Editor: George Pallis • gpallis@cs.ucy.ac.cy

Security Challenges for the Public Cloud

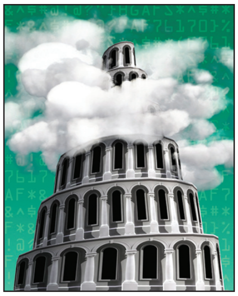

Kui Ren, Cong Wang, and Qian Wang • *Illinois Institute of Technology*

Cloud computing represents today's most exciting computing paradigm shift in information technology. However, security and privacy are perceived as primary obstacles to its wide adoption. Here, the authors outline several critical security challenges and motivate further investigation of security solutions for a trustworthy public cloud environment.

Cloud computing is the newest term for the long-dreamed vision of computing as a utility. The cloud provides convenient, on-demand network access to a centralized pool of configurable computing resources that can be rapidly deployed with great efficiency and minimal management overhead.[1] With its un- precedented advantages, cloud computing enables a fundamental paradigm shift in how we deploy and deliver computing services — that is, it makes possible computing outsourcing such that both individuals and enterprises can avoid committing large capital outlays when purchasing and managing software and hardware, as well as dealing with the operational overhead therein.[1]

Although cloud computing's benefits are tremendous, security and privacy concerns are the primary obstacles to wide adoption.[2] Because cloud service providers (CSPs) are separate administrative entities, moving to the commercial public cloud deprives users of direct control over the systems that manage their data and applications. Even if CSPs' infrastructure and management capabilities are much more powerful and reliable than those of personal computing devices, the cloud platform still faces both internal and external security and privacy threats, including media failures, software bugs, malware, administrator errors and malicious insiders. Noteworthy outages and security breaches to cloud services appear from time to time[2]: Apple's iPad subscriber privacy leak (http://techcrunch.com/2010/06/15/ipad-breach-personal-data/), Amazon S3's recent downtime (http://status.aws.amazon.com/s3-20080720.html), and Gmail's mass email deletions (www.techcrunch.com/2006/12/28/gmail-disaster-reports-of-mass-email-deletions) are all such examples.

Because users don't have access to the cloud's internal operational details, CSPs might also voluntarily examine users' data for various reasons without detection. Additionally, owing to hardware virtualization, multiple users can now share the same physical infrastructure, which runs their distinct application instances simultaneously. Although it increases resource utilization, this unique multitenancy feature also presents new security and privacy vulnerabilities for user interactions.[3] Hence, we argue that the cloud is intrinsically insecure from a user's viewpoint. Without providing a strong security and privacy guarantee, we can't expect users to turn control of their data and computing applications over to the cloud based solely on economic savings and service flexibility.

Here, we outline several critical security challenges, point out their importance, and motivate further investigation of security solutions that will help a trustworthy public cloud environment become a reality.

Data Service Outsourcing Security

As individuals and enterprises produce more and more data that must be stored and utilized (emails, personal health records, photo albums,

1089-7801/12/$31.00 © 2012 IEEE

tax documents, financial transactions, and so on), they're motivated to outsource their local complex data management systems to the cloud owing to its greater flexibility and cost-efficiency. However, once users no longer physically possess their data, its confidentiality and integrity can be at risk.

For the former concern, data encryption before outsourcing is the simplest way to protect data privacy and combat unsolicited access in the cloud and beyond. But encryption also makes deploying traditional data utilization services — such as plaintext keyword search over textual data or query over database — a difficult task. The trivial solution of downloading all the data and decrypting it locally is clearly impractical, due to the huge bandwidth cost resulting from cloud-scale systems. Moreover, aside from eliminating local storage management, storing data in the cloud serves no purpose unless people can easily search and utilize that data.

This problem on how to search encrypted data has recently gained attention and led to the development of *searchable encryption* techniques. At a high level, a searchable encryption scheme employs a prebuilt encrypted search index that lets users with appropriate tokens securely search over the encrypted data via keywords without first decrypting it. However, considering the potentially large number of on-demand data users and the huge amount of outsourced data files in the cloud, this problem is still particularly challenging because meeting performance, system usability, and scalability requirements is extremely difficult. In this context, numerous interesting yet challenging problems remain, including similarity search over encrypted data, secure ranked search over encrypted data, secure multikeyword semantic search, secure range query, and even secure search over nontextual data such as graph or numerical data.

Another important issue that arises when outsourcing data service to the cloud is protecting data integrity and long-term storage correctness. Although outsourcing data to the cloud is economically attractive for long-term, large-scale storage, it doesn't immediately guarantee data integrity and availability. This problem, if not properly addressed, can impede the successful deployment of a cloud architecture. Given that users no longer locally possess their data, they can't utilize traditional cryptographic primitives to protect its correctness.[4] Such primitives usually require a local copy of the data for integrity verification, which isn't viable when storage is outsourced. Furthermore, the large amount of cloud data and the user's constrained computing capabilities make data correctness auditing in a cloud environment expensive and even formidable. So, enabling a unified storage auditing architecture is important for this nascent cloud economy to become fully established; users will need ways to assess risk and gain trust in the cloud. From a system-usability viewpoint, such a design should incur very limited auditing overhead in terms of computation and bandwidth, incorporate cloud data's dynamic features, and preserve users' privacy when a specialized third-party auditor is introduced.[4]

Beyond storage correctness, other security issues arise related to cloud storage services. One noteworthy security notion is *proof of ownership*.[5] This technique aims to prevent the exposure of user data via the side-channels that results from cross-user de-duplication, which is widely used to save the space and bandwidth CSPs require. Other challenging security problems include assured data deletion and remote assessment of fault tolerance — that is, the remote detection of hard-drive failure vulnerabilities in the cloud.[6]

Computation Outsourcing Security

Another fundamental service enabled within the cloud paradigm is computation outsourcing. By outsourcing workloads to the cloud, users' computational power is no longer limited by their resource-constrained devices. Instead, they can enjoy the cloud's literally unlimited computing resources in a pay-per-use manner without committing any large capital outlays locally.

However, current outsourcing practice operates in plaintext — that is, it reveals both data and computation results to the commercial public cloud. This can raise big security concerns, especially when the outsourced computation workloads contain sensitive information, such as a business's financial records, proprietary research data, or even personally identifiable health information. Furthermore, the cloud's operational details aren't transparent enough to users. Consequently, various motivations can cause the cloud to behave unfaithfully and return incorrect results. These range from possible software bugs, hardware failures, or even outsider attacks to cloud servers deliberately being "lazy" to save computational costs. Thus, we're in great need of secure computation outsourcing mechanisms to both protect sensitive workload information and ensure that the computation results returned from the cloud are correct. This task is difficult, however, due to several challenges that the mechanism design must meet simultaneously. First, such a mechanism must be practically feasible in terms of computational complexity. Otherwise, either the user's cost can become prohibitively huge, or the cloud might not be able to complete the outsourced computations in a reasonable amount of time. Second, it must provide sound security guarantees without restricting system assumptions. Namely, it should strike a good balance between security guarantees and practical performance. Third, this mechanism must enable substantial computational savings at the user side compared to the amount of effort required to solve a problem locally. Otherwise, users have no reason to outsource computation to the cloud.

A recent breakthrough in *fully homomorphic encryption* (FHE) has shown the general results of secure computation outsourcing to be viable in theory. But applying this general mechanism to everyday computing tasks is still far from practical due to FHE operations' extremely high complexity, which can't yet be handled in practice. On a different front, researchers are working on mechanisms for specific computation outsourcing problems, such as linear programming via problem transformation,[7] genomic computation via specialized computation partition,[8] and efficient verification of large-scale biometric computations, all of which should provide much more practical efficiency than the more general solutions currently available.

Access Control

In many application scenarios, such as those in enterprises or organizations, users' access to data is usually selective and highly differentiated. Different users enjoy different access privileges with regard to the data. When data are outsourced to the cloud, enforcing secure, efficient, and reliable data access among a large number of users is thus critical.

Traditionally, to control the dissemination of privacy-sensitive data, users establish a trusted server to store data locally in clear, and then control that server to check whether requesting users present proper certification before letting them access the data.[9] From a security standpoint, this access control architecture is no longer applicable when we outsource data to the cloud. Because data users and cloud servers aren't in the same trusted domain, the server might no longer be fully trusted as an omniscient reference monitor[9] for defining and enforcing access control policies and managing user details. In the event of either server compromise or potential insider attacks, users' private data might even be exposed.

One possible approach to enforce data access without relying on cloud servers could be to encrypt data in a differentiated manner and disclose the corresponding decryption keys only to authorized users. This approach usually suffers from severe performance issues, however, and doesn't scale, especially when a potentially large number of on-demand users desire fine-grained data access control. Researchers have been working on how to realize a fine-grained access control design that fully leverages the cloud's computation resource richness.[9] Via this approach, data users would be able to securely delegate to the cloud most cumbersome user/data management workloads — such as handling frequent user access privilege updates in large dynamic systems — while still preserving the underlying data confidentiality against any unauthorized access.

Trustworthy Service Metering

As computing as a service increases in popularity, users employ cloud resources as a public utility to accomplish their tasks. To make the service profitable, CSPs charge users according to the resources they consume. For example, the Amazon Elastic Compute Cloud (EC2) charges users based on the time that their specified EC2 instances are in a running state, while Google AppEngine charges on the basis of how many CPU cycles a user application consumes. However, because users might have little or no visibility into the cloud infrastructure, they're often unable to directly connect their actual cloud resource consumption and the usage charges.[10] For example, although hardware virtualization lets multiple users run tasks on the same infrastructure without explicitly interfering with each other, many shared resources, such as memory, I/O, and network bandwidth, can't be perfectly isolated. Consequently, CSPs might incorrectly apply unexpected costs to a user's usage report when the true culprit might be possible software bugs or network congestion caused by other users running tasks on the same physical infrastructure.

So, how to guarantee service metering's trustworthiness is of critical importance if the utility-based computing paradigm is to be successful. A unified mechanism for securely and fairly measuring resource consumption is greatly needed and will benefit both cloud users and CSPs. First, a trustworthy service-metering mechanism will let users obtain verifiable guarantees on the amount of cloud resources they actually consume and thus help them trust the cloud more easily. Second, such a mechanism can serve as an indispensable arbitration means to resolve any unsettled dispute over charges between cloud users and a CSP. Third, an unbiased and independent service auditor can further employ such a mechanism to audit and quantify the quality of service the CSP promises in its service-level agreement, ensuring that the utility-computing-oriented service model is economical. Finally, trustworthy service metering will encourage more cloud adoption, leading to increased revenue for CSPs owing to improved overall resource utilization.

Multitenancy Security and Privacy

Multitenancy is an essential attribute of cloud computing.[1] To optimize resource utilization, CSPs often use hardware virtualization to hide a computing platform's physical characteristics. This lets multiple users run their distinct application instances simultaneously on the same physical infrastructure without seeing each other's data. Multitenancy increases use of the underlying hardware resources and, with virtualization, eases the management burden for CSPs, allowing for efficient and effective resource provisioning and re-allocation without the need for any upfront hardware purchase or setup.

Despite its benefits, this multitenant cloud environment also presents severe security threats and privacy vulnerabilities to both the cloud infrastructure and cloud users. Virtualized environments share similar functionalities with existing operating systems and applications in the physical environment, so software bugs and newly identified security vulnerabilities in these systems remain the primary threat to any virtualized multitenant environment. Considering the scale of cloud systems, the potential threat from these security risks can be even bigger compared to that for a nonvirtualized computing environment. Furthermore, for resource management in the cloud, different virtualized application instances must be constantly provisioned, allocated, or even migrated between multiple physical machines. Consequently, such dynamic features in the multitenant environment further exacerbate the problem's complexity and make achieving and maintaining consistent security difficult.[2]

Multitenancy also opens doors for potential privacy leaks. As mentioned previously, side-channel attacks present new risks to cloud users' information in the multitenant environment. In a recent study, researchers used engineering techniques to infer the virtualized resource allocation strategy from CSPs and successfully placed their virtualized application instance on the same physical machine as the target victim. They were then able to extract the victim's private information through traffic patterns and other side-channel information.[3] These results show that even in a strongly isolated multitenant environment, this new type of privacy leak is still a possible threat.

Multitenancy security and privacy is one of the critical challenges for the public cloud, and finding solutions is pivotal if the cloud is to be widely adopted. However, little work exists today that not only addresses these problems but also consistently and scalably maintains this dynamic computing environment's scalability.

Security Overhead and More

Although designing security into the cloud benefits users and CSPs, it inevitably increases overhead for both. For users in particular, such overheads could offset the cloud's economically appealing benefits and might conflict with their reasons for using the cloud in the first place. How to quantitatively explore the trade-offs between security overhead and cloud benefits is another interesting but important problem. Any solution to this question will help users make better-informed decisions before moving to the cloud.

We've described several critical security challenges for the commercial public cloud, but our list is by no means comprehensive. For example, although cloud computing provides literally unlimited computation powers while reducing costs, how to prevent malicious cloud users from abusing cloud resources is still an issue. Such abuses could include password/key cracking, malicious data hosting, or botnet command and control. Adopting stricter monitoring of cloud resource usage could be one way to mitigate this concern, but it's inevitably in conflict with legal users' privacy rights. Hence, new research is needed.

Security and privacy is one fundamental obstacle to cloud computing's success. In this context, we've discussed several critical security challenges that current research thrusts aren't yet addressing. This article is intended as a call for action to motivate further investigation of the many challenging security issues that will impact the public cloud's future. Clearly, much work remains for a trustworthy public cloud environment to become a reality.

Acknowledgments

This work is supported in part by the US National Science Foundation under grants CNS-1117111 and CNS-1054317, and by an Amazon Web service research grant.

References

1. P. Mell and T. Grance, "The NIST Definition of Cloud Computing," US Nat'l Inst. of Science and Technology, 2011; http://csrc.nist.gov/publications/nistpubs/800-145/SP800-145.pdf.
2. "Security Guidance for Critical Areas of Focus in Cloud Computing," Cloud Security Alliance, Dec. 2009; https://cloudsecurityalliance.org/csaguide.pdf.
3. T. Ristenpart et al., "Hey, You, Get Off of My Cloud! Exploring Information Leakage in Third-Party Compute Clouds," *Proc. 16th ACM Conf. Computer and Communications Security* (CCS 09), ACM Press, 2009, pp. 199–212.
4. C. Wang et al., "Privacy-Preserving Public Auditing for Storage Security in Cloud Computing," *Proc. 30th IEEE Int'l Conf. Computer Communications* (INFOCOM 10), IEEE Press, 2010, pp. 525–533.
5. S. Halevi et al., "Proofs of Ownerhip in Remote Storage Systems," *Proc. 18th ACM*

Conf. Computer and Communications Security (CCS 11), ACM Press, 2011, pp. 491–500.

6. K. Bowers et al., “How to Tell if Your Cloud Files Are Vulnerable to Drive Crashes,” *Proc. 18th ACM Conf. Computer and Communications Security* (CCS 11), ACM Press, 2011, pp. 501–514.
7. C. Wang, K. Ren, and J. Wang, “Secure and Practical Outsourcing of Linear Programming in Cloud Computing,” *Proc. 31st IEEE Int’l Conf. Computer Communications* (INFOCOM 11), IEEE Press, 2011, pp. 820–828.
8. R. Wang et al., “Privacy-Preserving Genomic Computation through Program Specialization,” *Proc. 16th ACM Conf. Computer and Communications Security* (CCS 09), ACM Press, 2009, pp. 338–347.
9. S. Yu et al., “Achieving Secure, Scalable, and Fine-Grained Access Control in Cloud Computing,” *Proc. 30th IEEE Int’l Conf. Computer Communications* (INFOCOM 10), IEEE Press, 2010, pp. 534–542.
10. M. Liu and X. Ding, “On Trustworthiness of CPU Usage Metering and Accounting,” *Proc. 1st Int’l Workshop Security and Privacy in Cloud Computing* (ICDCS-SPCC 10), IEEE Press, 2010, pp. 82–91.

Kui Ren is an assistant professor in the Illinois Institute of Technology’s Electrical and Computer Engineering Department. His research expertise includes cloud computing and security, wireless security, and smart grid security. Ren has a PhD in electrical and computer engineering from Worcester Polytechnic Institute. He’s a recipient of the US National Science Foundation Faculty Early Career Development (CAREER) Award, and is a senior member of IEEE and a member of the ACM. Contact him at kren@ece.iit.edu.

Cong Wang is a PhD student in the Illinois Institute of Technology’s Electrical and Computer Engineering Department. His research interests are in applied cryptography and network security, with a focus on secure data service outsourcing in cloud computing. Wang has an ME in communication and information systems from Wuhan University. He’s a student member of IEEE and the ACM. Contact him at cong@ece.iit.edu.

Qian Wang is a PhD student in the Illinois Institute of Technology’s Electrical and Computer Engineering Department. His research interests include wireless network security and privacy, and cloud computing security. Wang has an MS in communication and information systems from the Shanghai Institute of Microsystems and Information Technology, Chinese Academy of Sciences. He’s a co-recipient of the Best Paper Award from IEEE ICNP 2011, and is a student member of IEEE. Contact him at qian@ece.iit.edu.

cn *Selected CS articles and columns are also available for free at http://ComputingNow.computer.org.*