JISC

University of Southampton
School of Education

Technical Review of Using Cloud for Research

Guidance Notes to Cloud Infrastructure Service Providers

May, 2010

Introduction

Provisioning and maintenance of research computing facilities is a core part of the IT strategies of High Education Institutions (HEIs). Existing research computing facilities may include HPC clusters, campus grid infrastructure, and virtualised computational resources, which have been widely employed to support research activities across different disciplines. Additionally, investment is increasingly required for hardware procurement and system maintenance.

The recent success in the use of cloud in business has introduced a new model that can potentially be employed by HEIs to provide cost-effective research computing. Cloud computing, due to its commercial nature, is still new to research communities, and a number of questions arise about its applicability to the HEI sector. Would cloud computing be appropriate if adopted for research? What are the differences between a research cloud and a commercial cloud? How could HEIs adopt cloud computing for research?

This document provides guidance notes to potential cloud infrastructure service providers and helps identify the issues involved in adopting cloud computing for research activities in this early stage of development of the field.

Understanding Cloud Computing

There are different interpretations of cloud computing. In order to provide a common understanding from a technical perspective, we define cloud computing as an emerging business model that delivers computing services over the Internet in an elastic, self-serviced, self-managed, cost-effective manner with guaranteed Quality of Service (QoS).

There are four types of stakeholder in a cloud computing environment: end user, Cloud Service Provider (CSP), Cloud Tool Provider (CTP), and Cloud Service Vendor (CSV). Individual stakeholders have different views of the characteristics of cloud computing.

Essential Technical Characteristics

- **Elastic** and dynamic scaling in and out.
- **Self-service** provisioning and management
- **Self-management** and automatic scaling
- **Cost Effectiveness** and multi-tenancy on per-usage basis

Current commercial CSPs deliver cloud computing capabilities at three hierarchical service layers: the infrastructure layer (Infrastructure as a Service, IaaS), the platform layer (Platform as a Service, PaaS), and the application layer (Software as a Service, SaaS). Each cloud delivery layer targets a different group of stakeholders and provides specific self-service and self-management facilities.

A cloud system (IaaS, PaaS, and SaaS) can be deployed in three different models as well, as described in the following paragraphs.

- **Private Cloud –** Cloud services are owned and delivered only within an enterprise or organisation. Private clouds employ advances in virtualisation technologies and automated management technologies to enhance scalability and utilisation of local data centres while reducing administrative and management tasks. Private clouds can be built on existing on-premises computing infrastructures using open-source cloud software (Cloudware) or third-party

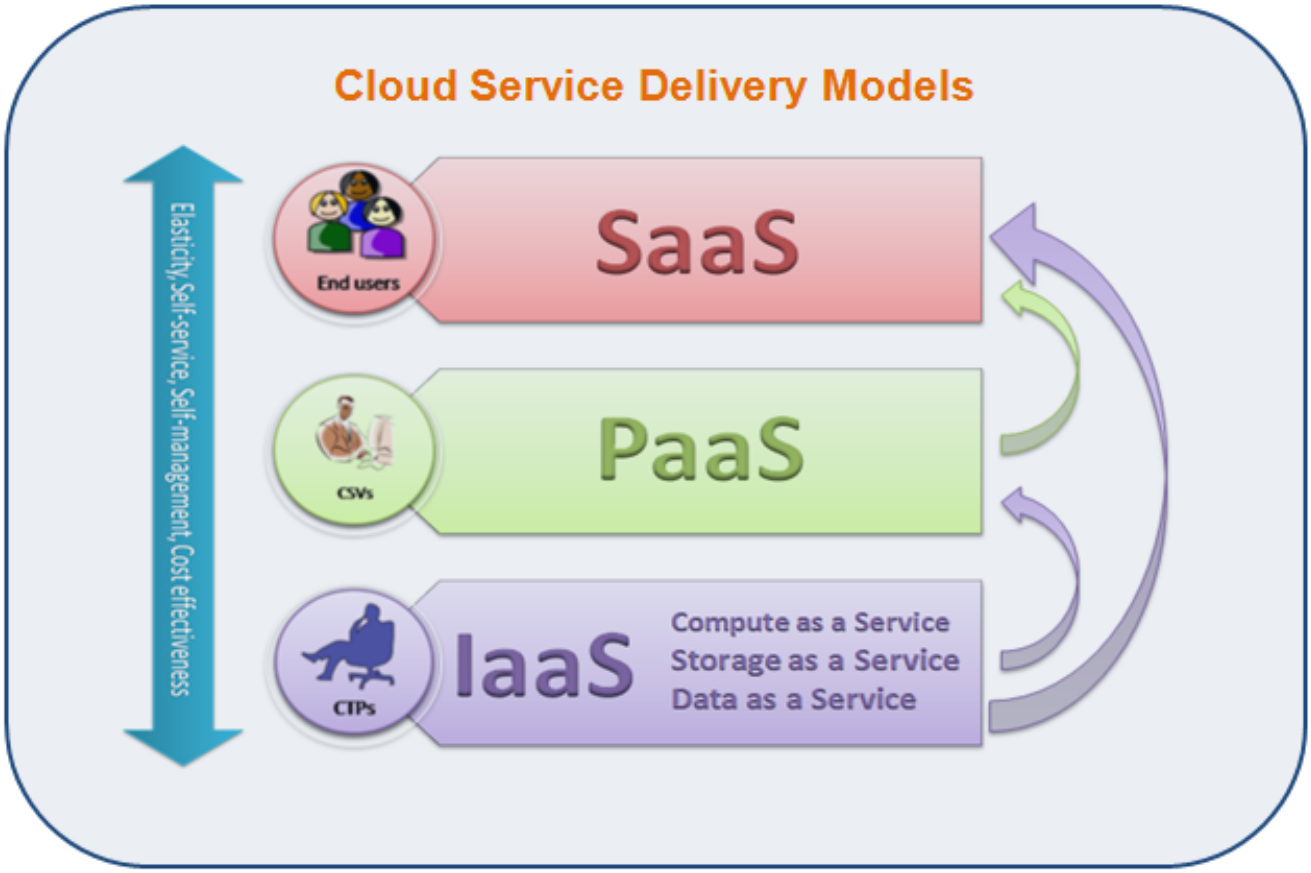

commercial offerings to meet the needs within an organisation.

- **Public Cloud –** A public cloud provides computing services that are publicly accessible through standard self-service APIs over the Internet. Public cloud services may be free for development and test purposes, or may charge on a per-usage basis in a production cloud environment. A private cloud can be easily turned into a public cloud by opening self-service APIs for public access. Because most public cloud services are paid for, service level agreement (SLA) enforcement is of great importance in a public cloud. A public cloud may offer a certain degree of control of virtualised resources (e.g. virtual images, virtual machine instances) while reserving full control over the local computing infrastructure.
- **Hybrid cloud** – A mixed deployment model employing both private and public infrastructures. The hybrid cloud is mostly used for offloading to a public cloud while maintaining the desired degree of control inside a private cloud. The hybrid cloud model is mostly adopted for maintaining sensitive data inside a local private cloud. In this report, we distinguish the cloud bursting model (see below) from the hybrid cloud model because cloud bursting does not necessarily involve local private cloud infrastructure. The hybrid cloud model may involve more than one cloud provided by more than one CSP, therefore in this model standardisation is of more importance than in other cloud deployment models.
- **Community Cloud** – Multiple organisations with common concerns, such as security, policy, interests, and missions, may share cloud infrastructures across administrative domains, forming a community cloud. Infrastructure constituents of a community cloud can be managed by partner organisations or by a third party.

Understanding Research Cloud Computing

The differences between commercial and research cloud usage can be best demonstrated by comparing their respective stakeholders. A typical cloudy research environment involves four broad active stakeholders, the HE institutions (HEIs), researchers, public CSPs, and standards bodies.

Rather than explicitly separating end user, developer, and tool provider roles, researchers are able to take on the responsibilities of all of these roles, and have full control throughout the cloud application development lifecycle. For example, HE institutions can build private cloud services upon existing IT infrastructure, while also acting as end users of public CSPs, to burst and offload to public cloud services fluctuating demands. Therefore, the HEI acts in a separate role from that of a public CSP. Interoperability is of utmost importance in the cloudy research environment so that shared research services can be deployed in a consistent manner across multiple HEI private cloud infrastructures.

Appreciating the Facts

Before starting to build a particular cloud strategy, there are some facts that might need to be appreciated.

- Cloud computing for research is still in its infancy. There is little support from either commercial public CSPs or open-source Cloudware vendors.
- Development of open-source Cloudware has only just started. Little software is available, and what is available is not user friendly.
- Maintaining an in-house cloud does not save on administrative tasks by comparison with an existing IT infrastructure, mainly because of the immaturity of open-source Cloudware.
- Current public CSPs may not provide functional management facilities, though third-party management tools are available.

Understanding Your Case Scenarios

In the TeciRes final report, extensive reviews of current practices using cloud in UK HEIs were categorised into six case scenarios. The four scenarios which apply to HEIs are as follows.

- **Cloud Bursting** – In this scenario, a HEI provides research computing services to researchers, while bursting and offloading to public cloud services when fluctuating requirements need to be met. Public cloud services are invisible to researchers. This scenario is differentiated from the hybrid cloud scenario because no private middleware is deployed on research computing facilities inside an HEI. A HEI

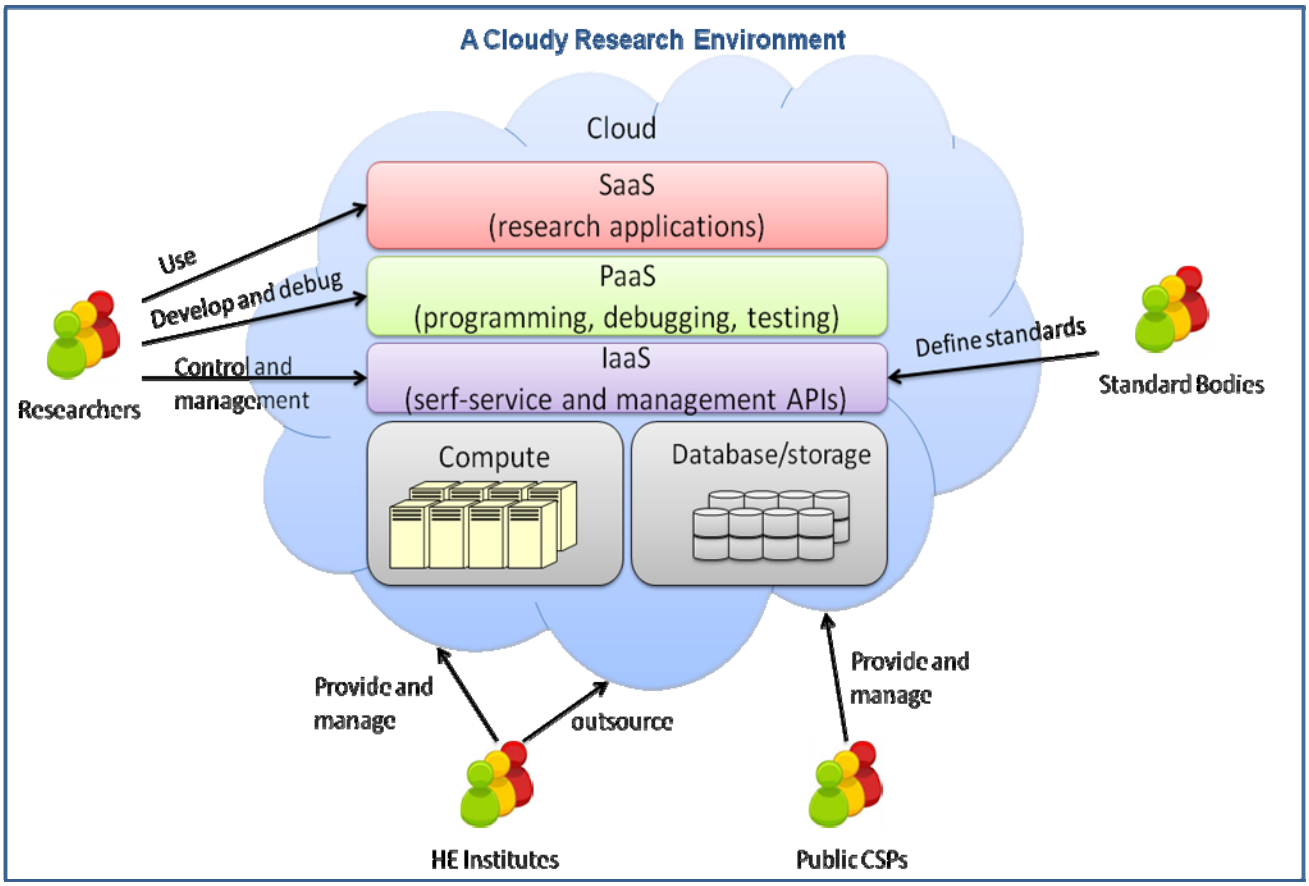

may also outsource to multiple CSPs. A Grid-Cloud bridge is an interesting example case in this scenario. In the UK, most HEIs have a campus cloud deployed, and are interested in delegating computational tasks to public cloud services when the local system is overloaded.

- **Private Cloud** – In this scenario, HEIs leverage virtualization technologies and open-source Cloudware to build a private research environment for researchers inside the organization. This scenario is widely adopted by HEIs to increase local resource utilisation.
- **Hybrid Cloud** – This scenario involves multiple clouds working together including, in particular, cloud services provided by public CSPs and HEIs. This scenario is ideal for HEIs to maintain a certain degree of control over their cloud infrastructures. For example, by keeping research data in local storage cloud services, HEIs maintain full control over data privacy.
- **Community Cloud** – This scenario is a research-specific vision that allows HEIs with common research scopes or common requirements to share cloud services and management tasks.

Given the scenario for a particular example case, the Tecires final report provides the detailed technical requirements that apply.

Understanding CSPs' Architecture

Both cloud-bursting and hybrid-cloud scenarios involves outsourcing to public CSPs. It is therefore important to evaluate their underlying architectures and their capabilities on offer. The evaluation should at least cover system architecture, network architecture, storage architecture, security strategies, and self-service and self-management capabilities.

A list of key questions about the CSP's architecture, as follows, will help an evaluation.

- Do you share with other tenants on a virtualised platform?
- If shared, what are the underlying virtualisation technologies being employed?
- How many and what operating systems are supported?
- Does the architecture allow self-provisioning and management of virtualised resources (i.e. virtual machines, virtual machine images, etc)?
- Does it provide self-management facilities allowing monitoring and accounting for actual applications usage?
- Does it allow auto-scaling?
- What types of interfaces (i.e. RESTful, command line interface, programming APIs, or Web portal) are provided which allow the management of the promised cloud capabilities?
- Are these interfaces standards compliant?
- Are programming abstractions or platforms provided?
- How is data privacy managed in the cloud?
- Can access controls be defined?
- What are the technologies which enable storage virtualisation?
- Does it support a Virtual Private Network (VPN)?

Rules of Thumb

How to choose public CSPs

- Always evaluate multiple (at least three) CSPs' offerings;
- Compare their system architectures and capabilities using a table or checklist (TeciRes final report gives an example in section 4);
- Try out before committing to a particular CSP;
- If online information of CSPs cannot accommodate your advanced requirements (e.g. a large-scale complex system), contact them.

Understanding the Existing Environment

Understanding the in-house computing environment is the first step in making decisions on the way to developing and delivering custom cloud solutions. Most HEIs maintain research computing facilities underpinned by various computing technologies. Grid computing, for example, is employed in a number of UK HEIs. Migrating from grid to cloud needs careful consideration given the fact that current open-

source Cloudware does not dramatically reduce maintenance effort as much as might be expected. An evaluation of the difficulties of migrating existing projects and software to be cloud-ready will be needed.

These are some questions about the current existing environment which are relevant.

- Does the existing system architecture use virtualisation technologies, and what are they?
- What is the local network architecture, and does it offer an Internet Virtual Private Network (VPN) service?
- What are the local storage architectures (i.e. storage area network or network-attached storage), and associated data backup and recovery processes?
- Are there any existing cloud-ready management tools available in the current system?

Building a Private Cloud

Given an understanding of the current, existing environment, selecting an appropriate open-source Cloudware may be the next step in building a private cloud solution. Only a few research projects on developing open-source cloud middleware software (Cloudware) have been initiated recently, and the choice of such software is currently limited.

The three most prominent Cloudware provisions are:

- **Eucalyptus Systems** – an open-source private cloud software enabling an organisation to establish an in-house cloud computing environment.
- **OpenNebula** – a standards-compatible open-source toolkit for building public, private and hybrid IaaS clouds using on-premises IT infrastructure. The toolkit is partially funded by the open-source Resource and Services Virtualisation without Barriers (Reservoir) project, funded by the European Union under FP7.
- **Nimbus** – a purpose-built cloud toolkit that turns a computer cluster into an IaaS cloud by employing virtualisation technologies.

The TeciRes final report lists the features of these open Cloudware provisions in a comparison table. Choosing an appropriate Cloudware, however, is more complex than choosing a public CSP. The performance of two private clouds with the same Cloudware may differ dramatically, depending on the local IT infrastructure and the applications being run. Therefore small-scale test-beds are recommended as always useful to try out at least two different Cloudware provisions. An evaluation within a real runtime environment with meaningful benchmarks will better support a final decision.

Rules of Thumb

How to build a private cloud

- Always evaluate your existing IT infrastructure firstly and understand your problems;
- Try out before committing to a particular Cloudware;
- Use a comparison table to list all capabilities that Cloudware candidate offer;
- Always start from a same-scale testbed;
- Keep yourself updated from experienced users and Cloudware vendors;

Leveraging Third-Party Tools

Apart from commercial CSPs and open-source Cloudware providers, there is a group of third-party companies that play an important role in supplying cloud-based management services. These services normally provide unified management capabilities and can be adapted for custom manageability purposes. The TeciRes project reviewed three cloud-based services which deliver self-management features to cloud tenants allowing control of remote cloud resources.

- **NimSoft** – the fastest growing provider of performance- and availability-monitoring software. Provides public and private monitoring solutions based on a unified monitoring architecture.

- **RightScale** – a management platform that enables the creation and management of scalable cloud applications across cloud environments.
- **CohesiveFT's VPN-Cubed** – claimed to be the first commercial solution that enables security control in a cloud and cross clouds.

An HEI may leverage these tools or incorporate them into custom solutions for its private cloud framework, making it more automated, manageable, and easy to use.

Making It Hybrid

Considering the relatively limited resources an HEI can typically offer, a private cloud can offload to a public commercial cloud in accommodating fluctuating demand, while delegating part of the administrative tasks to the public CSP. Although there is no example case identified in the TeciRes final report, Hybrid cloud seems to be a promising deployment model that contributes optimised local resource utilisation and selective control over cloud resources from a HEI's perspective.

Becoming a Community

Community cloud is a special cloud deployment model, involving multiple private cloud providers to share access. Users from one cloud are able to access the resources of the other. The community cloud therefore requires advanced user identity management and federated identity management. The main benefit of a community cloud deployment is to share the management tasks of a large-scale cloud system between multiple organisations. In order to ensure seamless sharing of cloud infrastructures, HEIs need to consider standards compatibility when choosing open Cloudware for their private clouds.

Conclusion

As this guidance note highlights, there are different considerations in choosing public CSPs and open Cloudware depending on the HEI-specific cloud strategy. Building a small-scale private cloud infrastructure within a HEI is a very attractive option to achieve a highly flexible and scalable research computing environment. The adoption of cloud computing in HEIs will also contribute to the evolution and development of cloud computing technologies.

For more Information

- **TeciRes Project**
  http://tecires.ecs.soton.ac.uk
- **TeciRes Final Report**
  http://tecires.ecs.soton.ac.uk/documents
- **Review of Using Cloud Computing for Research**
  http://www.jisc.ac.uk/fundingopportunities/funding_calls/2009/09/cloudcomputing.aspx
- **Review of Environmental and Organisational Implications of Cloud Computing for HE and Further Education**
  http://www.jisc.ac.uk/whatwedo/programmes/greeningict/environmentalreviewcloudcomp.aspx

This document was produced by Xiaoyu Chen, Gary B. Wills, Lester Gilbert, and David Bacigalupo from the Learning Societies Lab of the University of Southampton during the course of the TeciRes project funded by JISC.