BLACK CIRCLE PTY LTD

CLOUD COMPUTING:

A BRIEFING FOR THE BUSINESS ANALYST

Table of Contents

1 Cloud Computing 2
1.1 What is Cloud Computing? 2
1.2 Benefits and Opportunities 5
1.3 Cost Considerations 8
1.4 Risks 11
1.5 Outlook 16
2 Conclusions 20
3 References 22
4 Appendix A: Concepts & Vocabulary 30
5 Appendix B: Security Considerations 37
6 Appendix C: Internet Security 39
7 Appendix D: Contracts and SLAs 43

ABN: 97 081 830 499

GPO Box 2785
Canberra ACT 2601

fax: +61 3 6257 2081
www.blackcircle.com.au

1 CLOUD COMPUTING

1.1 WHAT IS CLOUD COMPUTING?[1]

In common with the Australian Government (DOFD 2011) this discussion adopts the US Government's National Institute of Standards and Technology's definition of cloud computing. It is (Mell & Grance 2010) "a model for enabling convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction…[It] is composed of five essential characteristics:

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.

Broad network access: Capabilities are available over the network and accessed through standard mechanisms that promote use by heterogeneous thin or thick client platforms (e.g., mobile phones, laptops, and PDAs). (The idea of ubiquitous network access "does not necessarily mean Internet access. By definition, a private cloud is accessible only behind a firewall. Regardless of the type of network, access to the cloud is typically not limited to a particular type of client" (Cloud Computing Use Case Discussion Group 2010)) .

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. There is a sense of location independence in that the customer generally has no control or knowledge over the exact location of the provided resources but may be able to specify location at a higher level of abstraction (e.g., country, state, or datacenter). Examples of resources include storage, processing, memory, network bandwidth,

[1] A detailed glossary of terms associated with cloud-based computing is at **Appendix A: Concepts & Vocabulary**.

and virtual machines (VM). In many cases privacy laws and other regulations require the cloud provider's resources to be in a particular location. The cloud provider and the cloud consumer must work together to adhere to those regulations.

Rapid elasticity: Capabilities can be rapidly and elastically provisioned, in some cases automatically, to quickly scale out and rapidly released to quickly scale in. To the consumer, the capabilities available for provisioning often appear to be unlimited and can be purchased in any quantity at any time.

Measured Service: Cloud systems automatically control and optimize resource use by leveraging a metering capability at some level of abstraction appropriate to the type of service (e.g., storage, processing, bandwidth, and active user accounts). Resource usage can be monitored, controlled, and reported, providing transparency for both the provider and consumer of the utilized service."[2]

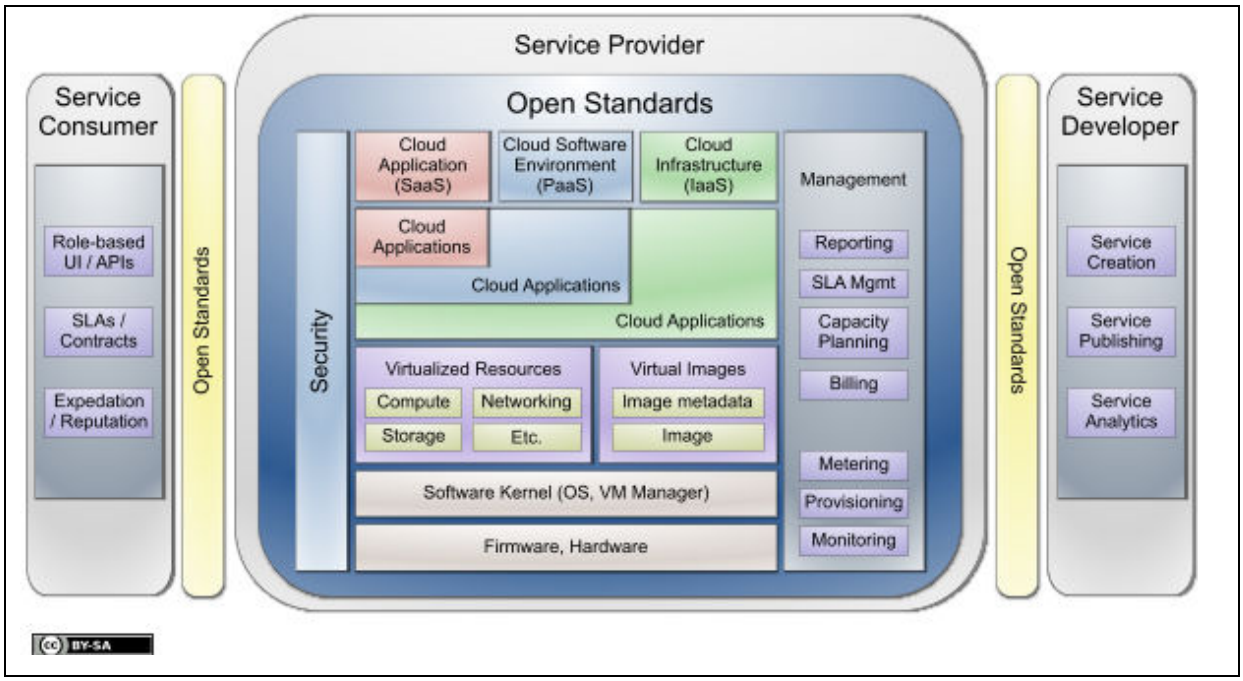

Figure 1: A model of Cloud Computing (Cloud Computing Use Case Discussion Group 2010)

[2] Other definitions are available. Those provided by Wang et al (2010, p. 3) and Grossman (2009, p. 23) are more or less equivalent and cover much the same ground. Others, such as Buyya et al (2009) are more technical in focus.

Figure 1 models elements of the cloud computing paradigm. What immediately becomes clear is that there is no prescription as to the nature of the players—consumer, provider or developer. As well, there are currently four deployment models recognised by the National Institute of Standards and Technology, USA (NIST) (Mell & Grance 2010): private, community, public and hybrid—this last comprises composites of the other three, and use " a superset of the technology used in a Community Cloud" (Cloud Computing Use Case Discussion Group 2010).

Even these models are not prescriptive. The service consumer may be an enterprise, an individual, or both (Cloud Computing Use Case Discussion Group 2010). The service provider may take many guises: for example, a private cloud "can be managed by a third party and can be physically located off premises. It is not necessarily managed and hosted by the organisation that uses it" (Cloud Computing Use Case Discussion Group 2010).

This all implies is that cloud computing is no one thing. Rather, it is a set of technologies and relationships that may be deployed in a wide variety of ways to meet different requirements. What then, is the touchstone? Certainly not resources—networks, servers, storage, applications, and services are all specifically mentioned in the NIST definition (Mell & Grance 2010) and it discusses three service models – in chronological order of evolution they are: *Infrastructure as a Service* (IaaS), such as Amazon Web Services, *Platform as a Service* (PaaS), such as GoogleApps, and *Software as a Service* (SaaS), such as Salesforce.com. Crucially, data has been conspicuous by its absence from the discussion so far. From this perspective, cloud computing may usefully be conceptualised as:

dynamic allocation of resources to manipulate data[3].

Contemplation of cloud computing accordingly requires a clear understanding of the problem to be resolved, so as to focus on configurations appropriate to manage the data relevant to the problem.

[3] Data here is taken in the widest sense, including text, numbers, images, video, sound.

1.2 Benefits and Opportunities

1.2.1 Cost Savings .

Cloud Computing may deliver significant cost savings derived from five areas:

- Savings may be made in the capital purchase or leasehold cost of ownership of hardware and software. Per transaction, the economies of scale available to the cloud provider can make it feasible for them to offer the service more cheaply, through *load balancing* and improved efficiency of resource use (Vouk 2008, p.243).
- Grossman (2008, p.24) suggests that cloud computing allows user to access capacity exactly as they need it—they pay only for the computing they need. Put another way, there is a saving on the opportunity cost of capital that would otherwise be invested step-wise in increasing capacity to meet growth requirements . Avoiding such stepped expenditures may be even more important during the early stages, when the venture may be speculative and cash may be limited (DeWire, quoted in Gray 2006, p. 173).
- The costs associated with buildings, plant and equipment required to house and condition the work area for ICT equipment may be reduced or eliminated.
- IT support costs including maintenance, backup, help desk services and the like may be reduced or be provided more cheaply by the cloud provider.
- Finally, much ICT equipment is prone to obsolescence over a short timeline, and cloud computing offers a model to avoid these costs (Alpern 2011).

Clearly the precise balance of these components and hence the actual savings to be made—if any—will differ from organisation to organisation, situation to situation. Determining cost savings as part of any specific justification for cloud computing must be done case-by-case: generic models can only be indicative of possible scenarios.

1.2.2 Higher quality of ICT services

Claims of unreliability have been made in respect of cloud computing, based on the premise that in-house staff can fix problems more rapidly (Alpern 2011). Yet a reputable provider will maintain their systems and procedures professionally. For smaller organisations and individuals, this may provide access to better disaster recovery procedures and facilities, stronger security measures (Grossman 2008, p.26), more frequent backups, faster networks, more reliable equipment, more sophisticated applications, than they could have afforded to deploy for themselves.

As well, cloud computing may give access to functions not generally feasible with conventional computing arrangements. Aymerich et al (2008, p. 115) give the example of searching over gigabytes of e-mail—help manuals, customer call records or other text-oriented stores also come to mind as possibilities here.

These impacts will tend to vary inversely with the size of the cloud consumer.

1.2.3 Increased Organisational Agility

Because it enables ICT to be deployed dynamically, cloud computing may assist organisations and individuals to respond to change more rapidly. Scenarios include:

- Dynamic allocation of ICT infrastructure—hardware, network, server capacity—as required to meet growth needs and peak load demands without bottlenecks (Aymerich et al 2008, p. 115), and
- Low usage or specialised applications may be made more accessible if there is no requirement to select, buy and install for only limited use.

The effect is to be able to scale up rapidly in response to growth, merger and acquisition activity, and to be able to rapidly implement applications to support business strategy (DeWire, quoted in Gray 2006, p. 173).

The importance of this will vary with the stability of the Cloud Consumer in terms of growth, merger and acquisition activity, and business strategy.

1.2.4 Mobility

Ubiquitous access implies that access anytime, anywhere. For staff who are dispersed, required to travel, to work out of normal business hours, or to work in a variety of time zones. For the organisation wanting to offer 24/7 access, or global access, cloud computing offers a variety of options to use to meet these demands.

1.2.5 Reduction in Environmental Impact

More efficient machine usage has the effect of improved efficiency per kilowatt-hour of power used, and consequent saving in greenhouse gas generation and fossil fuel usage. Two factors allow cloud computing to improve machine efficiency:

- Packing multiple sessions, possibly from different consumers (enterprises or individuals), into the same machine (multi-tenancy).
- Packing the computing capacity of multiple machines into a smaller number of machines (virtualisation).

As well, support for telecommuting may contribute through avoided staff travel.

The precise balance of these factors will differ in each situation: estimating the 'green' impact must be done case-by-case.

1.2.6 Focus on Business Strategy

Perhaps most significantly, the Cloud Security Alliance (2010, p. 6) suggests that cloud computing offers an unparalleled opportunity for organisations to focus on core competencies that offer real Competitive Advantage (Porter 2001). Two key factors contribute to this opportunity:

- The dynamic nature of resource allocation in cloud computing will allow a closer alignment between IT and business strategy; and
- By reducing or removing the distractions associated with in-house ICT management, the organisation can re-focus its resources and management effort to greater effect (Alpern 2011).

1.3 COST CONSIDERATIONS

Any brief market survey will show a bewildering array of service providers, with a wide range of pricing structures[4].

- **Zettagrid** (*Zettagrid – Shopping Cart* 2011) is based in Australia and offers a range of basic server and license packages charged by time period (month, quarter, year). Access to MS-Office Professional is $365 per annum per user; a web server can be deployed for $288 per month.
- **Web 24** (*Web 24* 2011) are also located in Australia and offer a wider range of infrastructure packages starting from $12.95 per month.
- **Cloud Central** (Cloud Central 2011) is an Australian player with a range of basic infrastructure: servers range in price from 3 cents per minute or $20 per month for the "nano" package through to $1.92 per hour or $1,280 per month for the "huge" package.
- **NEC** (*Cloud Computing – NEC Australia* 2011) started providing ICT infrastructure services in Australia four decades ago, with the provision of what was then known as "time-sharing" on mainframes. It is transitioning to the 21st century with its cloud-based offerings, which are charged by the month.
- By comparison, the USA firm **Gogrid** (*Cloud & Hybrid Hosting Pricing : GoGrid Pricing* 2011) offers two payment mechanisms: Pay-As-You-Go and Prepaid. Charging is per gigabyte for data transfer and storage; by hourly-usage for server usage. Dedicated services are available and are charged by time period (month or year). For a web server, the Pay-As-You-Go charge is 19 US cents per hour; Prepaid, the same web server is charged at 5-8 US cents per hour.

[4]. These examples were viewed in February 2011. Prices are in Australian Dollars unless otherwise indicated.

- At ***Amazon EC2*** *Pricing* (2010) there are several models for providing instances: On-Demand, Reserved, and Spot, and as well there is a (relatively) low-usage price that is free of charge. OnDemand usage of a Windows environment ranges from 3 US cents per hour to US$2.48 per hour for large-scale instances; Reserved prices run from 1-3 years and range (for equivalent capacity) from 1.6 US cents up to US$1.48, whilst the Spot price model offers a similar idea to “standby” pricing on airlines, and is currently ranges from 1.3 US cents to US$1.10 per instance.
- **Salesforce.com** (*CRM – The Enterprise Cloud Computing Company* 2011) provides an example of Software as a Service, with their Customer Relationship Management application available over the Internet at prices ranging from $1.00 per user per month (minimum $5) through to unlimited usage at $360 per month.

Some trends may be drawn from this heterogeneous group of examples.

1. There is a wide range of services, pricing structures, payment structures and providers available. Market testing in the light of specified business requirements will be essential to ensure that an appropriate provider is selected, and that an appropriate charge is incurred.
2. Undertaking such market testing will be complicated by the fact that cloud computing services may be obtained by suppliers across the world. If the provider's servers are located near a target audience then this may increase the chance of reasonable response times. Equally it may expose data to inappropriate laws and practices. Certainly it will increase the number of suppliers that could be selected, and consequent expense of selection.
3. With the complexity already evident, there is a real possibility that the charging structures for cloud computing could go the same way as telephony charging structures, which are reported by the Australian Consumers Association (Duncombe 2007) to be "too complex even for quantitative analysis and modelling", noting that "it took a Federal Court judge to adjudicate over which of two mobile phone plans was cheaper".
4. To add to the confusion, the industry is as yet young, and it is possible that metering methods, charging structures and the like will evolve significantly in the mid-term. This could significantly change the cost structure of any given implementation, and accordingly the desirability of choosing that option.
5. The market is dominated by infrastructure and generic productivity tools such as MS Office, Customer Relationship Management tools (eg: Salesforce.com) and the like. More specialised applications are not easy to find.

1.4 RISKS

1.4.1 Security

Whilst the Internet is not integral to the definition of cloud computing, in practice many configurations will use it: limitations and risks inherent in any Internet usage are accordingly relevant. **Appendix C: Internet Security** describes some of these risks around security and privacy, and indeed this is a key concern expressed by NIST (*Cloud Computing* 2011). Security, as Gray (2006) notes in his discussion of Web Services, "is a major, major issue." Given the focus on data as the core of the requested service, this comes as no surprise. Whether customer data subject to privacy legislation; commercially sensitive quotations or sales data; diary entries that could track an individual for fraud or assault; targets for industrial espionage; intellectual or industrial property such as designs; new market propositions undergoing piloting: at some point in the cloud data will leave the custody of individuals and resources under local control and pass to a remote third party.

Is this a practical problem? Once in the cloud, the data may move beyond national borders to jurisdictions where privacy standards are lower or data may be monitored by government agencies (Gellman 2009, Goodin 2010a). Then there is a real risk of what Choo (2011) refers to as "rogue providers", mining the data for their own purposes such as marketing and re-selling. The modular nature of cloud services means that there could be varying levels of security and privacy practiced by different elements of the service chain, unknown to the cloud consumer. Even reputable providers may not provide safety. Using Amazon's EC2 service as a case study, Ristenpart et al (2009) successfully extracted information from a target virtual machine through a side-attack launched from a virtual machine on the same physical machine. It appears that these types of vulnerabilitites are not completely eliminated: the breach of Sony's systems in 2011 that compromised the personal data of 100 million customers (Finkle & Baker 2011) was launched through the same service – Amazon's EC2 service (*Amazon server used to hack Sony PSP Network?* 2011).

Further, the author has personal experience of indadvertently accessing data, including identifying data such as names, job titles, phone numbers and e-mail addresses, through a SaaS implementation as recently as September 2011. To make it worse, the data related not only to individuals but also to a security firm, including access to a leave calender.

Such technical vulnerabilities are perhaps to be expected in new large-scale implementations of complex technologies. Further, when these vulnerabilities are exposed, industry response is generally swift and effective. When Tromer successfully demonstrated in 2005 a method for breaking the AES encryption process (Hardesty 2009), Intel rapidly responded with a chip redesign that eliminated the vulnerability in new machines—although this did nothing for the installed base of vulnerable chips. As well, non-technical issues such as jurisdictional control and dishonest or fraudulent providers remain. Certainly the perception is that these are real issues: Proofpoint (2010) reported that 49% of their sample of 261 agreed or strongly agreed that “SaaS and cloud computing solutions in the enterprise increase the risk of data leakage”, up from only 40.5% the previous year. Antonopoulos (2010b) confirms this, noting that only 14% of a separate sample trust IaaS for customer-facing functions. For the cloud customer the ramifications of any security breach are huge, and the potential for security and privacy breaches should not be underestimated. Ponemon Institute (*Ponemon Study* 2011) suggests that the average cost of data breaches rose nearly 6% in 2010, to US$7.2million per breach.

Tellingly, whilst Curtin University is aggressively implementing cloud-based capacity, in respect of core HR, financial and student data the CIO of the University, Peter Nikoletatos, has said "We might end up keeping the data from those applications on campus." (Winterford 2011).

1.4.2 Interoperability and portability

Look back at Figure 1 again, and a further significant risk becomes evident. The references to “Open Standards” clearly imply the possibility of “closed standards”—that is, proprietary standards and interfaces which would have the impact of locking consumers into specific functions or providers because of high switching costs. Interoperability and portability are the other two of the three key concerns expressed by NIST (*Cloud Computing* 2011).

The OpenStack project has been in place for 6 years now, and involves such major players as Citrix, NASA, AMD, Intel and Dell supporting the development of open standards for cloud computing. (*Open Stack Open Source Cloud Computing Software* 2010). In addition, a significant number of players (over 400 to date) in the ICT market have committed to the ‘Open Cloud Manifesto’ (*Open Cloud Manifesto* 2009), which encourages development and deployment of open standards to avoid exactly these concerns. The signatories include such major players as Adobe, Cisco, Hitachi, IBM, Lynx, Novell, RedHat, Siemens AG, Sun, Sybase and Symbian, as well as a host of cloud-oriented specialists. Whilst it is early days to determine whether this initiative will be successful, it is indicative of considerable commercial goodwill towards this issue As well NIST is undertaking significant activity in this respect, with creation of the Standards Acceleration to Jumpstart Adoption of Cloud Computing (SAJACC), a strategy to facilitate collaborative development of standards to support high priority security, interoperability, & portability requirements (Leaf 2010).

1.4.3 Uncertainty

Part of the promise of cloud computing is the low entry cost implicit in transaction-based changing. True enough as far as it goes, but the obvious corollary is that it is easy for transaction-based costing to total to an amount far in excess of the ownership-based model if there is sufficient growth in transactions. In effect, a known capped cost is traded for the uncertain possibility of a lower transaction-based charge.

A twist reported by the CIO of Cadbury, Ashley Peck (Grayson 2010) is the potential for employees to buy transaction-based computing power directly because of the low entry cost, bypassing traditional purchasing channels and controls in the process.

Further uncertainty flows from the specialised nature of cloud providers. If the required services must be sourced from multiple providers, then the system integration task falls to the consumer, who may not be sufficiently educated in these matters to manage the task adequately (Gray 2006, p. 176). In any case this would add to the costs of provider selection and cloud implementation.

1.4.4 Provider dependence

By outsourcing access to its data the consumer creates a new and high risk relationship. In the worst case, if the provider fails, then the consumer is completely exposed, unless and until they establish a new contract with a new provider (Gray 2006, p. 177). Choo (2011) draws attention to scenarios where business continuity may be lost on a temporary (not trivial) basis, ranging from natural disasters that are geographically remote from the consumer and accordingly not planned for, to unintended consequences of equipment seizure for law enforcement purposes.

Further, for the life of the relationship the consumer is exposed to the impact of the internal workings of the provider. For example (Gray 2006, p. 177), in the event that the provider upgrades software then the consumer must follow suit, willy-nilly, re-training its staff to a timeline to suit the provider, not their own business strategy, and in the event that the upgrades prove to have significant bugs then it is the consumer that is exposed. Brenner (2009) reports at least one instance of this situation arising.

Some provisions for managing provider dependence suggested by Gray (2006, p. 178) are incorporated in **Appendix D: Contracts and SLAs**. Perhaps the single most useful step to manage this risk is to have all data backed up to a third party.

1.4.5 Public Sector considerations

Public sector requirements for defence, national security and high legislative compliance, some of these implications increase these impacts. The Department of Finance and Deregulation (DOFD) recognises that cloud computing is “a new ICT sourcing and delivery model NOT a new technology”(DOFD 2011), it also recognises that the cloud will have significant effect on current thinking in this environment. Many of the issues parallel those seen in non-governmental sectors. Others are not: few commercial organisations have to deal with the realities of deploying emergency services to support communities hit by droughts, floods and cyclones whilst supported by location-independent ICT infrastructures.

1.5 OUTLOOK

1.5.1 Scalability

Grossman (2008, p. 25) posits that cloud-computing architectures have "proven to be very scalable", and refers to petabyte storage in cloud architectures as an example. The scalability of cloud technologies is further exemplified by the availability of parallel computing opportunities through cloud computing: for small organisations: this may offer an ability to crunch massive amounts of data for graphical, video, scientific or technological modelling purposes that was previously unavailable to them.

Whilst in principle cloud computing offers significant scaling opportunities, in practice the scalability of any specific situation must be assessed individually. Take a couple of commercial scenarios:

- Scaling a relational database up will only be successful if the internal architecture of the database has been designed for that—and even then may not be completely successful (think of running a petabyte-sized MS-Access database as an example). Adding more processing power may allow more and more users to be added, but this can then lead to slow degradation because of processing bottlenecks inherent in the design, or may lead to unanticipated catastrophic failure at a 'tipping point': in either situation, transaction-based costing may become very expensive.
- Clustered servers require multiple identical configuration with dedicated bandwidth between them—which clearly can't necessarily be guaranteed by a cloud provider (Scheier, RL 2009)

If these scenarios are to be avoided then applications and configurations must be optimised for scalability, which, if even possible, will add to the costs of the cloud implementation.

A further consideration is connection to the cloud provider(s). Grossman (2008, p. 25) notes that the remoteness of hosted cloud services can lead to “latency- and bandwidth-related issues associated with any remote application”, and there are industry reports supporting this (Scheier, RL 2009). Whilst service level agreements may specify response times, whether these are adhered to in practice is a different matter. Kennedy (2011) notes the criticality of implementation of the National Broadband Network if cloud computing is to become truly ubiquitous in Australia.

1.5.2 Growth

Curtin University is looking to move most of its IT infrastructure to a "lease or utility model" over the coming years. Explaining why, the CIO said, "Do we want to be an organisation that builds data centres? No. Our core business is education and research," (Winterford 2011). Many organisations agree. Gartner (quoted in Gregg 2010) suggests that the market size will grow to $US150 billion within the next two years (to 2013). Markets and Markets (*Cloud Computing Market* 2010) confirms this order of magnitude, predicting the global market to triple to $121.1 billion in 2015, with nearly three-quarters of the market relating to *SaaS* supporting Content, Communication and Collaboration offerings, and key players being Adobe Web Connect, Google Mail, Cisco WebEx, and Yahoo Mail. *The 451 Group’s Cloud Computing Outlook 2010 (* 2010) suggests that licensing costs and structures are significantly restricting wider-scale adoption: Information Week found that 63% of a sample of 504 ICT professionals mentioned security concerns and restriction by legacy applications preventing private cloud deployment (Biddick 2010). The clear implication is that growth could be even greater. Kiril (2010) recently reported that the Centre for Economics and Business Research (CEBR) has assessed economic impact to 2015 around EUR 763 billion. In fact, the future may already be here: 49%- of 13,000 business and technology executives from across the world surveyed by CIO.com claimed some form of cloud computing offering in operation at their business. (Glass 2010)

And locally, there is no doubt that cloud computing is gaining momentum. A sample of 600 CIOs in the Asia/Pacific region reported that 28% of them had already implemented some form of cloud computing, and that this is likely to double by the end of 2012 (Morris & Mortensen 2011).

The emerging pattern is for hybrid or 'converged' approaches to cloud computing: thus cloud federation between private and public clouds, or between different public clouds, will become more common (Morris & Mortensen 2011). Indeed, the variation in models being implemented by enterprises may be greater than realised by the ICT industry. Biddick (2010) reports that Indiana University has implemented a private cloud but that some manual steps have been retained in meeting requests, as faster turnaround is not generally required and human intervention in the process improves governance.

Turning to where this growth is likely to be, the Centre for Economics and Business Research (CEBR) report Kiril (2010) suggested there would be significant variation by industry across country: Germany saw a focus on banking and finance sector usage, whilst in France, Italy, Spain and the UK distribution, retail and hotels dominated, and associated job creation would be primarily in government, health and education sectors. There are many specific scenarios that could realise this outlook.

- Strong takeup by "federalised" retailers (for instance, Harvey Norman in Australia), enabling consistent real-time handling of stock, customer and sales data without the overheads of software and storage traditionally required, and will easy access to facilities for peak load management at Easter, Christmas, New Year and End of Financial Year periods.
- Small business may benefits from being able to access more powerful functions such as Enterprise Resource Planning systems without the cost barriers of traditional implementations. More important may be the ability to grow at their own pace, unconstrained by traditional step-wise purchasing and implementation of ICT (Glass 2010).

- The Royal Institute of Chartered Surveyors has suggested that property companies could benefit from cloud computing's support of ubiquitous access for their staff to data and applications when undertaking their core functions on their customer's sites (*Report sees Large Scale Adoption of Cloud Computing in the Property Sector* 2011).
- Large players such as banks with high variability in type of load over the course of the day (ATM transactions during daylight, data transfer runs at night, market transactions at times to meet world market opening hours) may benefit from the ability to balance their load more finely across the working day.
- In the confectionary industry Cadbury has been able to roll out 60-70 branded websites across the world a year using structured "factory mentality" processes supported by cloud-based capacity (Grayson 2010).
- The Department of Finance and Deregulation (DOFD) notes that in the government sector there have been numerous prototypes and implementations, including (DOFD 2011):

    - Implementation of virtualisation software in the Australian Bureau of Statistics to transition to a private cloud environment;
    - Treasury's Standard Business Reporting and Business Names projects have implemented private/community cloud capabilities
    - Department of Immigration and Citizenship is investigating the provision of an end-to-end online lodgement process on a cloud platform
    - The Australian Taxation Office's eTax, ELS and Tax Agent Board systems all employ cloud services.

Clearly, no one model is emerging. The particulars of the organisation, its sector, its industry and its chosen business direction will combine to create specific situations that may benefit from different combinations of elements of cloud computing technologies.

2 CONCLUSIONS

Cloud computing will require if anything a clearer understanding of the problem to be resolved – that is, more analysis, not less.

As discussed on page 4, contemplation of cloud computing can not be undertaken without a clear understanding of the problem or opportunity that it is intended to address, whether at the strategic (organisational) level or at the operational level. Alignment of this business need and the proposed cloud deployment may be tested using the model proposed by Henderson & Venkatraman (1999). Only once this alignment is achieved is it feasible to negotiate and contract for cloud services.

Cloud computing will require complex management to balance risk and reward

From the emerging trend towards hybrid models it appears that simple models—whether single provider or not, public-only or private-only—are not sophisticated enough to attain a reasonable balance between risk and benefit. Complex decision-making will be required in contemplating, planning and implementing more creative installations if cloud computing is to make good on its promise.

The usefulness of hybrid approaches taken together with the issue of provider dependence suggest that a modular approach may reduce risk. This could imply:

- Staging the implementation;
- Retaining some capacity inhouse; and/or
- Contracting with multiple providers to provide different elements of the services (in particular, backup services should be provided separately).

Whilst the organisation may be able to divest itself of some or all of its ICT operations, there will be a critical dependency on the new relationship(s) with cloud providers. To manage these relationships adequately will require the customer to have the following competencies available, and to cost them into the proposal:

- Strong contract negotiation skills and experience;
- Sophisticated understanding of cloud computing concepts and practices (as discussed at *Appendix D: Contracts and SLAs*
- Sound understanding of the complex security considerations (as discussed at *Appendix B: Security Considerations*).

Cloud computing is a high-risk approach to cost-saving

Cost savings will be unique to each situation and must be individually analysed. As well, costs charged by cloud providers may change significantly in the future, as the industry matures. The risk of increased charges must be taken into account when contemplating the financial aspects of a cloud computing proposal.

Even if the cost savings are real, reaping the benefits of cloud computing will almost certainly require dismantling current technologies, leaving the enterprise with little or no alternative arrangements in the event that the cloud provider fails to deliver the service or fails altogether. This must be taken seriously as a risk when designing an implementation, selecting a provider and negotiating the service level agreement.

However, cloud computing offers a means to start with minimal infrastructure costs, to expand incrementally in response to demand if required, and to avoid the costs of obsolesence. This offers a path to some short-term implementations, such as pilots or short lifespan services, or where there is considerable uncertainty in demand for new offerings, whilst also enabling the organisation to mature in its understanding of how to use this new deployment model to best advantage.

Cloud computing offers the opportunity to tighten analysis significantly

To detail requirements and analyse benefits adequately, to predict take-up and fine-tune interfaces, a pilot implementation may be required. However, this option has often been considered too difficult or too expensive to undertake. Cloud computing now offers a means to temporarily deploy appropriate infrastructure rapidly and to scale up as required—to meet the needs of analysis as well as operations, providing the opportunity to tighten the benefits analysis and planning beyond what has previously been available. To date, agile analysis has focused on requirements: cloud computing may offer the first options for agile analysis of benefits.

3 REFERENCES

Alpern, P 2011, 'Debunking 5 Myths on Cloud Computing', *Business Finance*, 22 June, viewed 23 June 2011, http://businessfinancemag.com/article/debunking-5-myths-cloud-computing-0622

Amazon EC2 Pricing 2010, Amazon Web Services, viewed 20 February 2011, http://aws.amazon.com/ec2/pricing/

Amazon server used to hack Sony PSP Network? 2011, Priyo, 14 May 2011, viewed 23 June 2011, http://www.priyo.com/tech/2011/05/14/amazon-server-used-hack-sony-p-25928.html

Antonopoulos, A 2010a, ''Cloud Security: Root of Trust', *Network World*, 2 February, viewed 20 February 2011, http://www.networkworld.com/columnists/2010/020210-antonopoulos.html

Antonopoulos, A 2010b, 'The ABCs of VoIP', *Network World*, 27 December.

Aymerich, FM, Fenu, G & Surcis, S 2008, 'An Approach to a Cloud Computing Network', *1st International Conference on the Applications of Digital Information and Web Technologies - ICADIWT 2008*, pp.113-118.

Biddick, M 2010, 'The Why and How of Private Clouds', *Information Week*, 5 June, viewed 20 February 2011, http://www.informationweek.com/news/hardware/data_centers/showArticle.jhtml?articleID=225300320

Brenner, B 2009, '5 Mistakes a Security Vendor made in the Cloud', *CSO Online – Security and Risk*, 30 September, viewed 20 February 2011, http://www.csoonline.com/article/503778/5-mistakes-a-security-vendor-made-in-the-cloud-

Buyya, R, Yeo, CS, Venugopal, S, Broberg, J & Brandic, I 2009, 'Cloud computing and emerging IT platforms: Vision, hype, and reality for delivering computing as the 5th utility', Future Generation Computer Systems, June, vol.25, no.6, pp.599-616.

Choo, KR 2011, 'Cloud Computing Risks', *Information Age*, January-February, ACS, pp.49-51.

Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update, 2009-2014 (CISCO) 2010a, CISCO, 9 February, viewed 9 January 2011, http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-520862.pdf

Cisco Visual Networking Index: Usage Study (CISCO) 2010b, CISCO, 25 October, viewed 9 January 2011, http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/Cisco_VNI_Usage_WP.pdf

Cloud Central 2011, viewed 2 February 2011, http://www.cloudcentral.com.au/

Cloud Computing 2011, NIST, viewed 21 February 2011, http://www.nist.gov/itl/cloud/index.cfm

Cloud Computing – NEC Australia 2011, NEC, viewed 20 February 2011, http://www.nec.com.au/Capabilities/Cloud-Computing.html?gclid=CLOqxcnYoKcCFUpypAodTD60eQ

Cloud Computing Use Case Discussion Group 2010, *Cloud Computing Use Cases*, version 4.0, viewed 21 June 2011, http://opencloudmanifesto.org/Cloud_Computing_Use_Cases_Whitepaper-4_0.pdf

Cloud & Hybrid Hosting Pricing : GoGrid Pricing 2011, GoGrid, viewed 20 February 2011, http://www.gogrid.com/?_kk=application%20cloud%20computing&gclid=CO2Z17PRoKcCFQTabgod9T3Ofg

Cloud Security Alliance 2010, *Top Threats to Cloud Computing V1.0*, Cloud Security Alliance, viewed 10 February 2011, http://www.cloudsecurityalliance.org/topthreats/csathreats.v1.0.pdf

CRM – The Enterprise Cloud Computing Company 2011, SalesForce, viewed 20 February 2011, *http://www.salesforce.com/au/*

Deering, S & Hinden, R 1998, *Internet Protocol Version 6 Specification*, The Internet Society, viewed 15 January 2011, http://tools.ietf.org/html/rfc2460

DOFD 2011, *Cloud Computing Strategic Direction Paper*, Draft, Commonwealth of Australia, Canberra.

Duncombe, S 2007, 'Phone and Internet Bundles', *Choice,* July, Sydney.

Finkle J & Baker, LB 2011, 'Analysis: Sony woes may cause some to rethink cloud computing', *International Business Times*, 6 May, viewed 31 August 2011, http://www.ibtimes.com/articles/142329/20110507/analysis-sony-woes-may-cause-some-to-rethink-cloud-computing.htm

Gellman R, 2009, *Privacy in the clouds: risks to privacy and confidentiality from cloud computing*, World Privacy Forum, viewed 10 February 2011, http://www.worldprivacyforum.org/pdf/WPF_Cloud_Privacy_Report.pdf

Glass, J 2010, 'Cloud Computing 'can benefit small businesses'', Experian, 6 Oct, viewed 20 February 2011, http://www.qas-experian.com.au/company/data-quality-news/cloud_computing_can_benefit_small_businesses__6208.htm

Goodin, D 2010a, 'Glitch diverts net traffic through Chinese ISP twice in two weeks', *Security*, 10 April, viewed 9 January 2011, http://www.theregister.co.uk/2010/04/10/bgp_glitch/

Goodin, D 2010b, 'Defcon speaker calls Ipv6 a "security nightmare"', *Enterprise Security*, 6 August 2010, viewed 9 January 2011, http://www.theregister.co.uk/2010/08/06/ipv6_security_nightmare/

Gray, P 2006, *Manager's Guide to Making Decisions about Information Systems*, Wiley, Hoboken NJ.

Grayson I 2010, 'Risks and rewards in cloud computing', *The Australian*, 22 June, viewed 20 February 2011, http://www.theaustralian.com.au/australian-it/the-hub/risks-and-rewards-in-cloud-computing/story-fn4hs56q-1225882469170

Gregg M 2010, *10 Security Concerns for Cloud Computing*, Global Knowledge, 3 November, viewed 20 February 2011, http://www.globalknowledge.com/training/whitepaperdetail.asp?pageid=502&wpid=689&country=United+States

Grossman, RL 2009, 'The Case for Cloud Computing', *IT Professional,* Mar-Apr, vol.11, issue 2, pp.23-27.

Hardesty, L 2009, 'Secure computers aren't so secure', MIT press release, 20 October, viewed 20 February 2011, http://web.mit.edu/newsoffice/2009/cryptography.html

Henderson, JC & Venkatraman, N 1999, 'Strategic Alignment: Leveraging information technology for transforming organisations', *IBM Systems Journal*, vol.38, no.2-3, pp.474-484.

Hyperconnectivity and the Approaching Zettabyte Era 2010, CISCO, 2 June, viewed 9 January 2011,

http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/VNI_Hyperconnectivity_WP.pdf

Kennedy, S 2011, 'Predictions 2011: the cloud gathers momentum', *The Australian*, 22 February, viewed 22 February 2011, http://www.theaustralian.com.au/australian-it/the-cloud-gathers-momentum/story-fn7uxtu0-1226005539451

Kiril 2010, 'Cloud Computing Adoption Can Generate EUR 763 Billion in Europe, by 2015, Cloud Dividend Report', *Cloud Tweeks*, 24 February, viewed 25 February 2011, http://www.cloudtweaks.com/2011/02/cloud-computing-adoption-can-generate-eur-763-billion-in-europe-by-2015-cloud-dividend-report/

Labovitz, C, Iekel-Johnson, S, McPherson, D, Oberheide, J & Jahanian, F 2010, 'Internet Inter-domain Traffic', *Proceedings of SIGCOMM 2010*, viewed 7 January 2011, http://conferences.sigcomm.org/sigcomm/2010/slides/S3Labovitz.pdf

Leaf, D 2010, *Overview: NIST Cloud Computing Efforts*, NIST, viewed 21 February 2011, http://csrc.nist.gov/groups/SNS/cloud-computing/documents/forumworkshop-may2010/nist_cloud_computing_forum-leaf.pdf

Lunsford, J 2010, *Limelight Financial Analyst Meeting 2010*, viewed 7 January 2011, http://www.limelightnetworks.com/resources/limelight-networks-financial-analyst-meeting-2010/

Cloud Computing Market 2010, 'Cloud Computing Market - Global Forecast (2010 - 2015) ', Markets and Markets, October, viewed 20 February 2011, http://www.marketsandmarkets.com/Market-Reports/cloud-computing-234.html

Mell, P & Grance, T 2010, *The NIST Definition of Cloud Computing*, NIST, viewed 21 June 2011, http://www.nist.gov/itl/cloud/upload/cloud-def-v15.pdf

Morris, C & Mortensen, C 2011, 'Hybrid Cloud on the Rise: A Key Strategy to Business Growth in Asia/Pacific', *The IDC Circle*, IDC, viewed 24 February 2011, http://www.idccircle.com/Portal/ExecutiveReports/ExecutiveReportLanding.aspx?from=sponsor&ReportId=172

Open Cloud Manifesto 2009, viewed 19 June 2011, http://www.opencloudmanifesto.org/Open Cloud Manifesto.pdf

Open Stack Open Source Cloud Computing Software 2010, Openstack, viewed 20 February 2011, http://www.openstack.org/index.php

Page, L 2010, '"Hyperbolic map" of the internet will save it from COLLAPSE', *Telecoms*, 10 September 2010, viewed 9 January 2011, http://www.theregister.co.uk/2010/09/10/hyperbolic_map_to_save_the_net/

Ponemon Study 2011, Symantec, 8 March, viewed 19 June 2011, http://www.symantec.com/about/news/release/article.jsp?prid=20110308_01&om_ext_cid=biz_socmed_twitter_facebook_marketwire_linkedin_2011Mar_worldwide_costofdatabreach

Porter, ME 2001, 'Strategy and the Internet, *Harvard Business Review*, March, pp.63-78.

Proofpoint 2010, *Outbound Email and Data Loss Prevention in Today's Enterprise, 2010*, viewed 20 February 2011, http://www.proofpoint.com/downloads/Proofpoint-Outbound-Email-and-Data-Loss-Prevention-2010.pdf

Report sees Large Scale Adoption of Cloud Computing in the Property Sector 2011, ITProPortal, 25 February, viewed 25 February 2011, http://www.itproportal.com/2011/02/25/report-sees-large-scale-adoption-cloud-computing-property-sector/

Ristenpart, T, Tromer, E, Shacham, H & Savage, S 2009, 'Hey, you, get off of my cloud: exploring information leakage in third-party computing clouds', *Proceedings of the 16th ACM conference on Computer and communications security – CCS '09*, 9-13 November, Association for Computing Machinery, Chicago, USA.

Scheier, RL 2009, 'Busting nine myths of cloud computing', *Infoworld*, 22 June, viewed 24 June 2011, http://www.infoworld.com/d/cloud-computing/busting-nine-myths-cloud-computing-260?source=fssr

The 451 Group's Cloud Computing Outlook 2010, 'The 451 Group's Cloud Computing Outlook 2010', Information Engineer, 7 February, viewed 20 February 2011, http://www.informationengineer.org/2010/02/06/the-451-groups-cloud-computing-outlook-for-2010.html

Tips for embracing cloud computing 2011, Help Net Security, 9 June, viewed 23 June 2011, http://www.net-security.org/secworld.php?id=11140&utm_source=feedburner&utm_medium=feed&utm_campaign=Feed%3A+HelpNetSecurity+%28Help+Net+Security%29

Vouk, MA 2008, 'Cloud Computing – Issues, Research and Implementations', *Journal of Computing & Information Technology*, Dec, vol.16, issue 4, pp.235-246.

Wang, L, Von Laszewski, G 2010, 'Cloud Computing: a Perspective Study', *New Generation Computing*, Apr, vol.28, issue 2, pp.137-146.

Web 24 2011, viewed 22 February 2011, http://www.web24.com.au/splash2.php

Winterford, B 2011, 'Optus launches cloud computing beta', *IT News*, 11 Feb, viewed 21 February 2011, http://www.itnews.com.au/News/166904,optus-launches-cloud-computing-beta.aspx

Zettagrid – Shopping Cart 2011, Zettagrid, viewed 20 February 2011, https://www.zettagrid.com/summary

4 APPENDIX A: CONCEPTS & VOCABULARY[5]

Access Management: The ability to write policies (typically in XACML) that examine security tokens to manage access to cloud resources. Access to resources can be controlled by more than one factor. For example, access to a resource could be limited to users in a particular role, but only across certain protocols and only at certain times of the day.

Agility: How quickly the *provider* responds as the consumer's resource load scales up and down.

API: see *Application Programming Interface*.

Application Programming Interface: An application programming interface (API) is a contract that tells a developer how to write code to interact with some kind of system. The API describes the syntax of the operations supported by the system. For each operation, the API specifies the information that should be sent to the system, the information that the system will send back, and any error conditions that might occur. APIs can be defined in specific programming languages or in more neutral formats. An API can also include the details of protocols (such as HTTP) and data formats. An API requires human intelligence to understand the semantics of the data and operations.

Audit and Compliance: The ability to collect audit and compliance data spread across multiple domains, including hybrid clouds. Federated audits are necessary to ensure and document compliance with SLAs and regulatory requirements.

Automation: What percentage of requests to the provider are handled without any human interaction.

[5] From (Cloud Computing Use Case Discussion Group 2010), unless indicated otherwise.

Automated Broker: An automated cloud broker that could dynamically select cloud providers based on business criteria defined by the consumer. For example, the consumer's policy might state that the broker should use the cheapest possible provider for some tasks, but the most secure provider for others.

Broker: A broker has no cloud resources of its own, but matches consumers and providers based on the SLA required by the consumer. The consumer has no knowledge that the broker does not control the resources.

Cloud bursting: Cloud bursting is a technique used by hybrid clouds to provide additional resources to private clouds on an as-needed basis. If the private cloud has the processing power to handle its workloads, the hybrid cloud is not used. When workloads exceed the private cloud's capacity, the hybrid cloud automatically allocates additional resources to the private cloud.

Community cloud: The cloud infrastructure is shared by several organizations and supports a specific community that has shared concerns (e.g., mission, security requirements, policy, and compliance considerations). It may be managed by the organizations or a third party and may exist on premise or off premise (Mell & Grance 2010).

Configuration Management: The ability to federate configuration data for services, applications and virtual machines. This data can include access policies and licensing information across multiple domains.

Customer service response: Whilst cloud computing is generally taken to be automated and self-service, this refers to the human interactions required when something goes wrong with the on-demand, self-service aspects of the cloud.

DaaS: see *Data as a Service.*

Data as a Service: Wang et al (2010) propose that Data as a Service (DaaS) is "data in various formats and from multiple sources…accessed via services by users on the network", and they cite examples such as GoogleDocs and Amazon Simple

Storage Service (S3). It is unclear from their definition how this is a service, unless their (unstated) assumption is that the provider owns the data rather than the customer.

Durability: How likely the data is to be lost.

Elasticity: The ability for a given resource to grow.

Federation: Federation is the act of combining data or identities across multiple systems. Federation can be done by a cloud provider or by a cloud broker.

Governance: Governance refers to the controls and processes that make sure *policies* (refer below) are enforced.

HaaS: see *Hardware a Service.*

Hardware-as-a-Service: Hardware-as-a-Service (HaaS): see Infrastructure-as-a-Service for a discussion of this concept.

Hybrid cloud: The cloud infrastructure is a composition of two or more clouds (private, community, or public) that remain unique entities but are bound together by standardized or proprietary technology that enables data and application portability (e.g. cloud bursting for load-balancing between clouds) (Mell & Grance 2010),

IaaS: see *Infrastructure as a Service.*

Identity Management: The ability to define an identity provider that accepts a user's credentials (a user ID and password, a certificate, etc.) and returns a signed security token that identifies that user. Service providers that trust the identity provider can use that token to grant appropriate access to the user, even though the service provider has no knowledge of the user.

Infrastructure-as-a-Service: For Infrastructure as a Service (IaaS), the provider maintains the storage, database, message queue or other middleware, or the hosting environment for virtual machines. Thus instead of running a virtual server on their own machine, the cloud consumer runs a virtual service on a virtual machine

(Gregg 2010) - the consumer uses that service as if it were a disk drive, database, message queue, or machine, but they cannot access the infrastructure that hosts it. NIST (Mell & Grance 2010) clarifies that, ”the capability provided to the 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 consumer does not manage or control the underlying cloud infrastructure but has control over operating systems, storage, deployed applications, and possibly limited control of select networking components (e.g., host firewalls).Wang et al (2010) suggest that *IaaS* is composed of Hardware-as-a-Service (*HaaS*) and *Software-as-a-Service* , but this second level taxonomy does not have great traction in the literature.

Integration: Integration is the process of combining components or systems into an overall system. Integration among cloud-based components and systems can be complicated by issues such as multi-tenancy, federation and government regulations.

Interoperability: Interoperability is concerned with the ability of systems to communicate. It requires that the communicated information is understood by the receiving system. In the world of cloud computing, this means the ability to write code that works with more than one cloud provider simultaneously, regardless of the differences between the providers.

Linearity: How a system performs as the load increases.

Load balancing: When elasticity kicks in (as new VMs are booted or terminated, for example).

Multi-Tenancy: Multi-tenancy is the property of multiple systems, applications or data from different enterprises hosted on the same physical hardware. Multitenancy is common to most cloud-based systems.

PaaS: see *Platform as a Service.*

Platform-as-a-Service: For Platform as a Service (PaaS), the provider manages the cloud infrastructure for the platform, typically a framework for a particular type of application. The consumer's application cannot access the infrastructure underneath the platform. More generally, NIST (Mell & Grance 2010) suggests that the capability this service model provides to the consumer is "to deploy onto the cloud infrastructure consumer-created or acquired applications created using programming languages and tools supported by the provider. The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, or storage, but has control over the deployed applications and possibly application hosting environment configurations." All phases of the System Development Lifecycle may supported, using Application Programming Interfaces APIs, gateways, and portals. Some providers doe not allow software created by their customers to be moved off the provider's platform (Gregg 2010)

Policy: A policy is a general term for an operating procedure. For example, a security policy might specify that all requests to a particular cloud service must be encrypted.

Portability: Portability is the ability to run components or systems written for one environment in another environment. In the world of cloud computing, this includes software and hardware environments (both physical and virtual).

Private cloud: The cloud infrastructure is operated solely for an organization. It may be managed by the organization or a third party and may exist on premise or off premise (Mell & Grance 2010).

Public cloud: The cloud infrastructure is made available to the general public or a large industry group and is owned by an organization selling cloud services (Mell & Grance 2010).

Reliability: How often the service is available.

SaaS: see *Software-as-a-Service.*

Service Level Agreement: A Service Level Agreement (SLA) is a contract between a provider and a consumer that specifies consumer requirements and the provider's commitment. Typically an SLA includes items such as uptime, privacy, security and backup procedures.

Single Sign-On / Sign-Off: The ability to federate logins based on credentials from a trusted authority. Given an authenticated user with a particular role, federated single sign-on allows a user to login to one application and access other applications that trust the same authority. Federated single sign-off is part of this pattern as well; in many situations it will be vital that a user logging out of one application is logged out of all the others. The Single Sign-On pattern is enabled by the Identity Management pattern.

SLA: see *Service Level Agreement.*

Software-as-a-Service: For Software as a Service (SaaS), the provider installs, manages and maintains the software. The provider does not necessarily own the physical infrastructure in which the software is running. Regardless, the consumer does not have access to the infrastructure; they can access only the application. NIST (Mell & Grance 2010) expands this a little: "[t]he capability provided to the consumer is to use the provider's applications running on a cloud infrastructure. The applications are accessible from various client devices through a thin client interface such as a web browser (e.g., web-based email). The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, storage, or even individual application capabilities, with the possible exception of limited user-specific application configuration settings."

Throughput: How quickly the service responds.

Trust: The ability for two parties to define a trust relationship with an authentication authority. That authentication authority is capable of exchanging credentials (typically X.509 certificates), and then using those credentials to secure messages and create

signed security tokens (typically SAML). Federated trust is the foundation upon which all the other secure federation patterns are based.

Virtual Machine: A virtual machine (VM): A file (typically called an image) that, when executed, looks to the user like an actual machine. Infrastructure as a Service is often provided as a VM image that can be started or stopped as needed. Changes made to the VM while it is running can be stored to disk to make them persistent.

VM: see *Virtual Machine*.

5 APPENDIX B: SECURITY CONSIDERATIONS

The Cloud Security Alliance (2010, p. 6) documents what it sees as the top security threats to successful cloud computing, in no particular order of significance. Some of the remediations suggested require action by providers, others require action by industry standards organisations and governments. Actions that could usefully be taken by a cloud consumer (enterprise or individual) are noted here for each of these threats; explanatory notes added in italics.

- Threat: Abuse and Nefarious Use of Cloud Computing
  - ✓ Monitor public blacklists for one's own network blocks *with a view to changing providers if necessary*
- Threat: Insecure Application Programming Interfaces
  - ✓ Analyze the security model of cloud provider interfaces *(and choose a provider only if they have a strong model.*
  - ✓ Ensure strong authentication and access controls are implemented in concert with encrypted transmission *and ensure this is specified in the SLA*
  - ✓ Understand the dependency chain associated with the API, *so as to be aware of the ramifications of specific weaknesses*
- Threat: Malicious Insiders *(ie: inside the cloud provider – and because of this all these remediations can only be implemented via the SLA))*
  - ✓ Enforce supply chain management and conduct a comprehensive supplier assessment, .
  - ✓ Specify human resource requirements as part of legal contracts.
  - ✓ Require transparency into overall information security and management practices, as well as compliance reporting: Gregg (2010) suggests that the former includes understanding the data classification and encryption protocols and practices, the latter includes training of the provider's staff.
  - ✓ Determine security breach notification processes.
  - ✓ *Finally, require automation of data handling (Gregg 2010).*

- Threat: Shared Technology Vulnerabilities
  - ✓ Promote strong authentication and access control for administrative access and operations.
  - ✓ Enforce *SLAs* for patching and vulnerability remediation.
  - ✓ Conduct vulnerability scanning and configuration audits - *enforce this through the SLA.*
- Threat: Data Loss/Leakage
  - ✓ Contractually demand *(through the SLA)* providers wipe persistent media before it is released into the pool.
  - ✓ Contractually specify *(through the SLA)* provider backup and retention strategies.
- Threat: Account, Service & Traffic Hijacking
  - ✓ Leverage strong two-factor authentication techniques where possible – *select providers only if they offer such techniques*.
  - ✓ Employ proactive monitoring to detect unauthorized activity *and require reports to be forwarded by the provider.*
  - ✓ Understand cloud provider security policies and SLAs.
- Threat: Unknown Risk Profile *(the risk profile of the provider, that is, and accordingly the remediation is to require ongoing clarification of the following matters and enforcment through the SLA))*
  - ✓ Monitoring and alerting on necessary information.
  - ✓ Disclosure of applicable logs and data.
  - ✓ Partial or full disclosure of infrastructure details (eg, patch levels, firewalls, etc.).

Antonopoulos (2010a) succinctly summarises the security issue as one of trust. Every security system is founded on a chain of trust, and if a link in the chain is broken, then the trust is broken. In the cloud, at some point the chain of trust MUST move outside the consumer's control—and accordingly beyond the point of trust.

6 APPENDIX C: INTERNET SECURITY

When purchasing in any service, the potential for service failure is unarguable, and reputable cloud service providers address these considerations in design and deployment of their operations. Thus, for example, the risk of site failure is real, and cloud providers are typically establishing multiple processing centres in different locations to address this issue (Aymerich, Fenu & Surcis 2008, p. 116; Buyyaa et al. 2008, p. 2). Equally, "cloud service providers must guarantee that data are processed automatically" (Aymerich, Fenu & Surcis 2008, p. 114): this will avoid processing errors, and will also ensure privacy from scrutiny by the provider's operational staff. A checklist of such items to be addressed in negotiating with a cloud provider is at **Appendix D: Contracts and SLAs.**

The risk…

However, this all still leaves unaddressed a more fundamental risk: the risk of 'Cloud failure'—that is, the risk that the Internet itself could fail as a platform. Is this risk real? The three examples descussed below show that evidence is already coming in to answer this question in the affirmative.

Security

The upgrade of Internet Protocol (IP) to Version 6 (Ipv6) has Joe Klein of the North American Ipv6 task force commenting that "We've… implemented a lot of new security features into IPv6…The problem with it is we're in a transition period and that's going to take anywhere from five to 10 years to fully implement..." (quoted in Goodin 2010b). That there may be bugs, buffer overflows and other obvious errors in the short-term that open up opportunities for hackers is clear. As well, however, this implies that there will be a lengthy period during which there will also be incompatibilities between firewalls, intrusion-prevention devices and other security mechanisms. Accordingly any individual site may have security reduced in order to operate around such incompatibilities.

Routing through insecure locations

Even with this matter resolved, no matter how strong the security a cloud provider offers at their data centre, eventually the data must be routed through the Internet to or from the end-user. The unspoken assumption behind claims of security is that transport of adequately encrypted data through the Internet is secure…yet that assumption is false. A fundamental design premise underpinning the Internet—and hence cloud computing—is that packets of data will be routed independently, with no control over the physical channel taken. This is the source of the Internet's robustness under attack, as no one channel can compromise all traffic: it also implies that there is no way of controlling where the traffic actually goes…and so it could be routed through countries that have very different standards of security and privacy than those contracted for, such as China and Pakistan. Goodin (2010a) describes incidents in which this situation occurred, as recently as 2010, and he also describes a demonstration of redirection of internet traffic to a compromised network in 2008. Clearly the Border Gateway Protocol (BGP), through which Internet traffic is routed, is open to abuse.

Routing Failure

Even more fundamentally, Dimitri Krioukov of the University of California observes that "we are already seeing parts of the Internet become intermittently unreachable, sinking into so-called black holes, which is a clear sign of instability". He further comments that "…the existing Internet routing architecture may not sustain even another decade"(quoted in Page 2010). The culprit is again the BGP: Krioukov is suggesting that the task of keeping the tables required by the BGP up-to-date is becoming unmanageable.

The Outlook

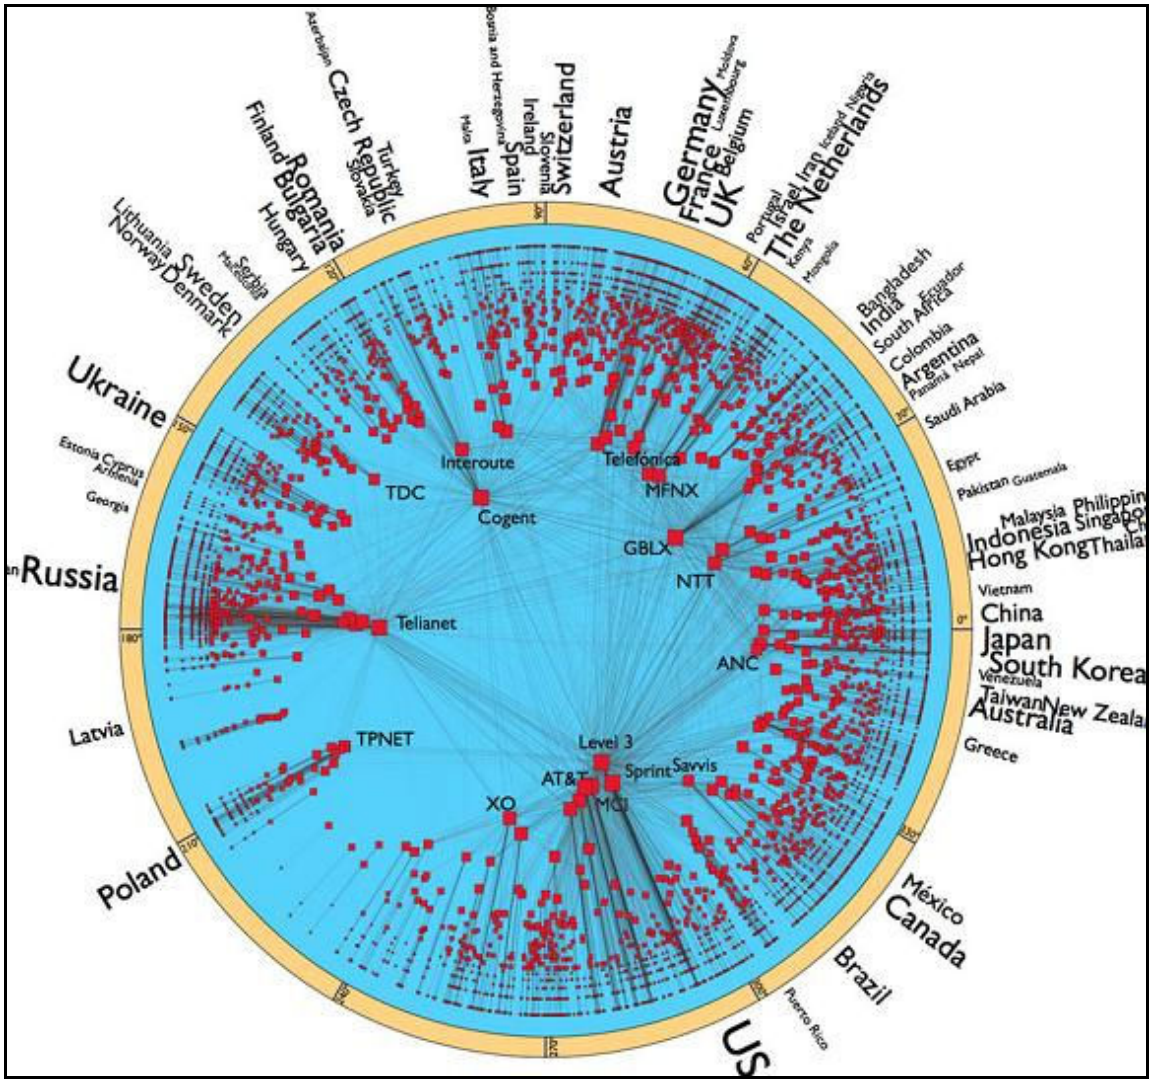

Figure 2: Map of Internet Traffic (Page 2010, p. 2)

Figure 2 maps Internet Traffic by geography, with volume indicated by proximity to the centre of the diagram. Overall, the volume of traffic on the internet has grown roughly 6 orders of magnitude over the last two decades, to around 15 exabytes per month and is increasing at a compound annual growth rate of 34% (Hyperconnectivity and the Approaching Zettabyte Era 2010). Mobile traffic is increasing (CISCO 2010a), whilst peer-to-peer traffic is dropping, leaving web and video dominating as the key areas of traffic and of growth (Lunsford 2010, Labovitz 2010, CISCO 2010b). CISCO (2010b) notes that "[a]s the traffic share of video rises,

so does the volatility of the traffic", further adding to the instability of the Internet as a platform.

Over and above this predictable increase in traffic from known sources, the explosion in IP addresses available through implementation of 128-bit addressing under version 6 (Deering & Hinden 1998)—up nearly 30 orders of magnitude from the 4 billion available under version 4—will open up a host of new opportunities to deploy IP-based devices into new situations and new applications, with seamless integration via the Internet to existing data centres and applications. Without a concommitant increase in the bandwidth available to support communication between this dramatically increased number of addresses, there will only be increasing traffic contention, and a further increase in intermittent transport difficulties.

Conclusion

Whilst cloud computing offers local-level economies and opportunities to enterprises and individuals, it is at the risk of becoming embedded in a single platform of incomplete reliability. In addition, realising the economies on offer will almost certainly require dismantling current technologies, leaving the enterprise with little or no alternative or backup arrangements. With traffic only set to increase, and possibly compound the problem, there is a real risk of data loss, security breach, and poor customer service associated with using this technology. It is critical that any decision to make use of cloud computing technologies be made in this light.

7 APPENDIX D: CONTRACTS AND SLAS

1. The threshold issue is to determine what the are business goals of the cloud initiative .
2. Select one or more provider*s* to meet these needs, keeping in mind that:
   - Grossman (2009) recommends being clear about what is required:
     - on-demand computing instances (roughly equivalent to the number of users),
     - on demand computing capacity(storage and processing available to a user),
     - or both;
   - to manage security concerns to best effect, plan for security, first;
   - the modular nature of cloud computing makes one-stop-shopping less critical;
   - to mitigate key provider vulnerability, identify and select alternative locations and/or providers to re-deploy in the event of service interruptions;
   - to mitigate provider dependence, a separate provider may be selected for backup services;
   - to manage the *SLA*, a separate provider may be selected for monitoring services;
   - as the market evolves, machine-readable SLAs placed with an *automated broker* may provide a fast effective way of obtaining the desired services dynamically and cheaply.
3. For each provider selected:
   - check their track record, financial status and stability;
   - obtain written quotation for the services sought;
   - understand the jurisdiction under which any contracts will be enforced;
   - for key applications, obtain rights to source code in the event of provider failure;

- negotiate ownership of data, the right to retain it when the relationship ends, and the format in which it will be returned (Gregg 2010, p.4).

4. Negotiate and contract an SLA, covering (Cloud Computing Use Case Discussion Group 2010, pp. 54-62):

- A set of services the provider will deliver;
- A complete, specific definition of each service;
- The responsibilities of the provider and the consumer;
- Clarification of relationship, rights and responsibilities between the provider and any brokers or resellers that may be involved;
- Specification of performance indicators for each service, covering such matters as *Throughput, Reliability, Load Balancing, Durability, Elasticity (* with limits such as the maximum amount of storage or bandwidth clearly stated*), Linearity, Agility, Automation* and *Customer Service Response Time* (including many requests the consumer can make, how much they will cost , and how soon the provider will respond;
- Specification of volume reporting, including any data necessary for calculating chargeback to specific areas within the customer's organisation;
- Legal ownership of data, and rights of access (Gray 2006);
- Geographic location of data, and requirement for multiple locations;
- Isolation of customer's data from other data in case of law enforcement action;
- Data retention, storage and deletion protocols: in the context of ephemeral virtual machines - Antonopoulos (2010b) points out that these must be negotiated carefully, particularly in respect of transaction logs (whether for security, regulatory or application purposes), as the logs must persist long after the machines disappear;

- Maintenance standards, procedures and timeframes;
- Disaster recovery and business continuity arrangements, including frequency and storage location of backups (Gregg 2010);
- Security protocols, data classification systems, data encryption standards and application (in transit and/or in storage), reporting of security incidents and support during such incidents to be offered by the provider (Gregg 2010, p.4), and other matters discussed at **Appendix B: Security Considerations;**
- Specificationof regulatory requirements to be adhered to, and protocols for demonstration of regulatory compliance and certification;
- An auditing mechanism to monitor the service;
- Protocol for reporting service level breaches by the provider to the consumer;
- The remedies available to the consumer and provider if the terms of the SLA are not met (Gray (2006, p. 178) suggests incorporating an escape clause);
- How the SLA will change over time;
- Expiry date: Gray(2006, p.178) suggests that this be short-term – less than 3 years, so as to be able to escape if the relationship does not work well.

5. Finally, IBM suggests actively managing the relationship and the risks (*Tips for embracing cloud computing* 2011):
   - Actively monitor service delivery to customers, including response times and outages: decide on appropriate frequency and coverage of this monitoring, and stick to the schedule;
   - Develop a plan for redeploying to alternative sites in the event that service delivery is compromised;
   - Train the in-house team in the plan, and rehearse it.