contracts and liability

The contracts and Service Level Agreements currently on offer are one-sided in favour of the service supplier and against the client. The style is very much lawyer-speak, there is a risk that HEIs, or non-specialist staff within them, will sign without realising all the implications.

One particular issue to note is that these contracts are established under US, rather than any UK law, and require conflicts to be settled in US courts. Our analysis highlights some areas where the contracts may be illegal under UK law, but may be permissible under their own jurisdiction.

Whether the contracts and SLAs offered are appropriate for any given research task will need to be considered on a case-by-case basis. It is particularly important that users consider the obligations that they have under the Data Protection Act and the Freedom of Information Act. Vendors are unlikely to negotiate changes for individual clients, but opportunities to enable more business with the UK HE sector as a whole may prove enticing.

Final comments

Staff skills are unlikely to be significantly affected by any shift to Cloud Computing. The individuals who are involved in research computing at present are technically-minded, and cloud presents just one more technological opportunity. Cloud Computing is sometimes held up as a panacea which can allow non-technical users (and non-traditional research computing users from the humanities) to easily access powerful facilities. This is not the case, for infrastructure clouds at least, where significant technical expertise is required to provision and administer these systems.

It is clear that institutions should consider their approach to Cloud Computing – is it acceptable for individual researchers to use these services without institutional consideration of the risks? Many of the issues are similar to those around individual researchers buying and configuring their own local resources; who is responsible, and whose reputation is damaged if there is an incident?

Further Information

This project's website, including the full report

www.jisc.ac.uk/whatwedo/programmes/researchinfrastructure/usingcloudcomp.aspx

Using Cloud for Research: A Technical Review

www.jisc.ac.uk/whatwedo/programmes/researchinfrastructure/cloudcomptechreview.aspx

Review of the environmental & organisational implications of cloud computing in HE & FE

www.jisc.ac.uk/whatwedo/programmes/greeningict/environmentalreviewcloudcomp.aspx

Cloud Computing Risk Assessment

www.enisa.europa.eu/act/rm/files/deliverables/cloud-computing-risk-assessment

