Institute of Advanced Engineering and Science

International Journal of Cloud Computing and Services Science (IJ-CLOSER)
Vol. 2, No. 5, October 2013, pp. 321 – 331
ISSN: 2089-3337

The Cloud's Cloudy Moment: A Systematic Survey of Public Cloud Service Outage

Zheng Li[*,], Mingfei Liang[**], Liam O'Brien[***], and He Zhang[****]**
[*]CRL Lab, National Information and Communications Technology Australia (NICTA), Canberra, ACT, Australia
[**]Australian National University (ANU), Canberra, ACT, Australia
[***]ICT Innovation and Services, Geoscience Australia, Canberra, ACT, Australia
[****]State Key Laboratory of Novel Software Technology, Software Institute, Nanjing University, Jiangsu, China

Article Info

Article history:
Received 15th, 2013
Revised 10th, 2013
Accepted 30th, 2013

Keyword:
Cloud Computing
Cloud Service Outage
Outage Lessons
Public Cloud Service
Systematic Survey

ABSTRACT

Inadequate service availability is the top concern when employing Cloud computing. It has been recognized that zero downtime is impossible for large-scale Internet services. By learning from the previous and others' mistakes, nevertheless, it is possible for Cloud vendors to minimize the risk of future downtime or at least keep the downtime short. To facilitate summarizing lessons for Cloud providers, we performed a systematic survey of public Cloud service outage events. This survey followed the standard and rigorous methodology applied for Evidence-Based Software Engineering. This paper reports the result of the survey, such as: (1) none of the Cloud vendors can avoid suffering from service outages; (2) Cloud service outages could happen at any time, and each Cloud vendor has experienced violation of its Service Level Agreements during the past years; (3) *Climate* and *Age* are two influential factors related to the outage locations; and (4) *Power Outage* and *Routing/Network Issue* are two common classes of Cloud service outage causes. In addition to those findings, our work generated a lessons framework by classifying the outage root causes. The framework can in turn be used to arrange outage lessons for reference by Cloud providers. By including potentially new root causes, this lessons framework will be smoothly expanded in our future work.

Corresponding Author:
Zheng Li
School of Computer Science, ANU and NICTA
7 London Circuit, Canberra City, ACT 2601, Australia
Phone: +61 430100823
Email: imlizheng@gmail.com

1. INTRODUCTION

Cloud computing has increasingly become popular in the present business scenario, with various benefits ranging from convenience to economy. Many organizations are using Cloud to automate their service delivering. However, there are thorny issues and risks in using the Cloud [1]. Among the numerous and different concerns [2], [3], [4], the risk of inadequate service availability has been identified as the top obstacle to adoption of Cloud computing [5], [6], [4]. Given the reality that it is hard to eliminate the downtime of the data center systems behind Cloud services [7], most of the current studies emphasized data/system backup from the perspective of Cloud customers. For example, employing multiple Cloud providers was suggested as being a plausible solution to very high availability service delivery [5].

Although there is no absolute means for preventing outage [7], it is still worthwhile for Cloud vendors to learn from the existing mistakes, so as to minimize the risk of future downtime or at least keep the downtime short [8]. Unfortunately, to the best of our knowledge, few comprehensive investigations into service outages can be found in the literature. In other words, there is a lack of systematic discussion about outage lessons from the perspective of Cloud providers.

As an initial step to summarizing lessons for Cloud providers, we performed a survey of public Cloud service outages by using the Systematic Literature Review (SLR) approach. This paper reports the results of that survey

Table 1. Research Questions

ID	Research Question	Main Motivation
RQ1	Which Cloud provider experienced service outage?	To identify the outage host.
RQ2	When and for how long did the service outage take place?	To identify the outage duration/frequency.
RQ3	Where did the service outage happen?	To identify the outage location.
RQ4	What is the root cause of the service outage?	To identify the outage reason.

together with the methodology of this systematic survey. Due to the limit of resource and time, our work only focused on the top five Cloud vendors. The collected outage data can then be viewed as the representative of all the existing Cloud service outage events. The corresponding data analysis was unfolded to answer four predefined research questions about the outage host, duration/frequency, location, and root cause.

The contribution of this work is twofold. First, based on the data analysis, we highlighted a set of findings (e.g. two influential factors related to outage locations, cf. Subsection 3.3.) when answering the predefined research questions. Second, by classifying the outage root causes, this study essentially generated a framework for accommodating outage lessons for Cloud providers.

The remainder of this paper is organized as follows. Section 2. briefly introduces the methodology used to perform this survey. Section 3. describes the survey result that answers the predefined research questions. Section 4. summarizes a set of sample lessons driven by the root cause classification. Conclusions and some future work are discussed in Section 5..

2. METHODOLOGY OF THIS SURVEY

Cloud service outage events have been generally reported as news or posts scattering over web media, technical websites, blogs, etc. To efficiently perform this survey, we borrowed SLR approach to collect, assess, and analyze the relevant outage reports. As the main methodology applied for Evidence-Based Software Engineering (EBSE) [9], SLR has been widely accepted as a standard and systematic approach to investigation of specific research questions. Although the study objects here are not academic publications, this survey may still benefit from the rigorous review process defined by SLR. Following the guidelines of SLR [10], we did this work mainly covering three steps:

- Identify research questions and prepare the survey.
- Collect relevant outage reports and extract data.
- Analyze the extracted data and report the result.

2.1. Survey Preparation

Similar to preparing an SLR, the preparation of this study generated a review protocol based on a pilot survey. Due to the limit of space, here we only highlight the research questions defined in the review protocol, as listed in Table 1.

2.2. Report Collection and Data Extraction

Unlike exploring various academic libraries in normal SLR, the outage report searching in this study only resorted to the Google search engine.

Furthermore, we employed three constraints for the outage report collection:

(1) This study only focused on public Cloud to make our effort closer to industry needs. Given the large number of players in the market [11], we further limited our concentration to a small set of top Cloud providers (cf. Subsection 3.1.).

Table 2. Top 5 Cloud Providers

Overall Rank	Cloud Provider	Number of Occurences
1	Amazon	24
2	Rackspace	19
3	Microsoft	18
4	Google	17
4	Saleforce.com	17

(2) Considering that the term "Cloud computing" started to gain popularity in 2006 [12], we only collected reports posted between 2007 and 2012. This study did not trace the old outage cases, though some Cloud services like Gmail or Hotmail have existed for a longer time.

(3) This study distinguished unplanned outages from all kinds of Cloud service downtime. In particular, we collected outage information only from the third-party media, so as to ignore the tiny issues that attracted little public attention.

In total, we collected 112 Cloud service outage events. By reading the outage details, we extracted useful data related to the pre-defined research questions for further analysis.[1]

2.3. Data Analysis

The primary data analysis here is to carry out quantitative statistics based on the qualitative data descriptions. Given particular phenomena, we tried to give further explanations or suggestions. Moreover, the root causes of public Cloud service outage have been classified and arranged into a lessons framework for Cloud providers.

3. RESULT OF THIS SURVEY

The survey results are organized and reported following the sequence of answers to the predefined research questions.

3.1. RQ1: Which Cloud providers experienced service outage?

As mentioned previously, numerous public Cloud providers have been increasingly available in the market. It is thus nearly impossible to collect the outage data of different Cloud services all at once. Therefore, we decided to concentrate on the top Cloud providers only. Since different Cloud rankings have been published by different parties at different time, it would be more rational to combine those various opinions. By trying to exhaustively explore the web media and technical websites, we firstly gathered 34 rankings of public Cloud vendors.[2] Then, we rearranged the listed vendor names according to their occurrence numbers. As such, we finally achieved an overall Cloud ranking combining the individuals. Note that we deliberately excluded the Cloud rankings shown in the personal blogs. We found that most of the blogs were either criticized for bias or commented as copies of the others.

Given the limited resource and time, we further narrowed our focus down to the top five public Cloud providers, as listed in Table 2. The collected data (cf. Figure 1) shows that each of the Cloud providers has suffered from considerable service outages, not to mention that there are also unreported downtime events. Such a phenomenon confirms the opinion that service outage happens to any Cloud provider sooner or later no matter how smart or successful the provider is [13], [8]. To reduce the risks of Cloud service outage, building redundancy could be a generic strategy for both Cloud vendors and customers [14].

3.2. RQ2: When and for how long did the service outage take place?

Through the outage distribution illustrated in Figure 1, we also show that the top five Cloud providers suffer from service outages nearly every year. Two exceptions are: there was no outage report of Salesforce.com in 2007,

[1]The extracted Cloud service outage data are shown online: `https://docs.google.com/spreadsheet/ccc?key=0AtKzcoAAmi43dEtPVVlIQ0NRblJiTV9SOGNJb2ttN0E`

[2]The third-party rankings of public Cloud providers are listed online: `https://docs.google.com/spreadsheet/ccc?key=0AtKzcoAAmi43dE1TaTJINUdFM0hqVVQ3dy0wX0M3R2c`

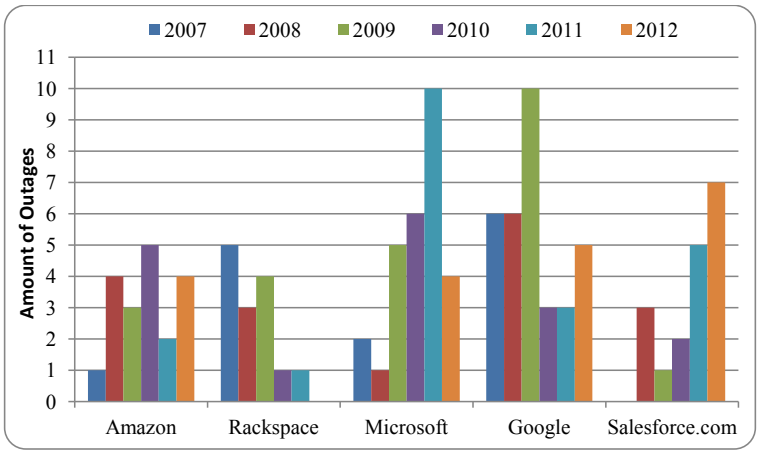

Figure 1. Outage distribution over providers and years.

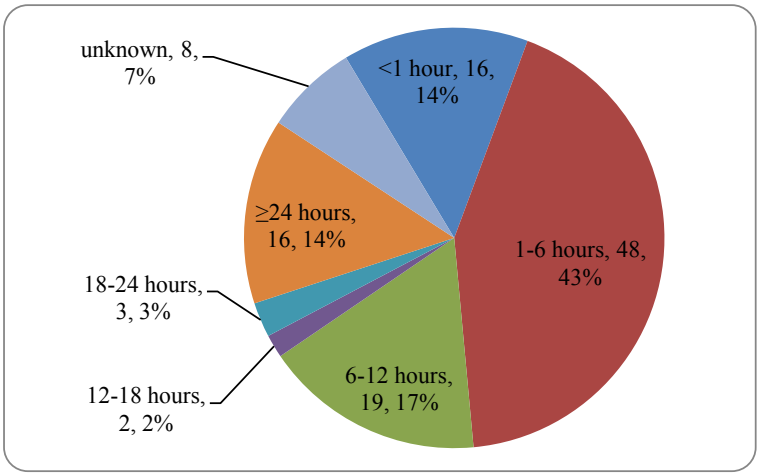

Figure 2. Outage distribution over six-hour scales.

and no report of Rackspace in 2012. The reason could be that Salesforce.com introduced its Platform-as-a-Service (PaaS) in late 2007 [15], while Rackspace started expanding its old data center at the beginning of 2012 [16].

Furthermore, Cloud service outages could happen at any month. In particular, Amazon's northern Virginia data center seems more likely subjected to power outage during June (in 2008, 2009, and 2012), when thunderstorms start appearing frequently across Virginia every year [17].

As for the outage duration, we grouped the collected events into a set of six-hour scales, as illustrated in Figure 2. Interestingly, the Cloud service outages lasted less than 12 hours and the others roughly follow the 80-20 distribution. In particular, we calculated several typical Cloud services' downtime and the corresponding availability, as shown in Table 3. It is clear that each Cloud provider has experienced violation of its Service Level Agreements (SLA) during particular service years. Note that the service downtime here refers to the sum of worst-case outage durations. For example, although Gmail only suffered an average of 10 to 15 minutes of downtime per month in 2008 [18], an unlucky user could have lost the service for around 71.5 hours according to our collected data.

Therefore, we suggest that the industry should help educate Cloud consumers to "expect the unexpected" before using outages. It is understandable that Cloud providers tend to particularly emphasize their SLA for the purpose of market hype. However, the ideal claims could mislead customers and in turn spoil the Cloud ecosphere. For example, users have considered that Cloud "downtime is completely unacceptable" [19], while the truth is that there is no absolute means for preventing downtime when running large-scale Internet services [7].

Table 3. Typical Downtime of Cloud Services

Cloud Service	Year	Downtime	Availability
Amazon S3	2008	$\sim$ 36.5 hours	$\sim$ 99.58%
Amazon EC2	2011	$\sim$ 8.5 days	$\sim$ 97.67%
Rackspace Storage	2011	$>$ 48 hours	$<$ 99.45%
Microsoft Azure	2012	$>$ 26.5 hours	$<$ 99.7%
Google Gmail	2008	$\sim$ 71.5 hours	$\sim$ 99.18%
Saleforce.com Heroku	2011	$\sim$ 104 hours	$\sim$ 98.81%

3.3. RQ3: Where did the service outage happen?

When extracting data, we found that a large proportion of outage events did not disclose their geographical locations, especially the Software-as-a-Service (SaaS) ones. Only Amazon- and Rackspace-related outage reports mostly specified the data centers where the service outages happened. Therefore, we mainly focused on Amazon and Rackspace to answer this research question. Given the reports specifying locations, 72.2% of Amazon outages (13/18) took place in its northern Virginia data center, while 72.7% of Rackspace outages (8/11) happened in its Dallas-Fort Worth (DFW) data center. By roughly investigating these two places, we summarized two influential factors related to the locations of Cloud data centers, namely *climate* and *age*.

Climate Influence: As mentioned previously, power outages occurred with thunderstorms three times (once a year) in Amazon's northern Virginia data center. In fact, thunderstorms are a frequent concern in Virginia, although northern Virginia experiences the least number of such storms [17]. Recall that "away from natural disasters" is one of the principles of site selection for building a data center [20], Amazon may have put this data center at risk from the beginning.

Age Influence: As one of the oldest Rackspace data centers, DFW data center has experienced a series of equipment failures [16]. Interestingly, the northern Virginia data center is also one of the Amazon's oldest. We are then concerned with two points for this phenomenon: on the one hand, old data centers may involve immature techniques and mechanisms from the beginning; on the other hand, a data center could gradually become vulnerable with equipment aging as time goes by. As such, a natural suggestion is that the Cloud data centers should be upgraded regularly.

3.4. RQ4: What is the root cause of the service outage?

Given the collected reports, it is impossible to identify the root cause of every Cloud service outage event. For example, Cloud providers may decline to supply technical details (e.g. Google News outage on September 22, 2009). Therefore, we only focused on the 78 out of 112 events with outage cause explanations. In addition, since an outage event may be a result of a combination of causes (e.g., Microsoft Office365 outage on November 13, 2012) or a close cause chain (e.g., Amazon EBS outage on October 22, 2012), we further broke the 78 outage explanations into 99 cause units.[3] By classifying those units (cf. Table 4), we show a set of typical and relatively frequent root causes of Cloud service outages, as illustrated in Figure 3.

In general, *Power Outage* and *Routing/Network Issue* are two common classes of Cloud service outage causes. The large amount of routing/network issues may not be surprising because the Clouds are inherently associated with intranet and Internet, but the power supply behind Cloud services seems more vulnerable than we expected. As for the details of power outage, it can be confirmed that Uninterruptible Power Supply (UPS) is not uninterruptible enough [21]; thunderstorm (lightning strike) is currently the only natural threat to power equipments; while interestingly, vehicle accident is not rare for being a reason of power interruption, and we thus suggest that the physical barrier security rule [22] should be applied not only to the Cloud data center buildings but also to the outside power infrastructures.

The other three common cause sub-categories are *Hardware*, *Software*, and *Human Mistake*. Each of the three categories covers more than one fifth outage events. In particular, the *Third-party Outage* refers to the scenario that an outage event happens due to other Cloud service outages (e.g., Heroku outage on July 10, 2012). This cause type suggests that not only Cloud consumers but also providers may suffer from the "Cloud ripple effect" [23].

[3]The breakdown of the outage causes are listed online: `https://docs.google.com/spreadsheet/ccc?key=0AtKzcoAAmi43dDdCUjhuTlZiaXJXWEZNQ1FGTTB2blE`

Table 4. Outage Root Cause Classification

Root Cause of Public Cloud Service Outage		
Power Outage	Direct Power Cut/Interruption	
	Hardware	Breaker
		Bus Duct
		Cable
		Electrical Ground
		Power Distribution Unit (PDU)
		Programmable Logic Controllers
		Transfer Switch
		Utility Distribution Network
	Human Mistake	
	Natural Disaster	
	Uninterruptible Power Supply (UPS) Issue	
	Vehicle Accident	
Routing/Network Issue	DNS Error	
	Hardware	Core Device
		Infrastructure
		Routing Device
	Human Mistake	Misconfiguration
		Misoperation
	Request Flood	
	Software	Bug
		Communication Error
		HTTP Error
(Other) System Issue	Database Error	
	Hack	DDoS Attack
		Virus
	Hardware	Chiller Failure
		Recent Change
		Server Down
	Human Mistake	Misconfiguration
		Misoperation
	Overload	Memory Leak
		Request Flood
	Software	Bug
		Recent Change
	Storage Error	
Third-party Outage		

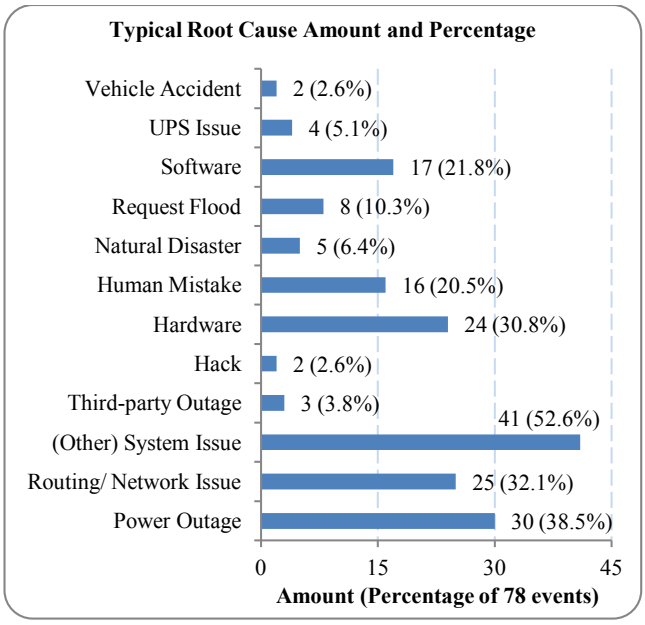

Figure 3. Distribution of several typical root causes. Note that, for the convenience of comparison between root causes, we summed up the outage events with the same sub-category name, although they could be under different primary cause classes.

Overall, the root cause classification shown in Table 4 can be viewed as a lessons framework for Cloud providers (cf. Section 4.). The large amount of lessons learnt from the existing Cloud service outage events can be naturally organized by using this framework. Through smooth expansion, we may continually develop this framework to cover new type of Cloud service outages and accommodate new lessons.

4. SAMPLE LESSONS FROM CLOUD SERVICE OUTAGE

Driven by the outage root cause classification (cf. Table 4), we have tried to summarize and rationalize the existing lessons from the Cloud service outage events. Due to the limit of space, here we only show some coarse-grained sample lessons that can be matched up to those primary cause classes.

4.1. Watch the Power

It has been claimed that "the worst, most sustained downtime has always been caused by power issues", while losing power in data centers will happen someday inevitably [24]. Therefore, we elaborate relatively more on the power-related lessons as follows.

4.1.1. Any single piece of power equipments can fail.

The failed equipment piece could incur cascading events and result in a large scale of power outage. For example, a failed bus duct prevented proper operation of Uninterruptible Power Supply (UPS) (cf. Rackspace outage on July 7, 2009); an electrical ground fault and short circuit in a major power distribution panel interrupted power to a particular Availability Zone (cf. Amazon outage on May 8, 2010); the failure of switches in the electrical infrastructure prevented transfer of electrical load between different power sources (cf. Rackspace outage on June 29, 2009); a breaker failure in the switch board affected all downstream power distribution units (PDUs) (cf. Rackspace outage on November 11, 2007); while failures in a PDU resulted in a portion of servers losing power (cf. Amazon outage on December 9, 2009).

Moreover, some power equipment failures could happen externally. For example, there could be unexpected power cut (e.g., Rackspace outage on December 5, 2007); and even small problems with utility distribution system (e.g., Amazon outage on June 14, 2012) can cause power outage.

4.1.2. Uninterruptible Power Supply (UPS) is not uninterruptible enough.

Given the possibly vulnerable power equipment, employing redundant/backup power systems would be a natural strategy to reduce power issues. An ideal mechanism could be "multiple power supplies in every server connected to 2 PDUs connected to 2 different generators" [25]. Considering the cost-benefit tradeoffs, one of the most practical and common efforts is to use UPS. Unfortunately, this study shows that UPSes could also become a huge single point of failure. Interestingly, UPS has been criticized for its deceptive designation: Once out of commission, UPS can be a solid barrier between Cloud service and generator power [21]. Therefore, regular testing should be further emphasized even for "uninterruptible" power equipments.

4.1.3. Backup power systems should be tested regularly.

Recall the Amazon outage on June 14, 2012, even with the correct setup of generator fallback, the power backup mechanism could still fail unexpectedly in some circumstance. It has been pointed out that the lack of regular testing is the backend flaw, although the power interruption could be the head of a cause chain [26]. In fact, regularly testing the backup power systems has been strongly suggested by industry [24].

4.2. Be Pessimistic about every Service Component

Given the listed root causes of public Cloud service outage (cf. Table 4), it is clear that various hardware and software issues can knock out Cloud services, not to mention the numerous routing/network problems (cf. Figure 3). As such, it would be valuable and necessary for people to realize and understand that "Clouds are made of components that can fail" [27]. As mentioned in Section 3.2., Cloud consumers should be ready to "expect the unexpected" outages, while Cloud providers should rethink and carefully build service levels that actually guarantee services from the perspective of consumers [28]. Note that understanding such a reality does not mean to ask people to passively live with it. On the contrary, being pessimistic about every service component requires both Cloud providers and consumers to fine tune their processes and responses to failures, by conducting full-blown load tests of their failover mechanisms [27], [29].

4.3. Minimize the Chain Reaction

As previously mentioned, we find that an event of Cloud service outage could often comprise a combination of causes or a close cause chain. Some discussions revealed that "the stress of failure will trigger a cascade of other failures" [27]. One of the logics behind this advice could be that human beings tend to make more mistakes under pressure. Inspired by the fire drills that help train people to deal with the event of an emergency, frequent load tests of failover plans may help engineers get familiar with what they need to do to reduce human mistakes when fixing Cloud service outages.

When it comes to the "Cloud ripple effect" [23] resulted from third-party outages, we may draw similar lessons for both Cloud service providers and consumers. In fact, the secondary or tertiary Cloud service providers are indeed customers of their primary providers. Interestingly, a consensus on surviving third-party outages is all about redundancy. Following the terminology from Amazon, the suggestion is to spread the load across multiple availability zones and across multiple regions [30], [31]; in general, the suggestion is to spread across multiple data centers and across multiple primary providers [30], [31]; a more aggressive suggestion is even to spread across public and private Clouds [32]. There is little doubt that such load spreading could be complicated and expensive, however, it would be a worthwhile mechanism if the Cloud services are serious about customer satisfaction [33].

4.4. Open the Outage Details

According to the collected outage data, we find that many events of public Cloud service outage did not disclose the details. It is natural that the "public Cloud" implies some loss of control and visibility from the customers' perspective. Nevertheless, delivering enough in-depth information has been identified crucial for customers especially during the outage [28], [34], [35]. More importantly, opening outage details would also be beneficial for Cloud service providers. There are two reasons for this.

Firstly, rapid and clear lines of communication have been proved a successful crisis management model. The existing case studies show that keeping customers updated can significantly drop off negative commentary [35]. Given the timely disclosure of an outage and the remedy activities thereafter, most customers would still forgive the Cloud service providers for their failings [36].

Secondly, disclosure of outage details can help boost the entire Cloud computing industry. A positive observation on Cloud service outages is that those unfortunate events also provided opportunities to learn from them [29]. By exposing what went wrong, each outage essentially acts as an education for Cloud service providers with how to prevent it from happening again or how to adapt when an outage occurs. For example, the existing Cloud failures have provided useful lessons in disaster planning and infrastructural designing for redundancy, which would reduce future risks and eventually make the Cloud stronger [8], [37].

5. CONCLUSIONS AND FUTURE WORK

Given the lack of comprehensive investigation into the frequent Cloud outage events, we performed a systematic survey of public Cloud service outages. The result of the survey confirms that Cloud service outages would be unavoidable no matter how smart or successful the provider is. As such, we suggest that Cloud consumers should consider the worst cases and carefully design corresponding strategies before employing Cloud services, while not treating SLA as a guarantee. As for the Cloud vendors, on the one hand, they should reduce the market hype and help educate users to "expect the unexpected" issues; on the other hand, it is necessary and worthwhile to learn from the previous and others' mistakes to minimize the risks of future downtime [8].

In addition to revealing findings based on the quantitative analysis, another main contribution of this study is that we finally established an education framework for learning from Cloud service outage. We also list a set of coarse-grained sample lessons to show that the established framework can help guide people to arrange and/or refer to the relevant knowledge.

The main limitation of this work is the completeness of the Cloud service outage data. On the one hand, we only focused on a small amount of public Cloud vendors. On the other hand, we only collected outage events reported by web media and technical websites. Therefore, our future work will further collect outage events of more Cloud vendors and gradually expand the lessons framework. Meanwhile, we will continue summarizing the outage lessons and arrange them within the aforementioned framework for reference by Cloud providers.

ACKNOWLEDGMENTS

NICTA is funded by the Australian Government through the Department of Communications and the Australian Research Council through the ICT Centre of Excellence Program. NICTA is also funded and supported by the Australian Capital Territory, the New South Wales, Queensland and Victorian Governments, the Australian National University, the University of New South Wales, the University of Melbourne, the University of Queensland, the University of Sydney, Griffith University, Queensland University of Technology, Monash University and other university partners.

REFERENCES

[1] G. Anthes, "Security in the Cloud," *Commun. ACM*, vol. 53, no. 11, pp. 16–18, November 2010.

[2] P. A. Boampong and L. A. Wahsheh, "Different facets of security in the Cloud," in *Proc. 15th Communications and Networking Simulation Symp. (CNS 2012)*. Orlando, FL, USA: Society for Computer Simulation International, March 26 - 29 2012, pp. 1–7.

[3] I. Iankoulova and M. Daneva, "Cloud computing security requirements: A systematic review," in *Proc. 6th Int. Conf. Research Challenges in Information Science (RCIS 2012)*. Valencia, Spain: IEEE Computer Society, May 16-18 2012, pp. 1–7.

[4] D. Sun, G. Chang, L. Sun, and X. Wang, "Surveying and analyzing security, privacy and trust issues in Cloud computing environments," *Procedia Eng.*, vol. 15, pp. 2852–2856, 2011.

[5] M. Armbrust, A. Fox, R. Griffith, A. D. Joseph, R. Katz, A. Konwinski, G. Lee, D. Patterson, A. Rabkin, I. Stoica, and M. Zaharia, "A view of Cloud computing," *Commun. ACM*, vol. 53, no. 4, pp. 50–58, April 2010.

[6] A. Rahimli, "Factors influencing organization adoption decision on Cloud computing," *Int. J. Cloud Comput. Serv. Sci.*, vol. 2, no. 2, pp. 140–146, April 2013.

[7] S. Bigelow, "The causes and costs of data center system downtime: Advisory board Q&A," http://searchdatacenter.techtarget.com/feature/The-causes-and-costs-of-data-center-system-downtime-Advisory-Board-QA, June 2011.

[8] Pingdom, "The major internet outages so far in 2008," http://royal.pingdom.com/2008/09/04/the-major-internet-outages-so-far-in-2008/, September 2008.

[9] T. Dybå, B. A. Kitchenham, and M. Jørgensen, "Evidence-based software engineering for practitioners," *IEEE Softw.*, vol. 22, no. 1, pp. 58–65, January 2005.

[10] B. A. Kitchenham and S. Charters, "Guidelines for performing systematic literature reviews in software engineering," Keele Univ. and Durham Univ. Joint Rep., Tech. Rep. EBSE 2007-001, 2007.
[11] CloudHarmony, "Public Clouds," http://cloudharmony.com/clouds, Feburary 2013.
[12] Q. Zhang, L. Cheng, and R. Boutaba, "Cloud computing: State-of-the-art and research challenges," *J. Internet Serv. Appl.*, vol. 1, no. 1, pp. 7–18, May 2010.
[13] C. Brooks, "Heroku learns from Amazon EC2 outage," http://searchcloudcomputing.techtarget.com/news/1378426/Heroku-learns-from-Amazon-EC2-outage, January 2010.
[14] J. Cogswell, "Building redundancy into your Cloud outage strategy," http://searchcloudcomputing.techtarget.com/tip/Building-redundancy-into-your-cloud-outage-strategy, January 2013.
[15] M. LaMonica, "Salesforce.com extends its application platform with Force.com," http://news.cnet.com/8301-10784_3-9778674-7.html, September 2007.
[16] R. Miller, "Rackspace Cloud will expand in Dallas," http://www.datacenterknowledge.com/archives/2012/01/06/rackspace-cloud-will-expand-in-dallas/, January 2012.
[17] B. P. Hayden and P. J. Michaels, "Virginia's climate," http://climate.virginia.edu/description.htm, February 2013.
[18] BBC, "Gmail down again for some users," http://news.bbc.co.uk/2/hi/technology/7934443.stm, March 2009.
[19] K. McLaughlin, "Microsoft Windows Live services suffer global outage," http://www.crn.com/news/applications-os/206900295/microsoft-windows-live-services-suffer-global-outage.htm, February 2008.
[20] M. Fontecchio, "Study ranks cheapest places to build a data center," http://searchdatacenter.techtarget.com/news/1238054/Study-ranks-cheapest-places-to-build-a-data-center, January 2007.
[21] R. McFarlane, "UPS – it's NOT uninterruptible," http://searchdatacenter.techtarget.com/news/1148907/UPS-its-NOT-uninterruptible, November 2005.
[22] C. Higbie, "Rules for designing the urban data center," http://searchdatacenter.techtarget.com/tip/Rules-for-designing-the-urban-data-center, April 2005.
[23] B. Darrow, "Heroku stung by amazon outage," http://gigaom.com/2012/06/15/heroku-stung-by-amazon-outage/, June 2012.
[24] J. Kaplan-Moss, "Lessons from rackspace's downtime," http://jacobian.org/writing/lessons-from-rackspace-downtime/, November 2009.
[25] Hacker News, "Cascading errors caused AWS to go down," https://news.ycombinator.com/item?id=4124719, 2012.
[26] Z. Whittaker, "Amazon explains latest cloud outage: Blame the power," http://www.zdnet.com/blog/btl/amazon-explains-latest-cloud-outage-blame-the-power/80094, June 2012.
[27] S. Hammar, "Lessons learned from the amazon web services outage," http://blog.apicasystem.com/2012/10/24/lessons-learned-from-the-amazon-web-services-outage/, October 2012.
[28] A. Karstens, "Lessons from the amazon outage: 5 ways that cloud providers must take responsibility for service levels," http://blogs.ixiacom.com/ixia-blog/amazon-outage-cloud-provider-service-levels/, May 2011.
[29] A. Rasmussen, "Lessons learned from cloud outages," http://apmdigest.com/lessons-learned-from-cloud-outages, July 2012.
[30] B. Butler, "5 tips for surviving a cloud outage," http://www.networkworld.com/news/2012/042712-cloud-outage-tips-258736.html, April 2012.
[31] A. Dave, "Learn to fail and avoid the next cloud outage," http://www.networkworld.com/news/tech/2013/021113-cloud-outage-266604.html?page=1, February 2013.
[32] D. Horovits, "AWS outage - moving from multi-availability-zone to multi-cloud," http://www.cloudifysource.org/2012/10/24/aws-outage-multi-availability-zone-multi-cloud.html, October 2012.
[33] J. Paterson, "Amazon EC2 outages - lessons not learned," http://www.channelweb.co.uk/crn-uk/view-from-the-channel-blog/2292069/amazon-ec2-outages-lessons-not-learned, September 2013.
[34] A. R. Hickey, "Amazon cloud outage highlights need for transparency," http://www.crn.com/news/cloud/229402233/amazon-cloud-outage-highlights-need-for-transparency.htm, April 2011a.
[35] J. Ainsworth, "Outages: Cloud customers cry out for communication," http://smartdatacollective.com/jamesainsworthalteriancom/39151/cloud-customers-cry-out-communication, August 2011.
[36] P. Wainewright, "Seven lessons to learn from amazon's outage," http://www.zdnet.com/blog/saas/seven-lessons-to-learn-from-amazons-outage/1296, April 2011.
[37] A. R. Hickey, "Amazon's outage will make the cloud stronger," http://www.crn.com/slide-shows/cloud/229402271/amazon-cloud-outage-10-lessons-learned.htm?pgno=9, April 2011b.

BIOGRAPHY OF AUTHORS

Zheng Li received his Degree of M.E. by Research from the University of New South Wales (UNSW). He is now a PhD student at the School of Computer Science at the Australian National University (ANU), and a graduate researcher with the Software Systems Research Group (SSRG) at National ICT Australia (NICTA). He is the author of more than 20 journal and conference publications. His research interests include empirical software engineering, software cost/effort estimation, machine learning, Web service composition, and Cloud computing.

Mingfei Liang holds a Master degree from Australia National University (ANU) and a Bachelor degree of Electronic Engineering from Shandong University. He is responsible for helping data collection for this paper. His research interests mainly include Cloud Computing and Artificial Neural Network.

Liam O'Brien has over 23 years experience in research and development in software engineering. He is a Software and Applications Architect with Geoscience Australia and was previously Chief Software Architect with CSIRO and a Principal Researcher at NICTA's e-Government Initiative. He is also a Member-at-Large of the Service Science Society Australia which he co-founded in 2010. He has previously worked as a researcher with Lero (Ireland), Carnegie Mellon University's Software Engineering Institute (USA), CSIRO (Australia) and the University of Limerick (Ireland). His main areas of research include enterprise architecture, software architecture, SOA, service science, software reuse, software modernisation, and cloud computing. He holds a BSc and PhD from the University of Limerick, Ireland. He is a member of the IEEE and IEEE Computer Society.

He Zhang is a Professor of Software Engineering in the Software Institute at Nanjing University, China. He joined academia after 7 years in industry, developing software systems in the areas of aerospace and complex data management. He has published 70+ peer-reviewed research papers in international journals, conferences, and workshops. His current research areas include software & systems process modeling and simulation, process enactment analysis and process improvement, embedded systems engineering, service-oriented computing, empirical and evidence-based software engineering. Dr. Zhang received his PhD in computer science from the University of New South Wales.