
Cloud Computing and the Power to Choose

By Rob Bristow, Ted Dodds, Richard Northam, and Leo Plugge

Some of the most significant changes in information technology are those that have given the individual user greater power to choose. The first of these changes was the development of the personal computer. The PC liberated the individual user from the limitations of the mainframe and minicomputers and from the rules and regulations of centralized system management. Individual users could install PCs on college and university campuses without obtaining the approval of the central IT organization. For the first time, individuals could choose their own word processor, programming language, utilities, and more. As the number of personal applications grew, more attention to user-friendliness became a key selling point.



Subsequent waves of change—the World Wide Web, mobile devices, Web 2.0 functionality, and virtualization—continued to benefit users as interlinked PCs and other personal devices delivered information and innovative new applications to individuals connected via global networks. In addition, the resulting industrialization of information technology allowed generic services used by individuals—services such as e-mail, chat, and the exchange of documents, pictures, and videos—to be provided on a massive scale. Today called *cloud services*, many of these—such as Wikipedia, Hotmail, and YouTube—were available before the phrase *cloud service* was even coined. What is new is the delivery of these

From the individual user's point of view, the only thing that matters is the power of choice.



services on an industrial scale through what might be called IT “factories.”¹ Yet the methods and techniques for delivering these services do not concern the user. From the individual user's point of view, the only thing that matters is the power of choice: the opportunity to choose among cheap (or free) competing services that are user-friendly, accessible from any location, and within higher education, potentially more reliable than campus services.

Colleges and universities around the world are thus discussing, planning for, and using cloud computing and cloud services. The rate of adoption varies from country to country, but the need for awareness and preparation is universal. This article will examine cloud issues—both opportunities and risks—by looking at examples from four countries: Australia, Canada, the Netherlands, and the United Kingdom.

Drivers

What are the driving forces behind the current interest in cloud computing in higher education?

- College and university IT departments are struggling to deliver scalable, secure, reliable, and cost-effective technology services in a time of shrinking budgets and growing demands for increased operational efficiencies.
- Students, faculty, and staff are bringing a rapidly changing array of consumer electronic devices to campus, and they expect ready access to easy-to-use mobile applications.
- Large commercial IT organizations are gaining significant economies of scale—in their infrastructure and service-delivery capacities—that individual college and university IT departments simply cannot match.
- College and university CIOs are becoming intrigued by the possibility that some—and perhaps most—of the services currently managed by the central IT department could be moved to the cloud. However, many CIOs lack direct experience using cloud services, fear the prospect of a new form of lock-in by service providers, and have deep and legitimate concerns about information privacy issues.
- Colleges and universities are looking to their CIOs and their IT staff for strategic advice associated with the institutional mission of teaching and learning and research, yet at the same time, CIOs are often distracted by operational issues associated with running core infrastructure.
- Many core IT services are being viewed as a utility service much like electricity or gas. As commercial providers enhance their capability and maturity in providing these core services, moving to a commercial provider becomes a more attractive option.

Examples

Some aspects of cloud services have become embedded in individual practice. For example, many students forward their college/university e-mail to their private Hotmail or Gmail account, which provides

much more storage space. They use their institutional e-mail account as a kind of stationery to show addressees that they are a member of a trustworthy institution—for instance, when they apply for an internship. Other examples include using cloud services to communicate with friends and colleagues (e.g., Facebook, Twitter), to collaborate on documents (e.g., Google Docs), to back up files online (e.g., Dropbox), and to exchange documents, video, or photos without the intervention of system managers. Researchers are also starting to leverage the elasticity that the cloud (or cloud computing) offers by using systems such as Eucalyptus and e-Science Central.

E-mail

Currently some 75 percent of colleges and universities in Australia and New Zealand have moved their student e-mail to the cloud.² Early in 2010, Macquarie University became the first higher education institution in Australia to also move its research, teaching, and administrative staff e-mail to a cloud provider.

Student e-mail provision has been adopted in a number of U.K. institutions, with more examining this option. One university, having already placed student e-mail in the cloud, is looking to embrace cloud solutions for all its IT requirements where appropriate. In Scotland, another university is going through the tendering process for placing its faculty and staff e-mail in the cloud, in addition to its student e-mail service.

In the Netherlands, the most popular scenario currently is to outsource student e-mail. Anything else is considered too risky due to concerns about privacy, ownership of the data, security, continuity, and so on. However, student e-mail is only one part of an institution's e-mail, reducing the economic benefits of outsourcing.

In Canada, the move to the cloud has progressed rather slowly due in large measure to strict legislation governing the storage of personal information outside the country. The legislative restrictions apply to U.S.-hosted cloud solutions that are institutionally *provided* even if they

are not institutionally *mandated*, but the legislation does not inhibit individuals from choosing to use the services. Indeed, most students and many employees make personal use of these services every day. The privacy statutes vary from province to province, making it difficult for would-be domestic cloud service providers to aggregate national demand under a standardized contract template. Nevertheless, higher education institutions are currently consulting with students and other community partners in an effort to find an acceptable solution that will lead the way to cloud services, beginning with student e-mail. Wilfred Laurier University plans to move student e-mail, calendar, and document storage to a cloud solution this summer. Dalhousie University is in the process of moving alumni e-mail services (for life) to Google; its next step is to explore migrating student e-mail to the cloud. Both institutions are motivated by a strategic direction to focus internal efforts on technologies that address core academic activities while beginning to disengage from nondifferentiating services that can be offered more efficiently elsewhere.

The University of Alberta is taking a strategic approach to moving e-mail services for faculty, staff, and students to Gmail. Institutional leaders have built consensus around a goal to adopt Gmail but have made adoption contingent on successfully completing a comprehensive Privacy Impact Assessment (PIA) and then negotiating a contract that is fully compliant with provincial legislation. The PIA is complete and has been accepted by the government agency responsible for information and privacy. Contract negotiations are ongoing. The PIA document has been shared with members of the

Canadian University Council of CIOs (CUCCIO), and assuming a successful conclusion to negotiations, the resulting contract will also be shared with CUCCIO members. These documents can then form the basis for standard templates that can be used across the country. This should result in faster and simpler adoption of some cloud services and serve as the foundation for negotiation with other providers.

Research

Another opportunity provided by cloud services is to support researchers in reducing the costs involved with computation. Only a small number of researchers need *capability computing*—that is, high-performance computing (HPC) systems with large numbers of cores. The majority of researchers are well served with *capacity computing*—that is, systems that share their computing power with several and up to many users. This capacity computing is exactly where cloud computing excels. Recently Microsoft and the National Science Foundation announced an agreement that will offer selected individual researchers and research groups free access to the Windows Azure cloud computing resources.³ This initiative opens up a whole new spectrum of opportunities for both researchers and institutions.

One specific use of the cloud for research is e-Science Central (<http://www.esciencecentral.co.uk/>) in the United Kingdom, developed by Professor Paul Watson and his team at Newcastle University. Building in part on experience gained through projects funded by the Joint Information Systems Committee (JISC)—projects such as my Experiment (<http://www.jisc.ac.uk/whatwedo/programmes/vre2/myexperiment.aspx>)—e-Science Central delivers cloud-enabled e-science capability to researchers across many disciplines. Researchers build workflows in a drag-and-drop interface from preexisting shared modules, or they write and contribute

their own workflows using Java, .net, or Python. The system utilizes a collaborative model inspired by Facebook, with users having the ability to form groups and share data and processes in fine-grained ways. The security model of e-Science Central works with this element of the system, allowing researchers to permit collaborators to see their workflows and to comment on them. The system also has a built-in blogging tool, enabling the development of conversations. Once an experimenter has set up the workflow, he or she initiates the process, at which time the system sends the computational work either to Newcastle's own servers or to the cloud, as appropriate. Although e-Science Central currently uses Microsoft's Azure platform, the system can also be configured to use Amazon EC2. It has about fifty regular users across a range of scientific disciplines, but Watson sees the system as scalable to thousands in due course. He also sees the potential for applications like e-Science Central to open up e-science to what he terms the "long tail" of researchers—those who at present are excluded, by geography or cost, from access to grid computing or HPC.

The Netherlands is currently building a national IT infrastructure for science and research as a cloud service. This infrastructure includes a virtual laboratory for e-science with generic functionalities to support a wide class of specific e-science application environments. The technology and resources come from different sources, such as SARA, BiG Grid, GigaPort3 and SURFnet, DAS-3, and Starplane. The software is developed in and by research domains such as high-energy physics, food informatics, medical diagnosing and imaging, biodiversity, and Dutch telescience.⁴

In Canada, although there has been some spot activity in the use of cloud services for research—for example, NEPTUNE Canada (<http://www.neptune.canada.ca/>) used Akamai video services to broadcast the launch of its undersea network—the main research focus continues to be grid technologies.

Capacity
computing
is exactly
where
cloud
computing
excels.



Challenges

Uncertain Definitions

The term *cloud computing* means different things to different IT professionals and to different institutions. It is often used to loosely describe a broad range of activities, ranging from outsourcing a specific activity to a single external provider (which many would argue is *not* cloud computing) to delivering a set of services from the cloud in such a way that users are not even sure where their data is being housed or where it is being processed.

Preliminary findings of three JISC studies confirm confusion about the terms *cloud* and *cloud computing* and also suggest that for some in the research computing community, *cloud computing* is just another term for *grid computing*. To this end, the grid community in Europe came together recently, under the auspices of the Open Grid Forum, for the Cloudscape II conference (http://www.ogfeurope.eu/pages/selecteddocument.aspx?id_documento=316b846a-0993-401c-9007-f63d2cba5ed9), exploring issues pertaining to both grid computing and cloud computing. There was a convergence of interest in both areas, since the need for standards and for interoperability between cloud systems was seen to be just as relevant as is this need in the world of grid computing.

Privacy

Broadly speaking, the issue of privacy is one that concerns campus legal office and senior management administrators more than end users. Students, many of whom use Gmail and Hotmail before coming to college, do not appear to be overly concerned with privacy or with the idea that companies like Google and Microsoft can access their e-mails. In 2008, for example, several Australian universities did an analysis of their student e-mail systems and found that more than 80 percent of students were not using the university-provided e-mail system and had already set up a forward to their Hotmail or Gmail accounts.⁵

Faculty and staff have proven to be a bit more vocal about the privacy issues. Particularly in the research domain, they have expressed concerns about sensitive research data going outside the bounds of the college or university. Yet though these concerns have been noted, some researchers readily use Gmail or Hotmail to move data to colleagues at other institutions as a way to get around firewalls or e-mail system attachment

limitations. Anecdotal evidence suggests that faculty and staff are becoming more comfortable with trusting some cloud providers with their sensitive data.

Interestingly, the corporate policies of Microsoft and Google are very strong

Privacy concerns campus legal office and senior management administrators more than end users.



on the issue of privacy. For example, in the early days of Australian universities moving their student e-mail services to cloud providers, many of the universities wanted to retain the ability to examine e-mail logs in cases where there was reasonable evidence to implicate a student for a breach of the university code of conduct. The services offered by Google and Microsoft did not offer this capability, and providing it proved to be challenging for both companies, since it contravened their internal company privacy policies. In a conversation between the Council of Australian University Directors of Information Technology (CAUDIT) and a number of cloud providers on the issue of privacy—and, in particular, the misuse of user data—the providers argued that the risks associated with misusing user data were significant. If, for example, it came to light that Microsoft or Google was misusing data, the higher education sector would immediately lose trust. The stakes for large, well-known cloud providers are significant: their reputation and brand are on the line.

Contractual and Jurisdictional Issues

Contractual and jurisdictional issues remain a key challenge when moving services to the cloud. When the contract is based around national law and when the data and transactions are occurring within the same country, the process is relatively easy. But when the data and

THE JISC STUDIES

The Joint Information Systems Committee (JISC) in the United Kingdom has launched three studies to investigate the cloud for higher education from a number of perspectives. As a whole, these studies will provide a comprehensive picture of the issues at the present time.

The first JISC study is a technical review of the cloud in terms of research requirements and is looking at the issues required to make cloud computing feasible, including security, distributed programming abstractions, scheduling, and benchmarking. The second study focuses on the organizational and human issues surrounding cloud computing for research, such as service-level agreements, skills and competencies, licensing, security, and the best mix of cloud and more conventional HPC for various research communities. The third study widens the scope to look at the implications of cloud computing for enterprise computing and also at the environmental implications of moving computing into the cloud.

Find out more on the JISC website (<http://www.jisc.ac.uk>).

With the use of cloud service providers, there can be a need to rearchitect applications.



transactions are happening offshore, the challenges increase dramatically.

Further, the implications of the USA PATRIOT Act lead to very real impediments for non-U.S. institutions when dealing with U.S. companies. For example, a college or university will encounter difficulties if it is collaborating with researchers in a country that is under embargo from the U.S. export clauses. In an increasingly connected world with collaborative research spanning many international boundaries, U.S. companies will need to work with their lawmakers to address these issues if they want large-scale acceptance of their offerings outside the United States. Some cloud providers, such as Microsoft, have recognized this as a significant problem and have shown a willingness to work with institutions to change the jurisdiction of contracts and to address as many of these concerns as possible.

Risk and Nonperformance

Strategies need to be employed to manage not only risk but also those situations when a cloud service provider suddenly and unexpectedly stops delivering services. For example in Australia, Curtin University has set up a strategic partnership among five providers with which it works closely so that services can be moved between partners as required. In another example, Victoria University is developing an innovative strategy to rearchitect its internal infrastructure so that integration with the cloud can be done in a more seamless and elastic manner. In this environment, data and processing can move transparently between the university infrastructure and the cloud provider. The university can then

readily address and control performance issues by adjusting, in real time, how much of the service sits in the cloud and how much stays on the university infrastructure.

An important aspect of risk associated with privacy and data is to correctly classify and value the intellectual property of the data

that is being entrusted to the cloud provider and then to build an appropriate indemnity into the contractual arrangements. Doing so not only ensures that the cloud service provider will correctly understand and value the data but also financially compensates the college or university in the event of a leak of sensitive data of significant value.

The other issue to be considered is that in moving services to the cloud, the institution no longer retains direct access and control. The college or university must ensure that it has strategies and processes in place so that should a cloud service stop working in the intended way, causing problems for other systems, the institution can manage the situation even though it cannot readily control or access the offending service.

Interoperability

To date, use of the cloud has largely focused on specific services, but as the practice increases, a looming issue is interoperability between clouds and the lack of standards. If an institution needs to quickly move a service from one cloud provider to another, there are few standards to ensure that the data associated with the service is readily portable between providers. Further, as adoption of the cloud continues, interoperability between services running on private, public, and hybrid clouds will introduce further complexity.

Some work has been done to develop open-cloud and interoperability standards.⁶ Still in their infancy, these standards are lightweight, and not all cloud providers have committed to them. A number of CIOs believe that this issue will become an increasing headache in the next years.

Network Capacity

One of the concerns with moving services to the cloud relates to latency and to the speed of the networks between the college or university and the provider. Peering R&E networks with cloud providers goes some way toward addressing this. The National Research and Education Network (NREN) in Australia—the Australian Academic Research Network (AARNet)—interconnects colleges and universities, typically with 10GB circuits in metropolitan areas and 1GB circuits in regional areas. The AARNet also domestically peers with a number of commercial providers—including Microsoft, Google, and Akamai—and internationally peers with other NRENs around the world. An enabler in terms of quality of service, the AARNet demonstrates the importance of leveraging NRENs to facilitate the use of the cloud.

In the Netherlands, network capacity and flexibility is addressed in the GigaPort3 program. The aim of the GigaPort3 project is to raise the SURFnet network infrastructure to a higher level and to integrate it seamlessly with other IT infrastructure facilities. The new network, SURFnet7, will increase the flexibility and manageability of lightpaths, implement Next Generation Ethernet, provide custom-size bandwidth cost-effectively and on time, extend the dynamic light-path service to allow applications to initiate a lightpath automatically, and further professionalize the international NetherLight network node. Part of the plans is to create seamless bridging between fixed and mobile networks for both research and education to use services anytime, anyplace.⁷

Rearchitecting

With the use of cloud service providers, there can be a need to rearchitect applications. Instead of a fat client, access typically changes to the browser utilizing small widgets. This can necessitate the rearchitecting of applications to make efficient use of the network and to keep the transactions at the back end.

Building useful, resilient, scalable,

and cost-effective interfaces to cloud systems for use by researchers or the enterprise remains as challenging as doing so for the grid or HPC and requires many of the same skills and approaches. Indeed, a new layer of abstraction seems to be emerging to cope with the opportunities to harness the power of the cloud for research computing applications. This “data-center operating system” will hide the technical implementation of different cloud offerings from the end user behind interfaces that allow rapid deployment and delivery of new services, whether for e-science or enterprise computing.

Staff

Over the coming years, the skills set required of the staff within the IT department to support and manage a portfolio that includes the cloud will differ from those required today. Staff will need to be more adept and sophisticated at managing strategic relationships with service providers, as well as being competent managers of services that span both the institution and the providers under contractual arrangements. Some IT staff may view the cloud as a threat to their future employment rather than an opportunity—a situation that will need to be managed.

Staff will need to be more adept and sophisticated at managing strategic relationships with service providers.



Perceptions

The JISC studies revealed other perception problems, including a reluctance by some IT leaders to give up the prestige that comes from owning and being responsible for large hardware installations. Another issue that has surfaced is that on initial inspection, using cloud computing for research looks to be more expensive than traditional approaches.

However, further examination shows that these more traditional ways of providing research computing hide many costs from end users—not the least being the cost of electricity and other overheads, including technical development and support, which is often delivered by people whose real job is research. Finally, the cost of moving to the cloud—in particular the need to rewrite algorithms to take advantage of new platforms—is also perceived as a barrier.

The Next Level

Thinking, planning, and working in the cloud requires colleges and universities to recognize that individual institutions are limited in their ability to scale their infrastructure and services to the levels of commercial providers, no matter how large and complex those institutions may be. In the *EDUCAUSE Review* article “Above-Campus Services: Shaping the Promise of Cloud Computing in Higher Education,” Brad Wheeler and Shelton Waggener defined three broad categories of “above-campus” sourcing models: commercial, institutional, and consortium sourcing.⁸ As the commercial providers of cloud services gain an increasing foothold in the sector, consideration needs to be given to institutional and consortium

sourcing, as well as to the value that national associations like CAUDIT, CUCCIO, EDUCAUSE, JISC, and SURF can provide to help higher education institutions.

The potential benefits of commercially sourced cloud services are compelling: reduced institutional costs, increased efficiencies through standardization, improved collaboration

across institutions, and the ability to shift management attention away from operational concerns over nondifferentiating services to those that are strategic to the academic mission. The risks, though not insignificant, are familiar territory for most CIOs: information security, negotiating and managing contracts and associated service levels, and potentially high switching costs stemming from lock-in by service providers. However, in the absence of collective and coordinated action to date, each institution that has deployed commercially sourced services has developed its own one-off contracts with providers, despite nearly identical needs across the sector. Sourcing student e-mail services through Google is a clear example.

Thus, one of the first things that national associations can do is to help draw up contracts with suppliers of cloud services and to help create an infrastructure that will facilitate the outsourcing of e-mail accounts. In Australia, CAUDIT has acted as a platform for collaboration, enabling early adopters to share their experiences with each other. Specifically, CAUDIT has facilitated not only the sharing of contracts between its members but also the subsequent refinement of the contracts to a template that is mutually acceptable for both the college/university and the cloud provider, reducing the legal costs for everyone. Likewise, in the Netherlands, SURF has drawn up a contract with Google to be a member of SURFederation, the Dutch authentication and authorization platform for all higher education and research. Microsoft is expected to follow soon. This important arrangement is a simple but very effective way to make the cloud services of different providers accessible and to give institutions the opportunity to choose which provider and which services best fit their needs. SURFederation is already used for a large variety of services from the cloud, from publishers to theaters.

But can these national associations do more? Can or should they play a more active role in the provision of cloud services

to their members? CAUDIT is in very early discussions with a cloud provider to explore this very notion. One possible scenario is for the national association to partner with the cloud provider to host the services inside the sector for its members on behalf of the cloud provider. Exactly how such a service would be provided remains a question of much debate, but an interesting possibility exists for the NREN to move up the stack and provide value-added services beyond the core network.

In the Netherlands, SURFdiensten is the broker organization for Dutch higher education institutions. It procures not only journal and software licences but also hardware and services for all higher education institutions in the Netherlands. SURFdiensten and SURFnet work closely together with the institutions in dealing with commercial service providers such as Google and Microsoft. The expertise of both organizations is thus improved and put to use for all SURF members.

Although at an earlier stage of development, CUCCIO is coordinating the early cloud activities of Canadian universities by leveraging its experience in setting up a Canadian Access Federation that offers federated identity services through eduroam and Shibboleth.

Meanwhile, the board of EDUCAUSE is exploring new leadership roles for the association, including that of *active steward*. For EDUCAUSE, active stewardship of above-campus activities would represent a challenging, nuanced, and vital role in service to the community. It is a role that would leverage the association's existing competencies while demanding that it stretch to develop new capabilities where the need for community leadership is acute. This could include policy development, directed advocacy, and most important, coordinated action that directly engages other member associations, community partners, and service providers.

Taking these ideas to the next level, what if there were an entity that acted as higher education's trusted broker in developing contract templates, negotiating

services, and managing effective relationships with commercial sourcing firms? The broker—or brokers—could fulfill the requirements of demand aggregation, resulting in optimal scale economies for commercial sourcing. Internationally, we already have excellent examples of community-driven organizations that are, in effect, trusted brokers: regional and national research networking societies. The success of these organizations (Internet2 in the United States, CANARIE in Canada, AARNet in Australia, TERENA in Europe, JANET in the United Kingdom, and SURF in the Netherlands) provides an attractive and instructive model of how colleges and universities have worked together to combine their collective networking service requirements, comprising contracting, standardization, service-level management, and support.

Consider the alternative: in the absence of regional networking organizations, each institution would have to negotiate its own bandwidth contracts and service agreements. Commercial networking providers would be required to negotiate hundreds of one-off contracts

What if there were an entity that acted as higher education's trusted broker?



with individual institutions, each contract being highly similar if not identical to the others. Instead, by aggregating multi-institutional demand with a single trusted entity—their national or regional research networking organization—colleges and universities achieved scale economies, best pricing, operational sustainability, and considerable leverage to mitigate switching costs and therefore encourage an environment of healthy competition.

In short, national and regional research networking organizations are

good models for thinking about cloud consortium services. These organizations are driven by the needs of their member institutions. They are rooted in a very high level of trust and transparency in their operation and governance. They have successfully developed the staff competencies to negotiate large-scale contracts, manage relationships with and performance by service providers, define and maintain outstanding levels of network service, and provide strong client support. These organizations work closely and productively with private-sector service providers to ensure that the needs of the higher education community are understood and met. In short, the existing competencies of these advanced networking organizations could enable them to move beyond networking and to act as a demand aggregator and broker for commercial cloud requirements and services. Indeed, in the province of British Columbia, Canada, the governing board of the regional advanced networking organization (BCNET) recently revised its organizational vision and mission to broaden the potential scope of activities beyond pure data networking.

Finally, there are also opportunities for groups of higher education institutions or a specially formed sector-owned organization to collaborate and create private or hybrid clouds to provide services. The Australian Research Collaboration Service (ARCS) is an example in which pooling the collective resources of the sector to provide a suite of services in a private cloud has provided real value to researchers. ARCS currently offers a service where any researcher in Australia can have a 25GB data store⁹ (more can be requested) that is in a private cloud

and that is available anywhere there is a reliable connection to the Internet. The service leverages the data storage infrastructure of a number of HPC facilities and universities in Australia; the researchers do not know where their data is being stored within the sector. Such initiatives are the beginning of a host of specific R&E services for which commercial providers are unlikely to provide a suitable solution.

The private cloud also offers a tremendous opportunity for colleges and

The cooperation of some IT departments may generate simply a puff, not a cloud.



universities to collaborate and “deduplicate” core services across institutions. Such arrangements, whether provided by an organization within the sector or in partnership with a commercial provider, could enable universities to realize significant financial savings. In Australia, this is already happening at the regional level, and over the next one-to-two years there is a reasonable likelihood of national initiatives being developed through CAUDIT.

On the other hand, although creating a private or shared cloud can help institutions maintain control over the services and can lead to some economy of scale, the result may still be a far cry from the advantages achieved by using services from companies like Amazon, Google, or Microsoft. Compared with these companies, the cooperation of some IT departments may generate simply a puff, not a cloud. More important, a private cloud can imply forced customership, a business model that lacks the driving force of competition to innovate and to provide users with the diversity to which they have become accustomed. This also implies that institutions should first of all reconsider their services: which services

should be provided by the institutions, which services can be outsourced, and which services can be drawn from the market by the individual user.

A Consortium Example: Kualiti Ready

Kualiti Ready (KR) is a community-source project chartered to provide a business continuity planning (BCP) service. It is also an example of higher education institutions organizing themselves to provide cloud services. Launched by the Kualiti Foundation in 2010, KR's ancestry

dates to software tools, developed at the University of California–Berkeley, that enable departments to document and maintain their local business continuity plans

and then to roll up those plans to a department, faculty, or institution level. Adoption of the tools quickly spread—first to most of the other campuses of the UC system, where the project became known as UC Ready, and later to dozens of other institutions in the United States and beyond, which downloaded early versions of the code to install and run locally. Recognizing the importance of the tools to higher education and the need for a sustainability plan, Berkeley and the UC system invested substantial additional resources to add requested functionality for pandemic preparedness and to bring the software into line with the Kualiti Foundation's technical and licensing standards.

Like other systems developed and supported by the Kualiti Foundation, KR can be downloaded at no cost from the Foundation's public website (<http://www.kualiti.org>). Importantly, in a first for the Foundation, this software is also being offered through a Software as a Service (SaaS) model on a subscription basis. Under the SaaS model, institutions need not incur the cost and operational overhead of their own infrastructure, version management, and technical support.

Instead, they pay an annual subscription fee that provides them access to the hosted service on the web, fully customizable to meet local needs, including soon the ability to access the service with their local identity credentials supported by a Shibboleth federated identity interface.

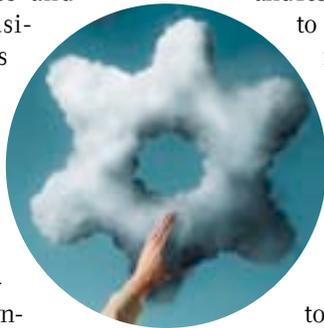
In this way, even the smallest institutions with modest budgets and minimal IT staff can affordably offer BCP capabilities to their campuses. During the first few years of the service, founding partners will provide SaaS hosting services, starting with UC Berkeley and Indiana University. Canadian hosting will be offered by the University of Toronto and the University of British Columbia. The need for nationally-based hosting is driven by Canadian privacy legislation, described earlier, which places strict limits on the information that can be stored outside Canada. But as a consequence, the community cloud is already international.

KR is a good early example of some key principles that are emerging to guide cloud developments:

- Aggregation of demand within the community cloud makes sense in circumstances where specific expertise or domain knowledge exists that differentiates the service from available public cloud offerings. A BCP tool tailored to the needs of colleges and universities, KR works as a community cloud offering, whereas massive-scale student e-mail is probably better provided through commercial sourcing.
- The service is offered by an existing, known, and trusted community organization. Clearly, the Kualiti Foundation has a relatively brief track record in providing direct services. However, the modest scope and relatively low risk of this first community cloud offering present an opportunity to build further competence within the Foundation in case more ambitious SaaS opportunities emerge.
- The service is a modest first step in an overall strategic direction without being part of a master game-plan. This is consistent with the notion of

learning-by-doing and then adjusting course based on experience. Without minimizing the importance of a BCP system, this is not an application that features high volumes of transaction processing or complex functionality. In short, a useful service can be affordably offered to the community, at minimal risk to all parties, while providing an excellent opportunity to learn, adjust, and improve.

- The service increases choice to the higher education community. The Quali Foundation is supported by and exists to serve its member institutions—colleges and universities. The business model (or perhaps more accurately, the service model) does not limit access to the SaaS offering but increases choice by also offering the free community-source-licensed software download. Institutions can pick the model that works best for their needs.



The time has come for colleges and universities to set a course for the adoption of cloud services and to consider what role community sourcing plays in their strategic direction. The Quali Foundation is just one example. Other consortia will almost certainly build and improve on that experience. Still missing from the picture are *consortium makers* of the type advocated by Wheeler and Waggener.¹⁰ The role of a consortium maker, operating at the higher education sector level, is to enable the formation of member-led consortia. The national associations have many of the competencies, plus the trust and experience with their communities, to consider taking on this leadership role.

Conclusion

Many colleges and universities are actively moving a variety of services to the cloud, with the pace and the volume of adoption increasing. Several CIOs have predicted that higher education institu-

tions will get out of the game of running the monolithic enterprise systems and will move the finance, human resources, and student systems into the cloud over the next five to ten years. Indeed, one CIO commented: “My boss used to constantly tell me I had my head in the clouds, now I tell him my strategy is in the cloud.”¹¹ Many CIOs do indeed view the cloud as a key strategy for the future—a strategy that will enable them to add greater value to their institution at a more strategic level, rather than constantly focusing on keeping the lights running.

Cloud services offer higher education and research institutions the power to choose: the opportunity to rethink which services are needed to support education and research and what will be the best way to deliver those services. Many services are readily available in the public cloud. Some services need to be procured through the institution’s IT department. Only a few services will require custom development, either alone or in partnership with other institutions. The final result will most likely be a loosely coupled, customized arrangement consisting of off-the-shelf systems and services based on proven technology.

With many in higher education today eyeing the potential of the cloud, the question now is not so much “Is cloud computing a good idea?” The key question to answer is: “What can we do with the cloud?” ■

Notes

1. For example, the Microsoft 700,000-square-foot data center in Chicago can accommodate 56 computing containers. Each container is filled with 1,800 to 2,500 servers. The Chicago data center is just one of two dozen centers that Microsoft plans to build. See the Microsoft Global Foundation Services website: <<http://www.globalfoundation-services.com/index.html>>.
2. CAUDIT Member Survey, 2010.
3. National Science Foundation, “Microsoft and NSF Enable Research in the Cloud,” press release 10-023, February 4, 2010, <http://www.nsf.gov/news/news_summ.jsp?cntn_id=116336&org=NSF&from=news>.

4. Gordon Cook, *Building a National Knowledge Infrastructure: How Dutch Pragmatism Nurtures a 21st-century Economy*, The Cook Report on Internet Protocol (Utrecht, The Netherlands: SURF, 2010), <<http://www.surfnet.nl/en/nieuws/Pages/Book%20Building%20a%20national%20knowledge%20infrastructure%20offered%20to%20Minister%20Plasterk.aspx>>.
5. Neil Thelander, IT Director, Queensland University of Technology, Australia.
6. See the Cloud Standards Coordination wiki: <http://cloud-standards.org/wiki/index.php?title=Main_Page#Cloud_Standards_Coordination>.
7. SURFnet, *GigaPort3 Annual Plan 2010*, December 22, 2009, <<http://www.surfnet.nl/Documents/GigaPort3-AP2010-EXT-1.0.pdf>>.
8. Brad Wheeler and Shelton Waggener, “Above-Campus Services: Shaping the Promise of Cloud Computing in Higher Education,” *EDUCAUSE Review*, vol. 44, no. 6 (November/December 2009), pp. 52–67, <<http://www.educause.edu/library/erm0963>>.
9. “ARCS Data Fabric: What It Is,” ARCS website: <<http://www.arcs.org.au/products-services/data-services/arcs-data-fabric-1/arcs-data-fabric>>.
10. Wheeler and Waggener, “Above-Campus Services.”
11. Quote from Allan Morris, Executive Director, Information Technology Services, RMIT University, Melbourne, Australia.

© 2010 Rob Bristow, Ted Dodds, Richard Northam, and Leo Plugge



Rob Bristow is Programme Manager for e-Administration, Joint Information Systems Committee (JISC), Bristol, United Kingdom.



Ted Dodds is Vice Provost for IT, the University of British Columbia, Canada, and is a member of the board of directors of CUCCIO, EDUCAUSE, and the Quali Foundation.



Richard Northam is General Manager, Council of Australian University Directors of Information Technology (CAUDIT), Canberra, ACT, Australia.



Leo Plugge is Executive Secretary, Wetenschappelijk TechnischeRaad, Stichting SURF, Utrecht, Netherlands.