Available online at www.sciencedirect.com

SciVerse ScienceDirect

www.compseconline.com/publications/prodclaw.htmComputer Law
&
Security Review

How reliable is cloud sourcing? A review of articles in the technical media 2005–11

Roger Clarke^{a,b,c}^a Xamax Consultancy Pty Ltd, Australia^b Cyberspace Law & Policy Centre, University of N.S.W., Australia^c Research School of Computer Science at the Australian National University, Australia

ABSTRACT

Keywords:

Cloud sourcing
Virtualised services
Cloud reliability
Software as a service (SaaS)
Platform as a service (PaaS)
Infrastructure as a service (IaaS)
Power outage

A review of articles in the technical media between 2005 and 3Q 2011 disclosed reports on 49 outages involving 20 cloud sourcing providers. Several of these were major events. Many caused difficulties for user-organisations' staff. Some caused lengthy suspension of services by user-organisations to their customers. A number of them involved irretrievable loss of data. Many user-organisations have failed to risk-assess their use of cloud sourcing, and are exposing their businesses to unmanaged risks.

© 2012 Roger Clarke. Published by Elsevier Ltd. All rights reserved.

1. Introduction

The term cloud computing has been in use since 2006, and was retro-fitted to a few services that were already in existence at the time. Many papers have been published about cloud architectures, cloud technologies, their uses, their impacts on business processes, their potential benefits, and the risks they give rise to. Adoption has been brisk. So it is now possible to gather a body of evidence in order to gain insights into what were previously theoretical investigations and prognostications.

The term cloud computing is applied to several somewhat different forms of service. Their common feature is that servers are 'virtualised'. This means that the concept of a server has ceased to mean 'a computer that runs processes that provide services to other computers', and has reverted to its original sense of 'a process that provides a service to other processes'. This time around, the process can run in any of a large number of computers (and probably an indeterminately large number of them), which can be widely dispersed across many locations and many networks.

Three forms of such 'virtualised' services are usually distinguished (Armbrust et al., 2009; Mell and Grance, 2009).

Software as a Service (SaaS) refers to services at applications level that a human end-user can recognise. Salesforce was the first, and other well-known examples include Zoho, Google Docs and Microsoft's Office 365; and contemporary versions of webmail services such as Hotmail, Yahoo and Gmail. Infrastructure as a Service (IaaS) refers to the rental of what appears to the user-organisation to be a bare machine plus an operating system. Platform as a Service (PaaS) lies between the other two, and provides the user-organisation with a more substantial suite of systems software over which customised applications can be installed or built. Well-known providers include Microsoft Azure, Google Apps, BitBucket and Heroku.

This study assesses the available evidence about the reliability of all three as-a-Service categories. It does this by identifying and assessing the media reports that have been published over a 6-year period, relating to outages of cloud services.

The article commences by briefly reviewing the organisational requirements for reliability of ICT services, which are at risk when cloud sourcing is used. It then explains the research method adopted in the study, and presents the evidence that was gathered. The final section draws together the inferences arising from that evidence.

2. Cloud reliability

In the July 2010 CLSR (26, 4), Svantesson and Clarke (2010) drew attention to risks for consumers in adopting SaaS. This paper focuses on the risks to organisations rather than consumers, and encompasses all three as-a-Service categories. It is a further development on the research reported in Clarke (2010), and follows a number of other papers on related topics in CLSR during 2010–11.

Of the technical disbenefits and risks identified in Clarke (2010), this study focused on several of the critical attributes that make up reliability – availability, robustness, resilience and recoverability (Avizienis et al., 2004), together with the contingent risks of major service interruption and data survival. These play forward to the business attributes of service quality, business continuity and business survival.

Misinformation about cloud computing continues to be reticulated in advertorials, and even in influential business magazines. A major mistake is to misrepresent the attribute of robustness of cloud sourcing as being equivalent to the attribute of robustness of conventional in-house facilities (e.g. Biery, 2011).

In the case of local infrastructure or services such as a desktop, LAN or workgroup server, an outage affects only those people who are local to it. When, on the other hand, every staff-member is dependent on the same infrastructure, the ‘one out, all out’ principle applies: the organisation’s business processes are frozen, and manual fallback arrangements are needed. Some applications are by nature shared organisation-wide, and hence co-dependency risks cannot be avoided but instead have to be managed in other ways. But cloud sourcing extends the co-dependency risk to services that were never subject to it before. After an organisation has adopted SaaS for its office applications, for example, a single server, database, network or power outage renders unavailable the office applications, office documents, mail-archives, appointments and address-books of every staff-member, not merely those local to the point-of-failure. “You have to think about ... not being able to do anything when, say, 10,000 workers are suddenly idled by a single tech outage” (Needleman, 2011).

The research reported on in this paper was performed in order to gain an early insight into the reliability of cloud services, and the impacts on user-organisations.

3. Research method

Articles were sought out that reported experiences of user-organisations, as distinct from those based on promotional material issued by technology providers. In preparation for this study, media monitoring was conducted on a sporadic basis from early 2009 onwards. Retrospective searches were then undertaken using the search-engines of selected, reputable IT media outlets, and Google News. These searches were conducted in November–December 2009 and February–March 2010, and updated and extended in August 2011 and November 2011. The primary search terms used were ‘outage’, in conjunction with ‘cloud computing’, ‘cloud’ and/or ‘SaaS’, ‘PaaS’ or ‘IaaS’.

Dependence on the single key term ‘outage’ risked missing reports of some events. Some reports about non-availability of

services might use other terms and not ‘outage’. It appeared to be even more likely that the word might be omitted from reports of intermittent or otherwise low-quality service, or about lost data, or about closure of a company or withdrawal of a service. In order to test the effectiveness of the search term, complementary searches were conducted using ‘downtime’, ‘unavailability’ and ‘crash’. Fewer items were located in this manner than was the case with ‘outage’, and all reports located this way also contained ‘outage’. A supplementary search was undertaken using ‘data loss’. This identified no additional events, but did locate one additional article of relevance. On the basis of these tests, the research method appears to be robust.

Using the above method, well in excess of 100 media reports were identified that were published between late 2005 and 30 September 2011. All were assessed, and this report summarises the findings. Access to the underlying research material is provided at the end of the paper.

The IT media are not in the business of maintaining a catalogue of outage reports. Rather, they write articles on events that they judge to be newsworthy. So it is to be expected that some events known to the IT media have not been reported, particularly where they were isolated, short, related to a little-known service, or otherwise seemed uninteresting. Conversely, events involving large services with many users are more likely to be reported.

Where multiple reports were found, generally only the most complete or clearest was used, unless the reports were cumulative, in which case each article that contributed significantly to understanding of the event was included in the collection. Articles that appeared to be preliminary or tentative were omitted from the collection. Qualifications, corrections and retractions were sought, although the only ones found were within the articles themselves, in the form of ‘Updates’.

Generally, the articles evidenced sufficient technical understanding on the part of the reporter, and some pointed to the suppliers’ own published explanations. The articles’ quality and reliability were assessed as being from moderate to high.

4. Summary of the articles

A total of 105 media reports were used, which between them identified 49 events, involving 20 providers, 10 of them SaaS, 5 PaaS and 5 IaaS providers. Some over-counting is possible in the SaaS category because some services may still have been running on conventional server-farms at the time of the event, rather than depending on the virtualised servers that are the hallmark of the cloud notion. On the other hand, some under-counting is inevitable, due to the dependence on only a sub-set of the very large number of technical media outlets. In addition, many events may not have come to the attention of the media, or may have come to attention but been judged to be insufficiently newsworthy to warrant an article.

4.1. SaaS

In all, 26 events of the 49 related to 10 Software as a Service (SaaS) providers. The earliest reports, in 2005 and 2006, concerned Salesforce, a Customer Relationship Management (CRM) app launched early in the century, which was the first

service to have the terms cloud computing and SaaS applied to it. The third event, in January 2009, had the effect of locking out 900,000 subscribers. It only lasted about an hour, but subscribers are highly dependent on it being ever-present.

The Web version of the Intuit accounting package had a 40-h outage on its accounting services in June 2010. Like Salesforce, Intuit has large numbers of small business users who are highly-dependent on its availability, including for inbound payments processing, particularly during business hours, but also into the evening.

A further eight SaaS providers were reported as having outages. Apple MobileMe had major problems immediately after launch in July–August 2008. Gmail was the subject of 7 clusters of outage reporting during the 3 years from mid-2008. Google Docs had outages in May 2011 and September 2011. One report in May 2011 related to Google Blogger. Microsoft Danger T-Mobile/Sidekick had a major problem in October 2009. The Swissdisk backup service was out in October 2009. Microsoft BPOS and Live services were the subject of 8 separate reports, and Microsoft's replacement Office 365 service had difficulties immediately after launch in July–August 2011 and again in September 2011, with a 3–4 h outage.

The September 2011 outages of Google Docs and MS Live and Office 365 were very serious, because, by that time, large numbers of organisations had switched their staff completely to these providers, and there was complete denial of service during c. 4-h periods.

Loss of data occurred in 4, or possibly 5, of the 26 events, including Apple MobileMe (July–August 2008), Microsoft Danger T-Mobile/Sidekick (October 2009), Swissdisk (October 2009), possibly Microsoft Hotmail (December 2010), and Gmail (February 2011). None of these cases involved wholesale loss of data (as has occurred occasionally with conventional ISP services). Nonetheless, for some organisations, some of these data losses may have been catastrophic and even survival-threatening.

A further report involved a withdrawn SaaS service. This was the drop.io file-sharing service, closed at 6 weeks' notice following takeover by Facebook. There have been many other closures of Web-based services, including Pownce and Fit-Finder (micro-blogging), Yahoo! Geocities (web-hosting), Yahoo! Photos, Yahoo! Briefcase (file-hosting), Google Video and Google Health. Some users were provided with no warning, or insufficient warning, or were otherwise unable to rescue their content.

4.2. PaaS

Platform as a Service (PaaS) is effectively a wholesale operation, whereby companies depend on the PaaS provider for the means to develop an application, and to deliver services to their customers using it. Of the 49 events, 7 events have been the subject of reports, 3 for Google Apps in June 2008, July 2009 and September 2011; and 1 each for Microsoft Azure in March 2009, BitBucket in October 2009, Heroku (now owned by Salesforce) in April–May 2011, and VMWare Foundry in April 2011. The BitBucket and the Heroku events resulted from outages of the underlying Amazon EC2 IaaS service.

All of these PaaS outages caused their users significant concerns. Several of them have been for very long periods, with BitBucket's outage lasting 19 h, and Heroku's 16–70 h.

None of these outages was reported as having resulted in loss of software or other data.

A special case was omitted from the set because of uncertainty about the appropriateness of categorising it as a PaaS outage. It is of considerable interest, however. Virgin Airline's Australian subsidiary, Virgin Blue, depends on Accenture's Navitaire for its reservation system. The service suffered a 21-h outage in September 2010, and there was no successful cutover to the intended warm-site alternative. The supplier of the controller for the solid-state storage array, NetApp, "positions Navitaire's use of the V-Series as a PaaS (Platform-as-a-service) private cloud service" (Mellor, 2010).

Because PaaS providers are high up the supply-chain, they are of least popular interest and hence the likelihood of under-reporting is higher than in the case of SaaS services.

4.3. IaaS

The remaining 16 events of the 49 related to 5 Infrastructure as a Service (IaaS) providers. As the earliest in the field (in 2006), and currently the largest provider, Amazon's EC2 processing service and S3/EBS data service have been prominent in media reporting, with 7 events in October 2007 (while still in beta, but including some data loss), February 2008 (which also blocked the New York Times site for 2 h), April 2008, July 2008 (8 h), June 2009, and much more seriously April 2011 and August 2011. In several of these events, some data was irretrievably lost.

Another prominent provider, Rackspace, has had 3 reported events, in July 2009, November 2009 and December 2009. XCalibre Flexiscale/Flexiant has had 3 reported events in August 2008, October 2008 and July 2009. EMC Atmos had 1 event reported in February 2010. An Australian provider, Ninefold, had 1 event reported in August 2011.

4.4. A vignette

The event that appears to be most worthy of reporting in slightly greater detail is the Amazon IaaS outage in April 2011. While reconfiguring the internal network within one US data center, engineers accidentally switched traffic through a low-capacity router. Many processes in the EBS storage system immediately started trying to replicate their databases onto storage devices connected to other networks. The already overloaded router was swamped. The attempts to copy data across to other networks increased. Services died. It took nearly 3 days to get the EBS data system working effectively again. The snowball effect meant that EC2 processing services were out for the first 12 h as well. A few databases were not able to be recovered. It took 9 days before an explanation was published.

Despite 5 earlier reported events, including an 8-h outage, it appears that many customers had placed complete reliance on Amazon's infrastructure, and had no contingency plans, with the result that their own services were out for between 1 and 7 days. Amazon's Service Level Agreements (SLAs) are highly restrictive, and it appears that they paid at most a small amount to a small number of customers directly affected by the initial failure. The many customers who had locked themselves into Amazon could do nothing apart from apologise to their own customers and settle any debts they owed under their own SLAs.

5. Inferences from the reports

This section extracts from the collection of media articles aspects that appeared to be common to a range of events.

5.1. Outages are not uncommon

If the total number of events reflects the number of outages that have actually occurred, then, even allowing for the fact that some were long, and some comprised multiple short outages, the total of 49 events is not large considering the number of providers that are now in the market. User-organisations may find it valuable to benchmark these reports against the incidence of outages experienced with in-sourced and conventionally outsourced services.

On the other hand, many providers have had significant outages early in their careers, some have had them at various stages of maturity, many customers are heavily dependent on these providers' services, the dependency is comprehensive across the organisation, and the loss of control means that it is much more difficult to devise and implement effective fall-back measures for outages of cloud sourced services in comparison with in-house services, at least at this immature stage of standardisation and inter-operability.

5.2. Outages arise from multiple causes

Three broad categories of primary cause are apparent:

- power outage
- networking reconfiguration
- the deployment of a new software version

5.3. Providers' safeguards are sometimes ineffective

In multiple cases, so-called 'Uninterruptible' Power Supplies (UPS) have failed to cope with power outages. The all-too-familiar explanation is given that the UPS has been discovered to be subtly dependent on the ongoing availability of the electricity supply that it was supposed to replace, and that is, by definition, unavailable.

A further safeguard that has not always been effective is the much-vaunted 'scalability' feature of cloud sourcing. Inadequate resources to cope with high traffic and processor demands were at the heart of a number of events, including the major Amazon outage of April 2011.

5.4. Failure cascades are prevalent

A feature of many of the incidents was a failure cascade (Buldyrev et al., 2010). This term refers to one foreseeable event triggering one or more others that the designers had not foreseen as a consequence. Many of the cascades have arisen because of the complexity of the architectures and the non-obviousness of many of the dependencies.

In one example, air-conditioning equipment at Rackspace failed to re-start after the electricity supply was re-connected. In most cases, however, the cascades occurred through an

explosion in network traffic, particularly where processes were designed to detect interruptions and replicate themselves and their associated databases across to another network segment. Some of the cascades are reasonably described as self-mutilation, or a DDOS attack from within, or, as one commentator describes it, 'auto-immune disease', i.e. "resiliency mechanisms can sometimes end up causing you harm" (Heiser, 2011).

5.5. Providers have had to be forced to be responsive

Reports on the early events suggested that many providers were unhelpful and even arrogant in their dealings with their customers. They were subjected to complaints directly, in forums and in the blogosphere, and these were reported in the technical media. Most providers then found it to be appropriate to adopt a more conciliatory attitude.

5.6. Providers have often been uninformative

Reports on the early events for many providers indicated a lack of information during and even after the event. Moreover, many of the reports that were provided were written or heavily worked over by non-technical staff, and were largely devoid of useful meaning, e.g. "a bug in our datastore servers", "near simultaneous switch failure", "memory allocation errors", "networking issues", "hardware failure", "a short", "maintenance issues", "an accident", "proactive efforts to upgrade to next generation network infrastructure", "a network event", "malformed email traffic", "data corruption", "network connectivity issues", a "networking interruption" and "a site configuration issue".

Under pressure from customers, commentators and the media, some providers have installed more effective incident reporting mechanisms, and some have installed more effective information publishing arrangements – variously called 'scorecard reports' and 'dashboards'.

Some providers publish reports on the causes of at least some events, on the consequences, and on the process of remediation. Some providers' event reports have long-term URLs, but many post any information they may provide in an ephemeral form or as blog entries that quickly shuffle well down a long web-page, are soon archived, or quietly disappear. Even at the very end of the period under review, there were ongoing calls for enhanced transparency (Winterford, 2011b).

5.7. Outages may affect important ancillary services

In several events, important ancillary services proved to be dependent on the same infrastructure whose failure had caused the outage of the primary service. In particular:

- in January 2009, Salesforce's 'public health dashboard' was out of action at the same time as the services on whose health it was meant to be reporting
- in May 2010, the normal channel for reporting incidents with Microsoft Live was blocked at the same time as the services whose failure users were trying to report
- in May 2011, a Google Blogger failure not only blocked the primary service, but also prevented Google from using

Blogger as a channel for disseminating information about the outage

- in September 2011, the Microsoft ‘status portal’ for all of the many non-functioning services was also off-line

It appears that in some cases the designers of nominally high-availability services have not been aware of the logic underlying robustness and resilience, or have placed low priority on these ancillary services.

5.8. *The direct impacts have sometimes been significant*

All of the incidents involved absence of service for a period of time for some users, and in some cases for all users. Several of the events involved loss of data, and in most of those cases users would have had no way of recovering the lost data because they are entirely dependent on the service provider.

5.9. *The indirect impacts have sometimes been even more significant*

Many organisations have become dependent on SaaS, and their staff are simply unable to perform their functions while the service is unavailable. This persists even after the service is re-launched, if their data is inaccessible or its integrity has been compromised. This applies to all of Salesforce (CRM), Intuit (accounting), Google and MS mail services, and Google, MS and Zoho office suites, but also to calendars, address-books and project management tools.

Many organisations have become dependent on IaaS, and the quality of the services that they provide to their own customers is undermined by cloud outages. Worse, the unavailability of services to customers often lasts longer than the outage of the underlying service. This is because the organisation has to switch to fallback arrangements, and then, after the service resumes, the integrity of the service and the database must be checked, and transactions handled under the fallback arrangements must be re-processed, prior to the resumption of normal service.

In the case of PaaS, more complex supply-chain structures may be in place. In cases like Heroku’s dependence on Amazon in April–May 2011, a single PaaS outage may take out multiple wholesale services, each of which takes out multiple retail services, each of which takes out multiple customers. The recovery processes are inevitably chained, and hence the period of the outages that the ultimate customers suffer may be multiples of the period of the PaaS outage.

5.10. *Few customers are recompensed*

When cloud computing was first being promoted, it was common to see propositions that SLAs were an unnecessary impediment to inexpensive service provision (e.g. [Buyya et al., 2009](#)). Where providers offer an SLA it is commonly in a heavily stripped-down form, and applies to only part of the service. The primary Terms of Service document, meanwhile, commonly attempts to avoid the offering of any warranties, or at least to minimise the scope and quantum of such warranties as are offered.

In the 49 cases that have been reported, few providers’ Terms of Service appear to have embodied any form of compensation. Some providers made token gestures by permitting customers to apply for discounts, subject to conditions, such as evidence of harm. In only 1 of the 49 events was it reported that the provider made material payments: “Rackspace said in a filing with the Securities and Exchange Commission that it would issue service credits totaling \$2.5 million to \$3.5 million” ([Brodin, 2009](#)).

Even where compensation is due, it is in most cases tied to the cost of the services during the period of the outage, and in some cases is offered as an extension to a pre-paid subscription period – which is of zero value if the customer abandons the service. In no case reported was compensation computed on a basis that reflected the business harm caused.

5.11. *User-side safeguards are commonly lacking*

There were few reports of user-organisations having applied multi-sourcing principles to their acquisition of cloud services ([Dibbern et al., 2004](#)). This involves running duplicate services on different cloud-providers, preferably in different geographical areas accessed via different network connections. A welcome sign of maturation, at least among service-providers, appeared in an article at the very end of the period examined: “During the recent EC2 outage, ... Autodesk wasn’t affected because it had plans in place to shift servers to alternative areas of supply” ([Thomson, 2011](#)).

Multi-sourcing is in any case currently hampered by the slow emergence of standards and inter-operability mechanisms. It is also challenging and expensive to implement real-time or near real-time replication across providers, in order to satisfy the pre-conditions for hot-site or warm-site fallback arrangements.

On the other hand, even some very simple and readily available approaches may be being overlooked. One is the use of SMTP to maintain copies of email on the organisation’s own site, even if the primary or working-copy of the organisation’s mail is in a cloud-service operated by Microsoft, Google or some other provider. More generally, however, it has long been a fundamental of business continuity planning that fallback arrangements must be in place in order to cope with contingencies ([Myers, 1996, 2006](#)).

6. Conclusions

Cloudsourcing is still in its infancy, so growing pains are to be expected. Providers need to monitor the experiences of their competitors, and invest in quality assurance, in prompt scalability, and in pre-tested rapid recovery processes.

User organisations are taking up the offers of the burgeoning population of SaaS, PaaS and IaaS providers. Many of them are not limiting their use to applications whose failure or non-availability would cause little harm, and hence they are exposed when outages occur such as those documented in this paper.

A significant proportion of user-organisations appear to have adopted cloudsourcing precipitately, without ensuring that the services will satisfy their business needs. Company directors have a clear obligation at law to ensure that risk

assessments are undertaken, and that risk management plans are in place. This is no longer pioneer territory. Clarke (2010) provided a framework for the activity. Golden (2011) recently suggested that careful attention to the commitments made by the cloud-provider in the SLA needs to be complemented by architectural features that will achieve a ‘fail-soft’ profile and service resilience, including:

1. Architecture of your application for resource failure
2. Architecture of your topology for infrastructure failure
3. Architecture of your deployment for provider failure

The evidence of these reports suggests that many company directors may be in breach of their legal obligations, and that their organisations need to re-visit their IT sourcing strategies, and to do so very quickly.

The seriousness of the situation is coming to be appreciated, with statements of concern by CIOs, headlines such as ‘Has the cloud bubble burst?’ (Winterford, 2011a), and advice in a leading financial newspaper about “a hybrid approach ... the need for multiple cloud providers and for local, non-cloud back-ups of essential data” (Menn, 2011). Cloudsourcing providers need to upgrade their service resilience, and to actively encourage more careful evaluation and adoption of their services, lest their reputations suffer because of customer difficulties and failures arising from cloud computing outages.

Research material

The following underlying research material is available (all in RTF format):

- a list of the **Searches** performed in order to locate media reports, at <http://www.rogerclarke.com/EC/CCEF-Searches-1109.rtf>
- a list of the **Media Reports** that were used, alphabetically by author and date, at <http://www.rogerclarke.com/EC/CCEF-ITMediaReports-1109.rtf>
- a chronological catalogue of **Events**, at <http://www.rogerclarke.com/EC/CCEF-ChronCat-1109.rtf>
- descriptions of the **Key Facts Reported about Each Event**, at <http://www.rogerclarke.com/EC/CCEF-IncidentTable-1109.rtf>

Roger Clarke (Roger.Clarke@xamax.com.au) is Principal of Xamax Consultancy Pty Ltd, Canberra. He is also a Visiting Professor at the Cyberspace Law & Policy Centre, the University of N.S.W., and a Visiting Professor in the Research School of Computer Science at the Australian National University.

REFERENCES

Armbrust M, Fox A, Griffith R, Joseph AD, Katz R, Konwinski A, et al. Above the clouds: a Berkeley view of cloud computing technical report no. UCB/EECS-2009-28. UC Berkeley Reliable

- Adaptive Distributed Systems Laboratory. February, 2009, at, <http://www.eecs.berkeley.edu/Pubs/TechRpts/2009/EECS-2009-28.pdf>; 2009.
- Avizienis A, Laprie JC, Randell B, Landwehr C. Basic concepts and taxonomy of dependable and secure computing. *IEEE Transactions Dependable and Secure Computing* 2004;1(1):11–33.
- Biery ME. Why your business should know more about the “Cloud”. *Forbes*, 2 November 2011, at, <http://www.forbes.com/sites/sageworks/2011/11/02/why-your-business-should-know-more-about-the-cloud/>; 2011.
- Brodin J. Rackspace to issue as much as \$3.5M in customer credits after outage. *Network World*, 6 July 2009, at, <http://www.networkworld.com/news/2009/070609-rackspace-outage.html>; 2009.
- Buldyrev SV, Parshani R, Paul G, Stanley HE, Havlin S. Catastrophic cascade of failures in interdependent networks. *Nature* 2010;464:1025–8. at, <http://www.nature.com/nature/journal/v464/n7291/abs/nature08932.html> (15 April 2010).
- Buyya R, Yeo CS, Venugopal S, Broberg J, Brandic I. Cloud computing and emerging IT platforms: vision, hype, and reality for delivering computing as the 5th utility. *Future Generation Computer Systems* 2009;25:599–616. at, <http://www.buyya.com/gridbus/papers/Cloud-FGCS2009.pdf> (January 2009).
- Clarke R. Computing clouds on the horizon? Benefits and risks from the user’s perspective. In: *Proc. 23rd Bled eConference, Slovenia, June 2010*; 2010. PrePrint at <http://www.rogerclarke.com/II/CCBR.html>.
- Dibbern J, Goles T, Hirschheim R, Jayatilaka B. Information systems outsourcing: a survey and analysis of the literature. *ACM SIGMIS Database* 2004;35:4 (Fall 2004).
- Golden B. Cloud computing and the truth about SLAs. *Networkworld*, 8 November 2011, at, <http://www.networkworld.com/news/2011/110811-cloud-computing-and-the-truth-252905.html>; 2011.
- Heiser J. Yes, Virginia, there are single points of failure. *Gartner Blog*, 30 May 2011, at, <http://blogs.gartner.com/jay-heiser/>; 2011.
- Mell P, Grance T. The NIST definition of cloud computing. *National Institute of Standards and Technology, Information Technology Laboratory*. Version 15, 10-7-09, at, <http://csrc.nist.gov/groups/SNS/cloud-computing/cloud-def-v15.doc>; 2009.
- Mellor C. NetApp and TMS involved in Virgin Blue outage. *The Register*, 28 September 2010, at, http://www.theregister.co.uk/2010/09/28/virgin_blue/; 2010.
- Menn J. Cloud creates tension between accessibility and security. *Financial Times*, 14 November 2011, at, <http://www.ft.com/intl/cms/s/0/6513a4d6-0a06-11e1-85ca-00144feabdc0.html>; 2011.
- Myers KN. *Total contingency planning for disasters: Managing risk – minimizing loss – ensuring business continuity*. New York: Wiley; 1996.
- Myers KN. *Business continuity strategies: protecting against unplanned disasters*. 3rd ed. Wiley; 2006.
- Needleman R. Was brief Google Docs outage a tremor or a tsunami?. *Cnet News*, 7 September 2011, at, http://news.cnet.com/8301-19882_3-20103034-250/was-brief-google-docs-outage-a-tremor-or-a-tsunami/; 2011.
- Svantesson D, Clarke R. Privacy and consumer risks in cloud computing. *Computer Law & Security Review* 2010;26(4): 391–7. at, <http://www.sciencedirect.com/science/article/pii/S0267364910000828> (July 2010).
- Thomson I. Autodesk shifts design apps to the cloud. *The Register*, 22 September 2011, at, http://www.theregister.co.uk/2011/09/22/design_resaped_by_cloud_autodesk_claims/; 2011.
- Winterford B. Has the cloud bubble burst?. *itNews*, 13 September 2011, at, <http://www.itnews.com.au/News/271814>; 2011a. has-the-cloud-bubble-burst.aspx.
- Winterford B. Transparency: a core tenet of the cloud. *itNews*, 20 September 2011, at, <http://www.itnews.com.au/News/272605>; 2011b. transparency-a-core-tenet-of-the-cloud.aspx.