ISSN: 2277-3754
ISO 9001:2008 Certified
International Journal of Engineering and Innovative Technology (IJEIT)
Volume 3, Issue 11, May 2014

Cloud Computing and Unknown Risks

Tariq Ahamad[1], Luay Assidmi[2]
College of Computer Engineering & Sciences
Salman Bin Abdulaziz University, KSA

Abstract: Cloud computing is 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. This cloud model promotes availability and is composed of five essential characteristics, three service models, and four deployment models Cloud computing is appealing from management and efficiency perspectives, but brings risks both known and unknown. Well-known and hotly-debated information security risks, due to software vulnerabilities, insider attacks, and side-channels. cloud computing exacerbates already difficult digital preservation challenges, because only the provider of a cloud-based application or service can archive a "live," functional copy of a cloud artifact and its data for long-term cultural preservation. This Articles discusses serious unknown risks that should be kept in mind before making our cloud potentially weak to perform to the best of expectations and keeps the data, privacy and applications including could itself safe and secure.

I. INTRODUCTION

Cloud is all about computer services, not products. Cloud computing is a general term for anything that involves delivering hosted services over the Internet. Cloud computing is not a technology but a model of provision and marketing IT services that meet certain characteristics [1]. These services are broadly divided into three categories: Infrastructure-as-a-Service (IaaS), Platform-as-a-Service (PaaS) and Software-as-a-Service (SaaS). The name cloud computing was inspired by the cloud symbol that's often used to represent the Internet in flowcharts and diagrams. Cloud computing is the delivery of computing services over the Internet. Cloud services allow individuals and businesses to use software and hardware that are managed by third parties at remote locations [2]. Examples of cloud services include online file storage, social networking sites, web mail, and online business applications. The cloud computing model allows access to information and computer resources from anywhere that a network connection is available. Cloud computing provides a shared pool of resources, including data storage space, networks, computer processing power, and specialized corporate and user applications [3]. While there are benefits, there are privacy and security concerns too. Data is travelling over the Internet and is stored in remote locations. In addition, cloud providers often serve multiple customers simultaneously [4]. All of this may raise the scale of exposure to possible breaches, both accidental and deliberate. Concerns have been raised by many that cloud computing may lead to "function creep", uses of data by cloud providers that were not anticipated when the information was originally collected and for which consent has typically not been obtained [5]. Given how inexpensive it is to keep data, there is little incentive to remove the information from the cloud and more reasons to find other things to do with it.

Security issues, the need to segregate data when dealing with providers that serve multiple customers, potential secondary uses of the data—these are areas that organizations should keep in mind when considering a cloud provider and when negotiating contracts or reviewing terms of service with a cloud provider [6]. Given that the organization transferring this information to the provider is ultimately accountable for its protection, it needs to ensure that the personal information is appropriate handled.

Attractive features and industry momentum make cloud computing appear destined to be the next dominant computing paradigm. Cloud computing is appealing due to the convenience of central management and the elasticity of resource provisioning. Moving critical information infrastructure to the cloud also presents risks, however, some of which are well-known and already hot research topics [7]. The much-discussed challenge of ensuring the privacy of information hosted in the cloud has resulted in an emerging breed of "cloud-hardened" virtualization hardware and security kernels [8]. Similarly, the challenge of ensuring high availability in the cloud has in part fuelled recent research on robust data center networking. Setting aside these known challenges, therefore, this paper attempts to identify and focus on several less well-understood—and perhaps less "imminent"—risks that may emerge from the shift to cloud computing [9]. In this research articles we have discussed stability risks due to unpredictable interactions between independently developed but interacting cloud computations; availability risks due to non-transparent layering resulting in hidden failure correlations; and preservation risks due to the unavailability of a cloud service's essential code and data outside of the provider.

The above list is also probably incomplete: it is likely that other important risks will emerge only as the industry continues its shift to the cloud. Nevertheless, I argue that it is worth proactively investigating longer-term risks such as these before they are certain or imminent, as the stakes may be high. Further, once any of these risks do become important, it may be too late to reconsider or slow the movement of critical infrastructure to the cloud, or to rethink the architecture of important cloud infrastructure or services once they are already perceived as "mature" in the industry.

II. INTERACTING SERVICES FIXITY ISSUES

In the last few years, cloud computing has grown from being a promising business concept to one of the fastest growing segments of the IT industry. Now, recession-hit companies are increasingly realising that simply by tapping into the cloud they can gain fast access to best-of-breed business applications or drastically boost their infrastructure resources, all at negligible cost. But as more and more information on individuals and companies is placed in the cloud, concerns are beginning to grow about just how safe an environment it is.

There are numerous security issues for cloud computing as it encompasses many technologies including networks, databases, operating systems, virtualization, resource scheduling, transaction management, load balancing, concurrency control and memory management [10]. Therefore, security issues for many of these systems and technologies are applicable to cloud computing. For example, the network that interconnects the systems in a cloud has to be secure. Furthermore, virtualization paradigm in cloud computing results in several security concerns. For example, mapping the virtual machines to the physical machines has to be carried out securely. Data security involves encrypting the data as well as ensuring that appropriate policies are enforced for data sharing [11]. In addition, resource allocation and memory management algorithms have to be secure. Finally, data mining techniques may be applicable to malware detection in clouds.

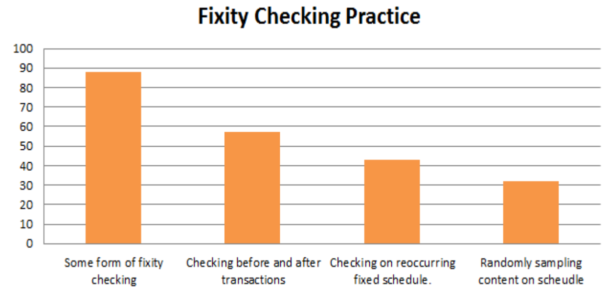

- 88% (49 of 56) of the organizations report that they are doing some form of fixity checking on content they are preserving.
- 57% (32 of 56) of the organizations are doing checks before and after transactions such as ingest.
- 43% (24 of 56) of the organizations are doing checks on some reoccurring fixed schedule.
- 32% (18 of 56) of the organizations are randomly sampling their content to check fixity.
- 46% (11 of 24) check fixity of content on at least a monthly basis.
- 21% (5 of 24) check fixity of content on at least a quarterly basis.
- 29% (7 of 24) check fixity of content on an annual basis.
- 4% (1 of 24) check fixity of content on a tri-annual basis.

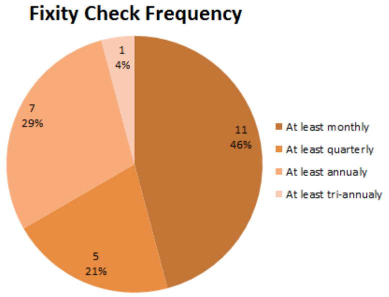

The rapid growth of the web, combined with growing interested in the cloud, has resulted in a competitive environment where developers and other stakeholders compete without adhering to a common set of standards for how browsers should function. More recently, social networks have emerged that vie to attract users to join their online communities. Many exploit features within the web browser to track user behaviour and target content based on browsing patterns [12]. Most of this tracking is done legitimately to attract new users and retain existing ones. There is, however, tracking that is done for nefarious purposes and it is revealing how much information can be obtained about a user based on a history of their browsing habits.

Cloud services and applications increasingly build atop one another in ever more complex ways, such as cloud based advertising or mapping services used as components in other, higher-level cloud-based applications, all of these building on computation and storage infrastructure offered by still other providers [13]. Each of these interacting, co-dependent services and infrastructure components is often implemented, deployed, and maintained independently by a single company that, for reasons of competition, shares as few details as possible about the internal operation of its services. The resource provisioning and moment-by-moment operation of each service is often managed by dynamic, reactive control processes that constantly monitor the behaviour of customer load, internal infrastructure, and other component services, and implement complex proprietary policies to optimize the provider's cost-benefit ratio [14].

Each cloud service's control loop may change the service's externally visible behaviour, in policy-specific ways, based on its neighbouring services' behaviour, creating cyclic control dependencies between interacting cloud services [15]. These dependency cycles may lead to unexpected feedback and instability, in much the way that policy-based routing in BGP is already known to lead to instability or "route flapping" in the much more restricted “control domain” of Internet routing To illustrate this risk, we consider a minimalistic, perhaps contrived, but hopefully suggestive example in following figure.

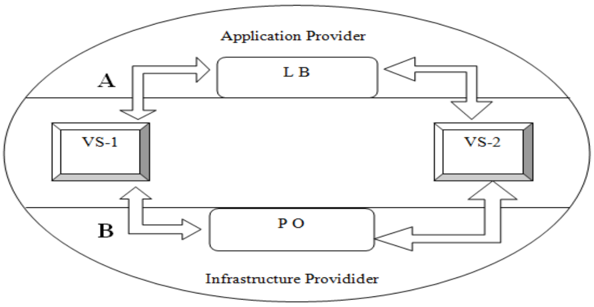

Application provider A develops and deploys a cloud-based application, which runs on virtual compute and storage nodes from infrastructure provider B. For simplicity, assume A leases two virtual nodes from B, and dynamically load-balances incoming requests across the web/application servers running on these nodes. Assume A’s load balancer operates in a control loop with a 1-minute period: after each minute it evaluates each server’s current load based on that server’s response time statistics during the past minute, and shifts more traffic during the next minute to the less-loaded server. Assume that A’s load shifting algorithms well-designed and stable assuming the servers in the pool behave consistently over time, like dedicated physical servers would.

Unbeknownst to A, however, suppose B also runs a control loop, which attempts to optimize the power consumption of its physical servers by dynamically adjusting the servers’ clock rates based on load. This control loop also happens to have a 1-minute period: after each minute, B’s controller measures each CPU core’s utilization during the past minute, then reduces the core’s voltage and speed if the core was underutilized or increases voltage and speed if the core was over utilized. Again, assume that B’s controller is well-designed and stable assuming that the servers’ load stays relatively constant or varies independently of B’s control actions.

Although both A’s and B’s control loops would be stable if operating alone, by the misfortune of their engineers (independently) picking similar control loop periods, the combination of the two control loops may risk a positive feedback loop. Suppose during one minute the load is slightly imbalanced toward virtual server 1, and the two control loops’ periods happen to be closely aligned; this will happen sooner or later in the likely event their clocks run at slightly different rates. A’s load balancer notices this and shifts some load away from the node in the next minute, while B’s power optimizer notices the same thing and increases the node’s voltage and clock speed. While either of these actions alone would lead toward convergence, the two in combination cause overcompensation: during the next minute, server 1 becomes more underutilized than it was over utilized in the previous minute. The two controllers each compensate with a stronger action—a larger shift of traffic back to server 1 by A and a larger decrease in voltage and clock speed by B—causing a larger swing the next minute. Soon all incoming load is oscillating between the two servers, cutting the systems overall capacity in half—or worse, if more than two servers are involved.

This simplistic example might be unlikely to occur in exactly this form on real systems—or might be quickly detected and “fixed” during development and testing—but it suggests a general risk. When multiple cloud services independently attempt to optimize their own operation using control loops that both monitor, and affect, the behaviour of upstream, downstream, or neighbouring cloud services, it is hard to predict the outcome: we might well risk deploying a combination of control loops that behaves well “almost all of the time,” until the emergence of the rare, but fatal, cloud computing equivalent of the Tacoma Narrows Bridge Comparable forms of “emergent misbehaviour” have been observed in real computing systems outside of the cloud context , and some work has studied the challenge of coordinating and stabilizing multiple interacting control loops, such as in power management.

Current approaches to solving or heading off such instability risks, however, generally assume that some single engineer or company has complete information about, and control over, all the interacting layers and their control loops. The cloud business model undermines this design assumption, by incentivizing providers not to share with each other the details of their resource allocation and optimization algorithms—crucial parts of their “secret sauce”—that would be necessary to analyse or ensure the stability of the larger, composite system.

III. CONCLUSION

Cloud computing is an emerging technology and It is important to measure the benefits of ubiquitous access against the concerns for privacy and preservation. This research article revealed various potentials for loss of privacy and corruption of data, along with methods for reducing or eliminating them. Further more work can be done to enhance the security and reliability of cloud computing with lesser threat of corruption or loss of data.

REFERENCES

[1] Amazon EC2 Crosses the Atlantic. http://aws.amazon.com/about aws/what new/2008/12/10 /amazon-ec2-crosses-the-atlantic/.

[2] Amazon S3 Availability Event: July 20, 2008. http://status.aws.amazon.com/s3-20080720.html.

[3] Amazon's terms of use. http://aws.amazon.com/agreement.

[4] An Information-Centric Approach to Information Security. http://virtualization.sys-con.com/node/171199.

[5] AOL apologizes for release of user search data. http://news.cnet.com/2100-1030_3-6102793.html.

[6] Armbrust, M., Fox, A., Griffith, R. et al. Above the Clouds: A Berkeley View of Cloud Computing. UCB/EECS- 2009-28, EECS Department, University of California, Berkeley, 2009.

[7] Ateniese, G., Burns, R., Curtmola, R., Herring, J., Kissner, L., Z., Peterson, and Song, D. Provable Data Possession at Untrusted Stores. In CCS. 2007.

[8] Fengzhe Zhang, Jin Chen, Haibo Chen, and Binyu Zang. CloudVisor: Retrofitting protection of virtual machines in multi-tenant cloud with nested virtualization. In 23rd SOSP, October 2011.

[9] Robert Gellman. Privacy in the clouds: Risks to privacy and confidentiality from cloud computing. World Privacy Forum, pages 1–16, 2009.

[10] Towards Secure and Dependable Storage Services in Cloud Computing Cong Wang, Student Member, IEEE, Qian Wang, Student Member, IEEE, Kui Ren, Member, IEEE, Ning Cao, Student Member, IEEE, and Wenjing Lou, Senior Member, IEEE- 2011.

[11] Cloud Security Alliance. (March 2012). Top Security Threads to Cloud Computing. [Online]. Available: https://cloudsecurityalliance.org/topthreats/csathreats.v1.0.pdf.

[12] J. Simth. (June, 2012). Benefits of Platform as a Service. [Online]. Available: http://www.hostreview.com/blog/120620-cloud-computing-with-benefits-of-paas-cloud-model.

[13] J. R. Vic Winkler, ―Securing the cloud,‖ Cloud Computing Security Techniques and Tactics, vol. 5, pp. 150-152.

[14] P. Mell and T. Grance, ―The NIST definition of cloud computing, National Institute of Standards and Technology. Retrieved 24 July 2011. Special publication 800-145.

[15] F. Gens, (September, 2008), Defining Cloud Services and Cloud Computing. [Online]. Available: http://www.cloudreviews.com/blog/what-is-hot-in-cloud-computing cloud computing