Cloud computing

Its development and security

J. Craig Mudge
CSIRO

Cloud computing is a mode of acquisition of infrastructure where a cloud service provides on demand computing and on-demand storage, accessible over the web, with a matching business model, namely pay-per-use. Lower cost, sometimes a factor of five, is achieved by automating operations in warehouse-sized data centres, sharing the hefty burden of cooling, physical security, and power backup, and the use of commodity components. As access to storage managed by cloud service providers gained market acceptance, the phrase "the cloud" came into use to refer to the location of information stored with a cloud service provider to be shared (business documents and photos) or copied for backup (against losing a phone or deleting information on a laptop). Security, and to a lesser extent privacy, concerns are the major barriers to the adoption of cloud computing, but the leading cloud service providers have responded to the point that sensitive information, such as medical records and credit-card-holder data, is now held in the cloud.

Introduction

The term cloud has come to have several meanings. At first, around 2007 in Australia, cloud computing referred to on-demand access to massive computing facilities owned by a cloud service provider and shared by all-comers. A few years later, the phrase "the cloud" came into use to refer to the location of information stored with a cloud service provider and accessed via a web browser – to be shared (business documents and photos), copied for backup (against losing a phone or deleting information on a laptop), or as part of an e-commerce distribution system (for example, eBooks and music).

The first meaning of the term cloud concerns procurement of information technology (IT) infrastructure, where cloud computing provides on-demand computing and on-demand storage, accessible over the web, with a matching business model, namely pay-per-use. Warehouse-sized data centres owned and operated by cloud service providers hold tens of thousands of PCs and benefit from economies of scale. The lower cost, sometimes a factor of five, of these cloud computing systems is achieved by automating operations, sharing the hefty burden of cooling, physical security, and power backup, and the use of commodity components (cheap disks, cheap microprocessor chips, mother boards, and Ethernet interconnect). Using commodity parts means that fault tolerance is essential, and this is done by run-time software hidden from users.

 http://doi.org/10.7790/ajtde.v1n1.7

Security concerns, and to a lesser extent privacy, are the major barriers to the adoption of cloud computing. Because security has always been an important concern for information managers, the protection of personal, scientific, commercial and intellectual property information must be assured by cloud service providers. The leading firms in the IT industry have responded to this need; their security policies and practices are strong and clearly communicated in their communication with customers. As cloud usage has become more pervasive, industry leaders have encouraged improved security practices by their customers, e.g., more use of encryption for data at rest and in transit, as well as making stronger practices easier to adopt, two-factor authentication being a good example.

Different deployment types of cloud computing

Cloud computing is a way in which a company or government can obtain its IT infrastructure (computer servers, storage, networks, and applications) as a service, obtaining those resources as a utility, just as we acquire electricity and water at home. A matching business model, pay-per- use, means that entities can replace their capital expenditures with operating expenditure. Self-service and payment by credit card further simplify ease of use and cost.

The US National Institute of Standards and Technology (NIST), identifies four types of cloud services:

1. Public cloud
2. Private cloud
3. Community cloud
4. Hybrid

The major providers of public clouds are Amazon Web Services, Google, and Microsoft Azure, with data centres of massive scale networked across the globe.

A private cloud delivers IT services to a single enterprise, generally across multiple lines of business, and is protected by being within corporate firewalls, and other security measures deployed in well-managed IT infrastructure.

A community cloud services the needs of multiple entities in a community that have common IT needs, for example, departments in a university. A contemporary example is NECTAR, the most recent Australian government eResearch infrastructure for Australian universities and publicly funded research institutions.

A hybrid approach allows a company to split its resources between a public cloud service and in-house infrastructure. In that case a company gets some of the cost benefits of scale but peace of mind from keeping its most sensitive data on its own premises.

 http://doi.org/10.7790/ajtde.v1n1.7

Service models

Three models, which vary by abstraction level, are used.

Software as a Service (SaaS). The capability provided to a consumer is access to a provider's applications running on a cloud infrastructure. The consumer does not manage or control the underlying cloud infrastructure of network, servers, operating systems, or storage. Salesforce.com, a customer resource management (CRM) function to manage a firm's communications with its customers, was an early example of SaaS. Since a consumer does not install the application software on her own machine, she does not have to manage upgrades and other administrative matters. Xero is a recent accounting software capability that is delivered in the SaaS model and includes cloud-based collaboration between a small firm and its accountant. Google Docs is Google's SaaS office suite.

Infrastructure as a Service (IaaS). The capability provided to a consumer is to provision processing, storage, networks, and other fundamental computing resources where the consumer is able to deploy and run arbitrary software, which can include operating systems and applications. The IaaS offerings from the main public cloud service providers are Google Compute Engine, Amazon Web Services (AWS), and Windows Azure.

Platform as a Service (PaaS). When the application deployed on the cloud infrastructure has been created using programming languages, libraries, services, and tools proscribed and supported by the provider, the service model is called PaaS. The consumer does not manage or control the underlying cloud infrastructure, but has control over the deployed applications, generally including configuration settings for the application-hosting environment. Windows Azure was Microsoft's first cloud computing offering and was delivered as PaaS. However, when some users found they needed access to lower level facilities in the cloud infrastructure as well as access to the Linux operating system, more flexibility had to be provided. This was done by adding IaaS to Microsoft's PaaS model.

Government clouds

Because of the special nature of government data, so-called government clouds have been built. The UK Government G-Cloud is an initiative targeted at easing the procurement by government agencies of commodity IT infrastructure. In the US, the CIA has let a government cloud contract for a private cloud housed on government premises, built to the specifications of the government agency, and operated by a public cloud service provider.

 http://doi.org/10.7790/ajtde.v1n1.7

Cloud bursting

Cloud bursting, a recently introduced term refers to the use of a public cloud service to handle overflow compute and storage needs during periods of a user's peak activity, such as end-of-quarter accounts processing in business. Data migration of a database to a newer system, or aggregated system, is another example, with the extract-transform, and load being a task suited to cloud bursting. In science an example would be image processing of photographic data after an annual flyover.

A disadvantage

Since cloud computing requires Web connectivity, if there is no Web access, or the Internet is down temporarily, a user has no access to her computation and storage facilities. However, Internet availability and reliability continues to improve. Internet access on newer aircraft allows cloud use by airline travellers.

A parallel programming framework

Parallel programming, and associated parallel execution of thousands of tasks, is an essential property of cloud computing at massive scale, since it enables certain functions core to web businesses, such as web search and spam detection in email, to be carried out efficiently. As a result, we now turn to a discussion of its mechanisms and benefits.

The parallel programming framework MapReduce, introduced by Google in 2004, provides a parallel programming style matched to many of the algorithms of interest in data intensive applications. It has an associated distributed file system Google File System (GFS), which also keeps multiple copies of data for redundancy. An open source implementation, Hadoop and Hadoop Distributed File System (HDFS), has been in development and production use for several years in Yahoo, IBM, and Facebook, for example, and is now quite stable and of enterprise grade. MapReduce/Hadoop also incorporate the run-time routines, referred to previously, which build reliable systems built from cheap components.

Parallel programming has become important for two reasons. First, those problems which are able to be decomposed into sub-parts for independent execution will take less time if they are executed in parallel. Second, microchip advances have reached a power limit. Although Moore's Law continues to hold, i.e., the number of transistors fabricated on a die continues to double every eighteen months, a power ceiling has been reached , which precludes increasing the clock frequency. As a result, the increase in transistor density at each generation is used to fabricate multiple processors (multi-cores) on each die.

MapReduce/Hadoop manages the parallel execution of those problems that require very little communication between their tasks executing in parallel. When there is zero communication between tasks, the problem is called "embarrassingly parallel", which is the case in much genomic analysis, divide and recombine approaches to processing large spatial data sets, Monte Carlo simulation, some machine-learning methods, and parameter sweeps in some numerical computational methods. Most computations run by Web companies are embarrassingly parallel and so are allocated across processors and cores. Indexing the web, search, ad placement, and analysing click streams and web logs all fall into this class, and are allocated across tens of thousands of machines in a data centre.

The task of security management in 2013

Businesses and individuals now have less control over security

Traditional security used to be the responsibility of a corporation's security function at work and a user's responsibility at home. We managed our own networks and computers, we installed our own firewalls and anti-virus software and we purchased products, both hardware and software, and thus we owned them. Two trends are changing this. The first is the cloud. Now our data, including email, calendars, presentations, and LinkedIn posts are on someone else's systems. The second trend is the new vendor-controlled hardware platforms we like so much, such as Kindle, iPad, and Android. We have very little control over these platforms. Taking the two trends together, we have less and less control over our security.

The scope of information security management

The scope of information security management is quite broad, covering both mechanisms to protect data and compliance with regulations laid down for companies and other entities, whether operating with financial, health, scientific, or operational data. The regulations and guidelines require adherence to the ISO 27000 standard. The scope of information security management includes any aspect of an institution, whether physical premises, physical equipment, software firewalls, web sites, outsourced data centres, or procedures which interact with data relating to a firm's customers, operations, intellectual property, or a firm's employees, and the business continuity of an entity. Such interactions with data includes its storage, computation, and transmission such that the data is not lost, changed, or disclosed to unauthorized persons or processes, whether in normal operations or abnormal.

 http://doi.org/10.7790/ajtde.v1n1.7

Blurring of private and work use of IT

Since the early 2000s, as mobile telephones became richer in function and as they became increasingly used for recreation and entertainment, there has been an increasing blurring of private and work use of information technology. This has made security management more complex.

A step increase in diversity of access and source of software

In a recent report from its newly launched Digital Productivity and Services Flagship, CSIRO provides the following information on smart devices and downloads:

> *40 billion apps have been downloaded to portable devices worldwide; 20 billion of these in 2012 alone. 3.65 million Australians use multiple smart devices to go online and use apps. 8.67 million – about half of adult Australians now own a smart device; by 2020 the average person will own six smart devices. $1.2 billion will be spent globally on smart devices in 2013.* (CSIRO 2013)

Security in public cloud computing

Data centre security and availability

Security begins with the physical security of data centres, the policy regarding personnel access, and the backup systems in place to allow continued operation in the face of failure of subsystems. The Telecommunications Industry Association, accredited by ANSI (American National Standards Institute), published an infrastructure standard for data centres, which defined four levels, called tiers. The simplest is a Tier 1 data centre, which is basically a server room, following basic guidelines for the installation of computer systems. The most strict level is a Tier 4 data centre, which is designed to host mission critical computer systems, with fully redundant subsystems, compartmentalised security zones controlled by biometric access, and independently dual-powered cooling equipment.

Security practices of major providers

Because the major cloud service providers publish their security and privacy policies and processes, users can have confidence in how their data is stored and cared for. [i] A useful exercise for users is to compare the security policies of their own internal IT function with those published by the cloud service providers.

Two-factor authentication

Two-factor authentication is an approach to authentication which requires the presentation of two authentication factors – a knowledge factor ("something the user knows") and a possession factor ("something the user has") – in order to achieve stronger authentication.

 http://doi.org/10.7790/ajtde.v1n1.7

The first factor, the knowledge factor, is the conventional password. The second factor can take several forms. Increasingly it involves a user's mobile phone, when a text message is requested by some Internet banking systems, for example. Before mobile phones were in common use, a piece of hardware called a security token, was the common possession factor. A personal token generates a one-time password which is valid for a short period of time, generally 30 seconds, long enough for a login to occur.

A widely used cloud storage and sharing system from Dropbox offers two-factor authentication, which is called "two-step verification for logging into an account". The possession factor uses a mobile phone.

Encryption

The major cloud service providers offer encryption for data at rest in their storage infrastructure and for data in transit, which is secured by Secure Sockets Layer (SSL) and AES 256-bit encryption.

(Distributed) Denial of Service (DDOS) attacks

DDOS are used in extortion attempts and in politically motivated attempts to damage an entity's web presence.

Denial of Service attacks are characterised by, for example, an attacker who repeatedly attempts and then abandons network connections to the defender's resources, thus using up network bandwidth, computing cycles, memory, etc. Single-point denial of service attacks (coming from a single network node, for example) are relatively easy to defend against once detected. Distributed attacks, often coming from widely-distributed botnets, or orchestrated collections of nodes – often the computers of innocent users compromised by viruses – are much harder to protect against.

Because of the large scale of the networking and load-balancing infrastructure of the main cloud service providers, they are able to handle DDOS attacks better than small data centres can.

Confidence in cloud computing continues to increase

Highly sensitive data is now routinely stored in public clouds. Amazon stores medical records and is compliant with the main US regulation, viz., HIPAA. It also complies with the Payment Card Industry (PCI) standard, which pertains to businesses that handle credit-card-holder data. As the Chief Information Security Officer for Amazon, Stephen Schmidt, noted when discussing the credit-card-based ecommerce business of Amazon: "*For some industries it is an absolute must-have. For instance, for Amazon.com to move onto AWS we had to be PCI compliant, because of the credit card transaction volumes.*"

Since patient health data is one of the most sensitive types of data, pharmaceutical companies have traditionally been reluctant to outsource their IT operations. However, Bristol Myers Squib is now sufficiently confident in the security of Amazon Web Services that they use complex patient data to design new drug trials with cloud computing.

Overall, in the several years of experience with public cloud computing, down time, not data breaches, have been the most significant security lapses.

Security and privacy regulations and guidelines

The Australian Prudential Regulation Authority (APRA) www.apra.gov.au is the prudential regulator of the Australian financial services industry. It oversees banks, credit unions, building societies, general insurance and reinsurance companies, life insurance, friendly societies, and most members of the superannuation industry. APRA's standards address a large number of matters including operational risk, outsourcing, and business continuity. [ii]

The Office of the Information Commissioner http://www.oaic.gov.au, formerly the Privacy Commissioner, is concerned with a firm's obligations under the *Privacy Act 1988* (Cth) to put in place reasonable security safeguards and to take reasonable steps to protect the personal information held by a firm from loss and from unauthorised access, use, modification or disclosure, or other misuse. There are data breach provisions as well, such as the preparation of a data breach policy and response plan.

ISO 27001, a certification under the International Organization for Standardization, is a security management standard that specifies security management best practices and comprehensive security controls following the ISO 27002 best practice guidance. Certification in the standard requires a company to systematically evaluate all information security risks, design and implement a comprehensive suite of information security controls and other forms of risk management to address company and architecture security risks, and adopt an overarching management process to ensure that the information security controls meet information security needs on an ongoing basis

Two other security standards or guidelines are worth noting. The Australian Federal Government Information Security Manual applies to federal and state government operations and is also recommended for Australian businesses. A clear presentation of security measures is provided by the Payment Card Industry Data Security Standard. The twelve requirements of the standard are easily understood, common sense steps that mirror security best practice. Regular independent outside audits and penetration tests should be carried out against this standard.

 http://doi.org/10.7790/ajtde.v1n1.7

Other considerations

Ownership of data

Responsible cloud service providers will clearly state their position on ownership of the data. Most will say that ownership remains with the user.

Security and software development

Software development is of interest for two reasons.

a. Poorly written software can inject new security vulnerabilities, and
b. The testing environment itself can expose a company to data loss and/or data alteration if not managed properly.

As security threats have increased, some organisations have begun paying attention to security in their application design. Around 2004, organisations began introducing development processes to address software development issues, variously called secure coding programs, software assurance, app sec, and secure development life cycles (SDLs). The first major process was Microsoft's Security Development Lifecycle, initially implemented internally, but openly available since 2004. [iii]

BYOD

Bring your own device (BYOD) means the policy of permitting employees to bring personally owned mobile devices (laptops, tablets, and smart phones) to their workplace, and use those devices to access privileged company information and applications. The benefits of the practice are numerous and include work efficiency and flexibility, allowing employees to have a more personalised approach to their work, and a lowering of costs to the organisation. However, BYOD has complicated security management.

It is good practice for a firm to require that personal devices are subject to the same constraints as those provided by the firm, including the use of remote-wipe capability to be invoked when a device is lost or stolen.

Discussion and Conclusions

The new method of procurement of IT infrastructure, akin to the way one procures electricity and water from a utility, namely accessing an elastic supply on a pay-per-use basis has been transformational. Startups, along with new divisions of existing businesses, no longer spend scarce capital on IT equipment. Because computational and storage resources in the cloud are accessible over the web, from a common-or-garden web browser, and a rich set of web standards has come into being, a high degree of interoperability between products and

 http://doi.org/10.7790/ajtde.v1n1.7

services of IT vendors has been made possible. The degree of interoperability is of historic proportions.

This interoperability has also meant that those users who have multiple devices (lap top, smart phone, and carry along pad, for example) can keep copies of their information consistent across their devices. This property also allows easy sharing of information and has transformed document management and distribution as a result.

Because security is the major barrier to adoption, much attention has been given to it by cloud service providers and the IT staff of firms who choose to use the cloud. In choosing cloud service providers, the most important factor should be the posture and mechanisms regarding security. Is the policy comprehensive and clearly presented? If a user wishes to strengthen the level of security management, by using two-factor authentication and encryption, for example, is the adoption of the stronger facilities straightforward? Furthermore, does the approach taken by a provider fit with the risk management function of an organisation and the risk framework used by that function.

Where an industry has defined rules for handling sensitive data, e.g., credit-card-holder information, or medical records, and has set up procedures for certification against those rules, obtaining such certification is recommended since it improves practices and provides a level of comfort to customers.

The US *Patriot Act*, which came in to being after the September 2001 terrorist attacks, gave the US government legal access to data held by US corporations whether the data was located in the US or abroad. At first this caused alarm among would-be customers of cloud service providers operating in Australia as subsidiaries of US companies. However, once it was known that the Australian Federal Police has equivalent power in data access to Australian-owned firms, further public discussion on a potential asymmetry of powers of access began to subside.

A practical dimension of privacy in cloud computing is the uneasiness experienced by some users of cloud systems which operate on a business model depending on knowledge of users for targeting offers, especially advertisements. Such knowledge is built from prior user behaviour, location, expenditure, web search history, and personal data supplied during registration for a service. Web search was the first major information service offered on the web. Although offered without charge, a search service generates an advertisement related to a user and her search and places it close by the answer to her query. As web businesses developed, other information and collaborative functions were added, most notably email, calendars, text editors, and spread sheets. Most services are offered without charge, but, of course, give the cloud service provider access to more user data that can be mined for clues

and data to improve targeting of offers. The commercial successes of web businesses show that users are willing to give access to their web behaviour data in return for services.

Bibliography

Australian Government Information Security Manual (ISM). Defence Signals Directorate (DSD). Available from: http://www.dsd.gov.au/infosec/ism/

Amazon Web Services. AWS Security Center. Available from: https://aws.amazon.com/security/

ATSE. 2010. "Cloud computing: Opportunities and challenges for Australia", J Craig Mudge, Principal Author. Report of a Study by the Australian Academy of Technological Sciences and Engineering (ATSE), 2010, ISBN 978 1 921388 15 6 Available from: http://www.atse.org.au/resource-centre/ATSE-Reports/Information-Technology/

CSIRO. 2013. Entering the App age. CSIRO Snapshot. Issue 22 / 1 May 2013

Dean, J. and Ghemawat, S. 2004. "MapReduce: Simplified Data Processing on Large Clusters". Jeffrey Dean and Sanjay Ghemawat. OSDI '04: 6th Symposium on Operating Systems Design and Implementation.

Defence Signals Directorate. "Cloud Computing Security Considerations". Available from: http://www.dsd.gov.au/infosec/cloudsecurity.htm

European Network and Information Security Agency. Procure Secure. A guide to monitoring of security service levels in cloud contracts. Available from: www.enisa.europa.eu/activities/Resilience-and-CIIP/cloud-computing/procure-secure-a-guide-to-monitoring-of-security-service-levels-in-cloud-contracts/at_download/fullReport

Google. Security and privacy overview. Google. Available from: http://support.google.com/a/bin/answer.py?hl=en&answer=60762

McAfee, A. and Brynjolfsson, E. 2012. "Big Data – the management revolution", *Harvard Business Review*, Oct 2012

Mell, P. and Grance, T. 2011. The NIST Definition of Cloud Computing. National Institute of Standards and Technology. Gaithersburg, MD. September 2011. Available from: http://csrc.nist.gov/publications/nistpubs/800-145/SP800-145.pdf

Microsoft, 2013. "Technical Overview of the Security Features in the Windows Azure Platform." Available from: http://www.windowsazure.com/en-us/support/legal/security-overview/

National eResearch Collaboration Tools and Resources project (NeCTAR) . Available from: nectar.org.au/

NIST Guidelines on Security and Privacy in Public Cloud Computing (SP SP800-144)

Payment Card Industry Data Security Standard. Available from: https://www.pcisecuritystandards.org/security_standards/index.php

Endnotes

[i] The relevant documentation for Amazon, Google, and Microsoft is to be found at

Amazon Web Services. AWS Security Center. https://aws.amazon.com/security/

Google. Security and privacy overview. http://support.google.com/a/bin/answer.py?hl=en&answer=60762

Microsoft, 2013. "Technical Overview of the Security Features in the Windows Azure Platform." http://www.windowsazure.com/en-us/support/legal/security-overview/

[ii] PPG 235 – Managing Data Risk 11 December 2012 is an APRA guide which seeks to assist the organisations APRA regulates in "managing data risk" by outlining what APRA regards to be sound information management practices.

[iii] Some ten or so years ago, Microsoft changed its investment priorities on security. It decided to move its emphasis from making its software products and services totally secure to training its internal developers to write secure code. It then made the training available to its customers.

Cite this article as: Mudge, J. Craig. 2013. 'Cloud computing. Its development and security'. *Australian Journal of Telecommunications and the Digital Economy* 1 (1): pp.71.1 – 7.12. DOI: 10.7790/ajtde.v1n1.7 Available from: http://telsoc.org/journal

 http://doi.org/10.7790/ajtde.v1n1.7