SOUTH EAST EUROPE
Transnational Cooperation Programme

Programme co-funded by the EUROPEAN UNION

secovia

A brief introduction to Cloud Computing

Jointly for our common future

A brief introduction to Cloud Computing

Cloud computing is not a single technology, but rather an **approach** for providing computer resources including pure computing power, complete computing infrastructures, applications or single processes or functionalities over the Internet. The "Cloud" in "Cloud Computing" is that set of hardware, networks, storage, services, and interfaces that together deliver the particular desired computing as **a service**. Cloud services include the delivery of software, infrastructure, and storage over the Internet (either as separate components or a complete platform) based on what the user has asked for.

According to the generally accepted Cloud computing definition[1] first proposed by the National Institute of Standards and Technology in the United States, there are five essential characteristics, which are **common among all cloud computing services**:

1. **On-demand self-service**: A consumer can unilaterally provision computing capabilities, such as server time and network storage, as needed automatically without requiring human interaction with each service's provider.
2. **Broad network access**: Capabilities are available over the network and accessed through standard mechanisms by heterogeneous platforms.
3. **Resource pooling**: The provider's computing resources are pooled to serve multiple consumers using a multi-tenant model, with different physical and virtual resources dynamically assigned and reassigned according to consumer demand.
4. **Rapid elasticity**: Capabilities can be rapidly and elastically provisioned to quickly scale out and rapidly released to quickly scale in. To the consumer, the capabilities available for provisioning appear to be unlimited and can be purchased in any quantity at any time.
5. **Measured Service:** Cloud systems automatically control and optimize resource use by leveraging a metering capability. Resource usage can be monitored, controlled, and reported.

Any provision of a computing service that has these five characteristics is an example of Cloud Computing, be it within a small organisational data-center or ranging across several continents and aggregating 500,000[2] computers like Amazon.

As a model for delivering computing services Cloud computing holds significant benefits for both suppliers and consumers. The "**On demand"** and "self provisioning" characteristics mean that, as organisations need computing resources, they can be commissioned while **rapid elasticity** means not only as that demand grows so can the amount of resources a user can commission but more importantly when they are not needed they can be decommissioned or released. This is important when tied to the fact that the use of cloud computing resources are

[1] NIST SP 800-145, *The NIST Cloud Computing Definition*
[2] http://huanliu.wordpress.com/2012/03/13/amazon-data-center-size/

measured and **metered** meaning that users pay only for the resources they actually use. For an organisations these characteristics allow a major shift from the **Capital Expenditure Model** where expensive computing infrastructure investments are made "upfront" and amortised over several years, to the **Operational Expenditures Model** where organisations can pay for resources as they are used and amortise expenses in the current operating year. This is important because it reduces the amount of capital needed to satisfy peak computing needs and frees up capital for other activities. According to IDC, [3] 81% of 479 organisations currently using cloud technologies in 2012 claimed that this is a major reason for their choice to use adopt those technologies. Very closely related is the fact that Organisations embracing cloud require less data center space, reduced data center staffing needs and spend less in the associated utilities required to power and cool equipment. From the provider point of view the same characteristics of Cloud Computing are attractive. Resource pooling means many applications can be virtualised and exist in a multitenant enviroment on a reduced set of machines, again with reduced capital investments. It means providers are not obliged to keep specific machines for specific clients active and can plan deployment and carry out maintenance of equipment in an optimised manner, providing the same amount of computing resources as in ASP models at a fraction of the cost. Providers can also provide more detailed Service Level Agreements and improved integrated Security with their clients and provide tools to monitor. Where technology progresses the characteristics of cloud that provide rapid elasticity mean substitution of new APIs and more modern equipment can be seamlessly implemented. These characteristics translate into better user satisfaction, better customer lock in and increased profits.

Cloud Computing – Benefits and Challenges

Benefits

Cost Efficiency: Traditionally, companies and governing bodies invest a lot in building their IT infrastructure paying for servers, software and multiple users' license fees. Their capital cost is getting even greater for maintaining and upgrading this infrastructure. While these investments are capital intensive, they are usually under-utilized most of the time and become obsolete as technological evolution takes place. At the same time, Cloud computing offers services at much cheaper rates and it eliminates the need for investing on maintenance and upgrades since these are responsibilities of the Cloud Provider. Moreover, Cloud's charging models (as 'pay-as-you-go') are flexible and enable Cloud Consumers to purchase the exact required resources at any given instance.

Quick Deployment and Scalability: Once a company decides to operate based on Cloud computing, it can deploy its system and make it fully functional within minutes. This is not the case when the choice of building stand-alone IT infrastructure is made. Additionally, with Cloud a company can start with a deployment as small as it is required given its current needs in terms of computing power and storage. When needed, the deployment can scale up to meet greater demands and scale back down when these are lower.

[3] IDC 2012 Quantitative Estimates of the Demand for Cloud Computing in Europe and the Likely Barriers to Up-take, SMART 2011/0045

Reliability: Usually Cloud providers use redundant data centers in multiple sites and thus they can recover from a local disaster providing reliability and Quality of Service to their clients. Small companies and government bodies usually cannot build their own multiple redundant data centers to achieve such reliability.

Backup and Recovery: Companies which rely on Cloud based services do not need to devise complex disaster recovery plans. It is the Cloud providers who are taking care of these issues and have the expertise to react faster. Moreover, since all data are stored in the Cloud, it is much easier to back them up and restore them than using physical devices for this task.

Globalized Access and Collaboration Facilitation: People worldwide can access the Cloud provided that they have an Internet connection. Consequently a company using Cloud can utilize workforce from all around the world. Additionally, Cloud services facilitate collaboration allowing employees, independently from their location, to work simultaneously on shared documents and applications. Finally the access to the Cloud usually is permitted by a plethora of electronic devices having access to the Internet including mobile devices (smartphones and tablets) allowing people to work even in remote places.

Competitiveness: Cloud provides SMEs with enterprise-class IT services and infrastructures which otherwise would be inaccessible. This gives them a competitive advantage over their big, established competitors. Additionally, since they do not have to worry about purchasing and maintaining state-of the art IT infrastructure, they can focus on innovation.

Easiness to learn: A company's employees tend to learn faster to work with Cloud applications since usually they are already familiar with them as they use them in their private lives. This is the case for popular Cloud applications like GMail and Google Drive.

Environmentally friendly: Since the Cloud resources are shared between cloud consumers and used by them only when there is the need, their utilization is more efficient and less energy is spent while they are idle. This is not the case for traditional IT infrastructures which cannot scale-down when they are not used.

Challenges

Security and Privacy: Security and Privacy are probably the two biggest concerns regarding Cloud computing and at the same time two of the more important reasons for slowing down cloud computing adoption. Since a company's or a government body's data are stored in the Cloud they are scattered in various sites. Essentially the Cloud consumer gives away data and information, which may be private, sensitive and confidential. The Cloud provider is responsible for maintaining and protecting them and thus has to be highly reliable. On top of that, in the case of government organizations, storing their data outside their national borders is usually not allowed by law. As a countermeasure a hybrid Cloud can be used where sensitive data are stored in the Cloud consumer's own data center and its access on the Cloud is permitted. Naturally security mechanisms between the Cloud Provider and the Cloud Consumer must be robust and carefully designed.

Technical Issues: While Cloud infrastructures are usually managed by well established IT companies, dysfunctions like outage and downtime may occur. On top of that, since an Internet connection is a precondition for accessing Cloud services, networking and connectivity problems may make these services unavailable.

Vulnerability: Since information is stored in the Cloud and its exchange is done over the Internet, the Cloud based services are exposed and as a result are prone to external threats like malicious users and hackers.

Lack of Standards and Provider Lock-in: While Clouds have well-documented proprietary application programming interfaces (APIs), there are not yet standards for them and thus different providers Clouds are usually non-interoperable. As a result the transition from one cloud provider to another is extremely complicated and expensive.

Service Models

Cloud Computing is traditionally offered according to three different usage models Software as a Service (SaaS), Platform as a Service (PaaS) and Infrastructure as a Service (IaaS). There are new service models for specific IT services like billing, security or network functionalities, but they are beyond the scope of this brief introduction.

SaaS is the model where the Cloud provider gives access to running applications already deployed on a provider's cloud infrastructure. Applications are typically available through web-based interfaces such as a Web browser on client devices ranging from PCs, tablets, smartphones and other portable devices. The user cannot influence the underlying network, servers, operating systems or storage and in most cases has no (or limited) control over the application itself. Salesforce[4] or Workaday[5] are good examples of widely used SaaS applications.

Example: SaaS – Google Apps

Google Apps[6] is a cloud-based suite which offers a variety of tools for different users and entities to connect and collaborate from any place, anytime and on any device. It offers solutions for business, education, non-profit and government. Especially for the e-government sector, it offers publishable shared calendars with smart scheduling features, a single point of access to and sharing of files with editors allowing for simultaneous document editing by different users, customizable policies for access to the applications features and easy and fast web site building, among others. Moreover, the Apps Marketplace[7] provides a series of tools for accounting and finance, project management, customer management, among others.

PaaS is the model where the provider allows users to deploy their own applications onto a cloud infrastructure using programming languages and tools developed/supported by the provider. The consumer does not manage or control the underlying network, servers, operating systems, or storage, but can control over the applications themselves and in some cases the application environments. Google App Engine[8], Force[9] or Amazon Web Services[10] are the market leading PaaS examples today.

[4] http://www.salesforce.com/
[5] http://www.workday.com/
[6] http://www.google.com/enterprise/
[7] http://www.google.com/enterprise/marketplace/
[8] https://developers.google.com/appengine/
[9] http://www.force.com/
[10] http://aws.amazon.com/

Example: PaaS –Windows Azure

Windows Azure[11] is provided by Microsoft and comprises a quite popular PaaS solution. It provides an on-demand platform with computation and storage capabilities, support of a wide variety of languages and environments (both Microsoft and third party ones) for developers to build, enhance, deploy and manage applications across Microsoft-managed datacenters. Moreover, it offers multiple data management services allowing storage at relational SQL databases, NoSQL table stores, and unstructured blob stores. The platform is further enhanced with backend capabilities for mobile applications, business analytics options (including hadoop and SQL reporting), media services as well as caching, messaging and integration mechanisms. Its customer base includes industrial, business, educational and public sector. Examples of governmental use of the Windows Azure is the Recovery and Reconstruction Support Program Database[12] launched by the Japanese Ministry of Economy, Trade, and industry for offering access to approximately 500 support programs to disaster victims (both citizens and companies), local government employees and others.

IaaS is the model where the provider delivers processing, storage, networks, and other fundamental computing resources directly to the user who deploys and runs heir own software stack including operating systems and applications. The consumer does not manage or control the underlying cloud infrastructure but has control over the operating systems, storage, deployed applications, and possibly limited control of some networking components like firewalls or load balancers.

Example: IaaS - Amazon EC2

Amazon EC2[13] provides a virtual computing environment, which allows use and management of a variety of computational and storage resources as well as operating systems through web service interfaces. It offers the flexibility of choosing among preconfigured template images or creating new ones fitting the needs of the customer as well as elasticity in terms of capacity adjustments depending on current or expected needs. Instances can be launched in separate Availability Zones in different geographic areas and/or countries, allowing for protection of the customer's applications from single location failures. The Service Level Agreement (SLA) commitment for availability reaches 99.95% for each Region. Instances can be created and used either on-demand (mainly for short term) or through reservation (for long term) and can be adjusted through users' bidding for unused EC2 capacity.

Especially for the eGovernment sector and taking into consideration its specific regulatory and compliance requirements in the U.S.A., a Region[14] (the AWS GovCloud (US) Region) has been designed to allow U.S. government agencies and customers to use the Cloud for more sensitive workloads.

Architecture Components

The Cloud Computing reference architecture[15] described below is intentionally generic. It does not describe any particular implementation model but is rather intended to show what actors exist and what kind of activities they are performing.

[11] http://www.windowsazure.com/
[12] http://www.microsoft.com/casestudies/Windows-Azure/Japanese-Ministry-of-Economy-Trade-and-Industry/Japanese-Earthquake-Recovery-Website-Keeps-Citizens-Informed-About-Vital-Programs/710000001712
[13] http://aws.amazon.com/ec2/
[14] http://aws.amazon.com/govcloud-us/

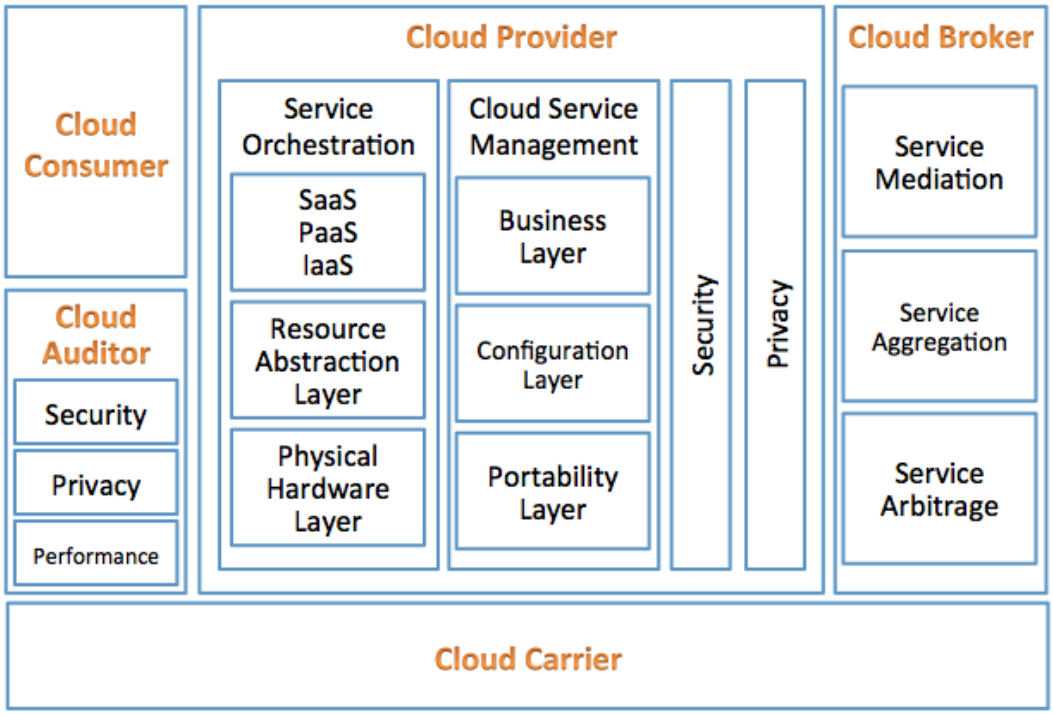

Cloud Consumers

A Cloud Consumer is the person that uses service from Cloud Providers. A cloud consumer contacts the cloud provider examines and chooses the services provided by the provider and stipulates service level agreements with the provider.

Consumers of SaaS can be organizations that provide their members with access to software applications, end users who directly use software applications, or software application administrators who configure applications for end users. SaaS consumers can be billed based on the number of end users, the time of use, the network bandwidth consumed, the amount of data stored or duration of stored data. Cloud consumers of PaaS can employ the tools and execution resources provided by cloud providers to develop, test, deploy and manage the applications hosted in a cloud environment.

PaaS consumers can be application developers who design and implement application software, application testers who run and test applications in cloud-based environments, application deployers who publish applications into the cloud, and application administrators who configure and monitor application performance on a platform. PaaS consumers can be billed according to, processing, database storage and network resources consumed by the PaaS application, and the duration of the platform usage. Consumers of IaaS have access to virtual computers, network-accessible storage, network infrastructure components, and other fundamental computing resources on which they can deploy and run arbitrary software.

The consumers of IaaS can be system developers, system administrators and IT managers who are interested in creating, installing, managing and monitoring services for IT infrastructure operations. IaaS consumers are provisioned with the capabilities to access these computing resources, and are billed

[15] This Architecture and the definition of the actors is “loosely” based on the National Institute of Standards and Technology’s organisational framework regarding cloud computing

according to the amount or duration of the resources consumed, such as CPU hours used by virtual computers, volume and duration of data stored, network bandwidth consumed, number of IP addresses used for certain intervals.

Cloud Consumers in Public Administrations

Across Europe administrations are already using cloud across at all levels of government both with and without a central government strategy. SaaS, PaaS and IaaS cloud applications, services and infrastructure already being used in government to manage their administrations core business processes, to carry out office productivity, to manage social media and citizen relationship services as well as to control core IT services. We see different roles and usage models

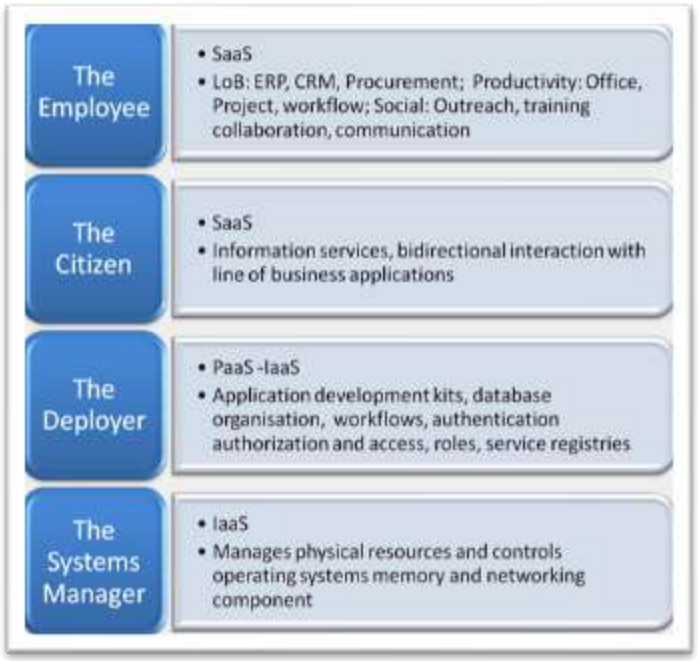

The Government Employee - *SaaS Usage*

In government we find general "line of business" employees using a wide variety of applications according to the SaaS model. SaaS applications include asset management, business intelligence, ERP, CRM, GIS, PR and marketing applications as well as others. Productivity applications include office tools, document management, project and planning suites or workflow applications. Social Media applications range from public outreach, to citizen engagement, personnel recruitment, training, work group collaboration and brainstorming.

The Citizen - *SaaS Usage*

We consider the citizen as a government user, because s/he is the consumer of many services and the advocate of one of the demanding "anywhere anytime" access requirement for cloud computing. Starting from the definition of cloud computing we will recall that the essential characteristics of a cloud deployment included on-demand self-service, broad network access, resource pooling, rapid elasticity and measured service. The citizen exerts a stress test of the "broad network access" category, interacting with the service/infrastructure from a number of different client devices (computer, smart phone, tablet, etc.). Citizens are using a number of information services but also having bidirectional interaction with line of business applications (income tax declarations, school reservation, etc.) and have persistence of their personal information and records in the cloud system.

The Deployer - *Paas Usage*

The government developer is focused on preparing the applications to be deployed for the other users. He is testing integrating and deploying applications. The Developer does not manage or control the underlying cloud network, servers, operating systems, or storage, but has control over the deployed applications and possibly application hosting environment configurations. The typical government Deployer is creating

application development routines, templates and development kits, working on Data structures, and database organisation, workflows, security services like authentication authorization and access, roles, service registries and similar “platform” activities.

The System Manager - *IaaS Usage*

The systems manager provides the government organisation computational processing powers, RAM, instance storage/disk space networks, and the other computing resources that are needed. The Systems Manager where the consumer is able to deploy and run arbitrary software, which can include operating systems and applications. The Systems Manager can manage or control the underlying cloud infrastructure and has control over operating systems, storage, deployed applications, and control of select networking components like gateways and firewalls.

Cloud Provider

A Cloud Provider is an organisation, which makes services available to Consumers, Brokers and Auditors. They manage the infrastructure required to provide services, run the cloud software that provides the services, and deliver the services through the network.

The SaaS provider deploys, configures, maintains and updates the operation of the software applications on a cloud infrastructure so that the services are provisioned at the expected service levels to cloud consumers. The SaaS provider has the responsibility to manage and control the applications and infrastructure. Consumers on the other hand have limited influence on the application enviroment and configuration.

The PaaS provider manages the computing infrastructure and runs execution stack, runtime software database, and other middleware software that provides the components of the platform, such as components. The PaaS Cloud Provider normally either provides its own development, deployment and management software deployment and development kits and management tools. The PaaS Cloud Consumer has control over the applications and possibly some the hosting environment settings, but has no or limited access to the infrastructure underlying the platform such as network, servers, operating systems or storage.

The IaaS provider owns and provides the physical computing resources including servers, networks and storage. The Cloud Provider maintains the service interfaces, virtual machines and virtual network interface software necessary to make computing resources available to the IaaS Consumer. The IaaS Consumer has access to the operating system and network. The IaaS Cloud Provider, on the other hand, has control over the physical hardware and cloud software that makes the provisioning of these infrastructure services possible, for example, the physical servers, network equipment, storage devices, host OS and hypervisors for virtualization.

A Cloud Auditor is any independent party that conducts assessment of cloud services, information system operations, performance and security of cloud provider’s implementations verifying conformance to standards, SLAs or legal obligations. A cloud auditor can evaluate the services provided by a cloud provider in terms of security controls, privacy impact, performance, system integrity, and/or availability of the system and its information. Audits are usually performed by examining activity logs but in some provider environments real time evaluation interfaces are made available.

A Cloud Broker is any entity that manages the use, performance and delivery of cloud services, and negotiates relationships between Cloud Providers and Cloud Consumers. A cloud broker is an entity that manages the use, performance and delivery of cloud services and negotiates relationships between cloud providers and cloud consumers. Cloud Brokers may combine a set of services from one or more cloud providers (*Service Aggregation)* or repackage/rebrand cloud services (*Service Intermediation)* so as to make new service offerings. These are typically a fixed set of services while an entity that mixes and matches services at or near run-time choosing the most available, efficient or cost effective services from multiple cloud providers is engaging in *Service Arbitrage.*

A Cloud Carrier is any intermediary that provides connectivity and transport of cloud services from Cloud Providers to Cloud Consumers.

Deployment models

Cloud computing is offered in different forms: public clouds, community clouds, private clouds, and hybrid clouds, which combine both public and private.

A Public cloud infrastructure is available for open use by the general public. It may be owned, managed, and operated by a business, academic, or government organization, or some combination of them. It exists on the premises of the cloud provider. It is connected to the Internet through broad network access. Users connect through the public Internet using any of a number of protocols

Private cloud. The cloud infrastructure, or a dedicated portion of it, is provisioned for exclusive use by a single organization comprising multiple consumers (e.g., business units). It may be owned, managed, and operated by the organization, a third party. The private cloud is reachable as a LAN extension to the servers in enterprise A's data center. How is this reachability realized? A secure Virtual Private Network (VPN) tunnel is first established between the enterprise data center and the public cloud. This tunnel uses public IP addresses to establish the site-to-site VPN connection. The VPN gateway on the cloud service provider side uses multiple contexts—each context corresponding to a specific private cloud. Traffic from enterprise A is decrypted and forwarded over to an Ethernet switch to the private cloud for enterprise A. A server on enterprise A's internal data center sees a server on private cloud A to be on the same network or some combination of them, and it may exist on or off premises.

Community cloud. The cloud infrastructure is provisioned for exclusive use by a specific community of consumers from organizations that have shared concerns (e.g., mission, security requirements, policy, and compliance considerations). It may be owned, managed, and operated by one or more of the organizations in the community, a third party, or some combination of them, and it may exist on or off premises.

Hybrid cloud. The cloud infrastructure is a composition of two or more distinct cloud infrastructures (private, community, or public) that remain unique entities, but are bound together by standardized or proprietary technology that enables data and application portability. The Cloud infrastructure is planned so that some activities like storage or executing complex algorithms can be carried out on supplementary independent clouds where the specific needs arise or where load exceeds the capacity of the cloud infrastructure (cloud bursting).

Whatever the configuration, it is evident that no longer buying and maintaining infrastructure on their own by adopting a Cloud strategy can completely change the way governments and their agencies are using technology to service citizens, other levels of government, and their suppliers. Across Europe Authorities are already leveraging IT cloud resources supplied by companies like Google or Amazon are already supplying IT resources in the cloud, eliminating many of the complex constraints from the traditional computing environment, including space, time, power, and cost.

1 References

Cooper, B. F., Silberstein, A., Tam, E., Ramakrishnan, R., & Sears, R. (2010). Benchmarking Cloud Serving Systems with YCSB. Santa Clara, CA, USA: Yahoo! Research.

Dumitras, T., & Shou, D. (2011). Toward a Standard Benchmark for Computer Security Research. Carnegie Mellon University.

Mell, P., & Grance, T. (2011, Sept). NIST CSRC Special Publications. Retrieved from NIST Computer Security Division: Computer Security Resource Center: http://csrc.nist.gov/publications/nistpubs/800-145/SP800-145.pdf

IBM Institute for Business Value, The power of cloud, Driving business model innovation, 2012. Available at: http://www.ibm.com/cloud-computing/us/en/assets/power-of-cloud-for-bus-model-innovation.pdf

Moore Stephens International, The benefits and challenges of cloud computing, 2013. Available at: http://www.moorestephens.com/cloud_computing_benefits_challenges.aspx

CIO magazine, 2011 Cloud Computing Survey, 2011. Available at: http://mkting.cio.com/pdf/CIOCloudSummary.pdf

techsoupglobal.org, 2012 Global Cloud Computing Survey Results, 2012. Available at: http://www.techsoupglobal.org/sites/default/files/TechSoup_Global_Cloud_Report_Executive_Summary.pdf