What Does the Future Hold for Cloud Computing?

David Zage, Sandia National Laboratories
Dustin Franklin, Sandia National Laboratories
Vincent Urias, Sandia National Laboratories

Abstract. The topic of cloud computing continues to generate significant interest among information technology and government decision makers even as they hesitate to adopt cloud solutions due to security concerns. The authors define the risks associated with various cloud deployment models and identify key solutions that can create easy-to-use, secure cloud deployments.

1. Introduction

Even though recent reports have begun to forecast diminished interest in cloud computing [16], large numbers of services are still migrating to the cloud and further infrastructure is being dedicated to platforms and solutions. While a large amount of cloud research has focused on utilizing the power, flexibility, and potential cost savings of cloud computing platforms, reports such as the Department of Homeland Security Roadmap for Cybersecurity Research [9] and previous research [4, 13] have expressed the explicit need for continued security analysis of cloud computing solutions. In polls, over 70% of government decision-makers [2] and 80% of IT executives [3, 14] identify security and ease of deployment as the primary obstacles to cloud computing adoption.

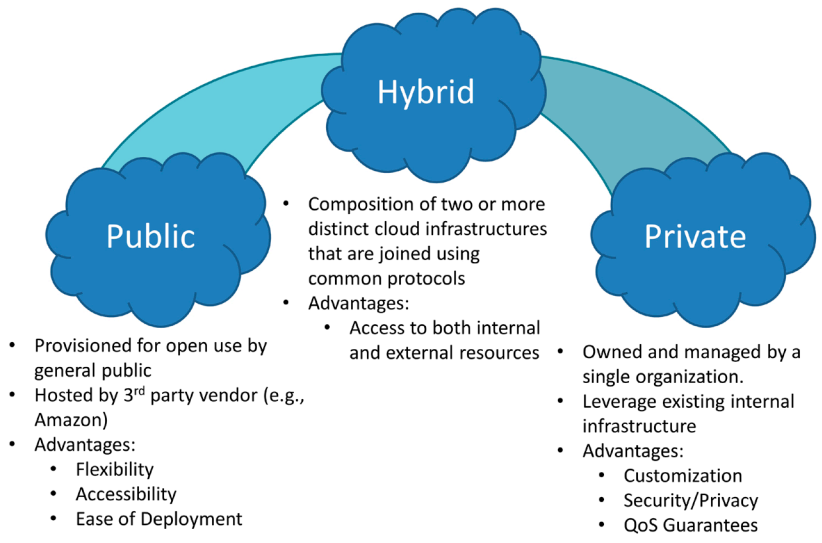

Figure 1 - Three deployment models for cloud computing. General descriptions and the advantages of each type are listed below each model.

Cloud services deal with amounts of data, users, and service heterogeneity that have never been seen before. These issues combine with the desired ubiquity of cloud accessibility to create a field that is ripe for vulnerabilities and a potential playground for adversaries. There has been work towards securing cloud deployments, but security is still typically an afterthought as companies focus on maintaining service availability. Many security solutions are ports of classic paradigms such as firewalls to web services. These ports enhance the security of services running on the cloud, but they do not increase the intrinsic security of the cloud service itself. This article analyzes the various deficiencies that continue to hinder cloud deployments and presents three key areas of work fundamental to improving cloud services. The discussion of these threats and solutions is colored by the authors experience in creating and using large-scale cloud infrastructures.

Cloud Deployment Models

In the field of cloud computing, three cloud deployment models seen in Figure 1 have emerged: 1) public, 2) private, and 3) hybrid clouds. A public cloud deployment is typified by the hardware and cloud components being hosted by a third party provider and the cloud being used by multiple users. A user has no control over the hardware layer and varying levels of control over other components of the cloud stack seen in Figure 2. Public cloud deployments were developed to optimize the cost of computation and storage and allow massive computing jobs to be performed for a fraction of their former costs. This cost savings typically comes at the expense of security; thus causing increasing interest in private cloud deployments. Private cloud deployments are owned and managed by a single organization. For example, a company's IT organization may choose to deploy a Infrastructure-as-a-Service (IaaS) cloud usable by anyone in the company. A private deployment enables the owners (e.g., corporate IT) and users to have control over the full stack of cloud system components.

The next logical evolution in cloud deployments leverages both the cost savings of public clouds and the potential security gains of private clouds by combining them into a hybrid cloud service. We believe this is the future of cloud computing, but will highlight issues that must be taken into account before deploying a hybrid cloud. While these models are designed to handle many of the same tasks and thus share a common set of threats, there are also security challenges unique to each due to the exposure faced by components of their respective infrastructure. We now look at threats common to all cloud deployments.

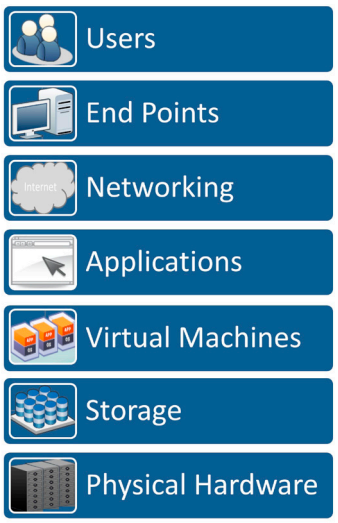

Figure 2 - The operational stack for typical cloud deployments.

2. Common Threats to Cloud Computing

The utility of cloud computing must be weighed against the threats it faces, which fall into three categories:

a. Failure to maintain security
b. Loss of availability
c. Reduction in usability

Security issues typically result from an adversary attempting to acquire data, knowledge, or persistent access to a system. Losses in availability stem from an adversary trying to deny legitimate user access to a cloud for the purposes of annoyance, delaying progress of work, or interfering with real-time responses to critical situations. Reductions in usability can affect the continued viability of a cloud service and the cloud paradigm as a whole. An adversary can cause systematic faults in a cloud service and ultimately diminish a user's faith in that cloud service. The threats affect the three deployment models universally and no one model stands out as being inherently better than the others.

Failure to Maintain Security

Security failures are the most conspicuous threats to cloud computing. Journalists regularly detail significant losses of data integrity and confidentiality originating from targeted attacks. Kapersky Lab reported that in 2013, 35% of businesses have lost data due to flawed system security [6]. The widespread availability of internet access has made hacking a global enterprise, allowing adversaries to work in areas where they face minimal penalties and have significant incentives.

Cloud deployments are complex systems of networked components that must work seamlessly in order to secure the large amounts of information they contain. A single misconfiguration or an unpatched vulnerability is sufficient to lead to exploitable security holes and allow an adversary access to the entire infrastructure. For example, in 2011, Sony Entertainment was the victim of a series of attacks on various pieces of their infrastructure in which personal information, including credit card information, was stolen from millions of accounts [11]. Exact details of the attack vector are unpublished, but the seriousness of the breach became apparent when the attackers released stolen data that indicated that Sony had been storing user information and passwords in plain text [8]. For their lack of security in credit card processing, Sony faced many repercussions, including a £250,000 fine by UK's Information Commissioner's Office. Although it is easy to blame the attackers, companies with valuable information need to be cognizant of the threats they face.

Including users into the chain of trust in cloud deployments has caused many security vulnerabilities. Good user security practices such as enforcing strong password selection, avoiding spearfishing, and testing web interfaces for cross site scripting attacks are necessary. While many of these vulnerabilities and practices are well known, it is important to make note of them as they continue to impact cloud deployments.

Loss of Availability

No matter the type of cloud, a user wants data to be accessible at any time and place. Availability is reduced as networks and machines fail, poorly deployed cloud solutions run into bottlenecks, and cloud services face distributed denial of service (DDoS) attacks. This is a particularly pressing issue for public clouds as they provide Service Level Agreements (SLAs) that are based on their contractual ability to provide computing power, storage, and services. Large sums of money stand to be lost when providers fail to meet their SLAs.

A concern that bridges both security and availability is the capability of monitoring the state of data during its lifetime in the cloud. Problems in data provenance include determining if data resides in locations that follow the same regulations a business must enforce (e.g., HIPAA), if the cloud service has stored the entirety of the data, if the data remains uncorrupted, and if the data is truly removed once it has been deleted by the user. A user that uploads his or her data to a third party may be forfeiting inherent rights to the control of their data, which could then be changed or viewed without notifying the user [5]. While this concern may appear easier to manage in private cloud deployments, cloud solutions are often adopted without understanding the full risk profile. IT typically lacks the tools necessary to understand how the complex pieces fit together and can provide minimal assurances for the end-user [14].

Reduction in Usability

If the expected performance of a cloud service is not up to a user's expectations or its SLA guarantees, the user may change or discontinue usage. Repeated negative experiences result in the user losing faith in the cloud computing paradigm. Additionally, corporations desire to have the ability to quickly deploy private and hybrid clouds with minimal effort (e.g., testbed-as-a-service), but the technology for doing this in an efficient, repeatable manner is still not mature enough for this to be a reality.

3. Threats to the Different Cloud Deployment Types

Given the general threats discussed in Section II, this section looks at cloud deployment-specific vulnerabilities.

Public Cloud Deployments

Common to all public cloud systems is the lack of control over the physical storage of data. In clouds which give the user minimal access to the cloud stack, such as software-as-a-service (SaaS) clouds, security of the data is reliant almost entirely on the practices used by the cloud provider, over which the user has little purview. In less restrictive systems, such as IaaS clouds, users are allowed to create and deploy virtual machines, which provide greater user customization and control of data storage. Several cloud service providers provide preconfigured virtual machines for their users to minimize user effort and eliminate obvious security flaws.

While virtualization is a great enabling technology, there have been actual attacks demonstrated where a virtual machine can be compromised and used to bypass system protections, enabling attacks on the rest of the cloud infrastructure [7, 18]. Although cloud providers do not provide information on other clients running on the same physical hardware, adversarial virtual machines can be used to find and attack specific services [15]. It should also be noted that cloud providers have minimal incentive to provide secure virtual machines. One could imagine cases

in which unscrupulous providers may distribute virtual machines containing flaws (e.g., networking issues) that result in extra computation and network usage, simultaneously earning money for the provider while costing the customer.

There are also large marketplaces where third-party companies and users exchange virtual machines. Users need to be wary of these virtual machines as they are often poorly secured [12, 1]. These preconfigured solutions have been found to contain unpatched code, share credentials with other virtual machines, and, worst of all, even Trojan software. Another less obvious problem with virtual machines that users must be aware of is they often lack the necessary randomness to create truly secure cryptographic keys, especially when the virtual machines are running on the same hardware.

Users of public clouds may suffer from attacks that are not directed at them. Any network outage that affects the cloud will restrict the cloud usability. Clouds regularly face DDoS attacks attempting to bring down a single service hosted on the cloud. When these attacks succeed, they have the side effect of affecting all of the other services hosted by the cloud provider. Innocent users of the cloud can expect that a successful attack will cause downtime for their services even though they are not the target. One potential solution to dealing with such issues is through the use of hybrid solutions, enabling at least partial data access even during downtime.

Private Cloud Deployments

The primary driver behind private cloud adoption is concern over the sensitivity of the data to be stored and processed. Businesses desire cloud solutions that can leverage excess internal capacity while minimizing the potential for data leakage. Additionally, internally run clouds can have significant advantages in availability and accessibility over public deployments.

Many companies and governmental groups are constructing private cloud infrastructures inside their network perimeter. This setup can often be easier (and more comforting) for them to deploy as it uses many traditional system security mechanisms. For example, existing web security (e.g., firewalls) and permission management infrastructure can be used to secure the system. While these mechanisms provide protection from outside an organization, they are not configured to protect resources from mismanagement and malicious insiders. It is extremely unlikely that every piece of data should be available to every department and all people. Relying entirely on access control at the network perimeter can lead to dissemination of data to an adversary that has entered the network through another route.

Private cloud deployments are also vulnerable to attacks designed at interrupting availability, such as DDoS attacks. Public cloud providers can prepare for such attacks by investing in redundant capacity that scales to handle excessive traffic as needed. A private cloud will not have the same growth capability or mitigation techniques and the system may fail when targeted, leading to unavailability and system downtime.

Ultimately, the security of a private cloud depends on the capabilities of the organization deploying it. Currently, the deployment procedures for clouds are opaque, with multiple services running and numerous layers of abstraction between each of the services and between the services and their administrative layers. Often, when configuring and using services like OpenStack, one of the numerous web services and authentications will fail silently. While there has been significant work done by the community to improve the deployment and administration of tools like OpenStack (such as using common configuration management/deployments tools like Puppet), there are really no standard solutions. Differences in networks (such as topology, IP Space, VLANs, etc.), in hardware (vendor, raids, etc.) and underlying virtualization tools all provide complexity and variation from the norm. When a failure does occur, determining where the failure occurred and why is similar to finding a needle in a haystack. Typically, a system administrator will embark on a debugging mission that might result in a functioning system or attempt to start over with alternative configurations. Currently, there are no tools that can give an administrator the data and insight into where/what might have failed.

Hybrid Cloud Deployments

Hybrid cloud deployments have the potential to offer many of the positive aspects of both public and private cloud deployments in a single service, but they also face unique challenges. A primary concern is understanding the composability and resulting security posture of the hybrid system. Given a secure private and public cloud deployment, (provably) aggregating these together to create a secure service is currently an open problem. Most hybrid solutions are joined by easier-to-understand higher level protocols (e.g., user programs) and not at lower levels (e.g., a cross-cloud database). Clearly identifying the interfaces and connectivity patterns between the public and private components is a critical first step towards creating a secure service. Not only must the security of the system be analyzed, mitigation plans for availability or security issues affecting either portion of the cloud must be in place.

Another major concern not present in other deployment scenarios is the accurate disseminate and tracking of data between the multiple components of the cloud. While it might be attractive to use the private portion of the cloud to store HIPAA data and the public for non-sensitive data, it must be ensured that the data will not commingle. Having a write-once, read-only cloud is of little use.

A final challenge that must be addressed in the creation of hybrid solutions is configuration management. While the potential to have heterogeneous solutions is beneficial in reducing dependence on any one piece of software, the cloud services must be easy to set up. This necessitates the fusing of data from both the public and private cloud to create a common interface for deployment and management.

4. Enhancing the Security of Cloud Deployments

In order to mitigate some of the security issues discussed, we present three promising solutions: A) enhanced deployment techniques that are automated and repeatable, B) full stack cloud introspection, and C) enhanced cloud storage solutions leveraging multiple providers. Each of the solutions mitigates a distinct security vulnerability that is found in the cloud infrastructure. They can provide much higher levels of confidence

in the security posture of the infrastructure at the cost of some additional management challenges.

Enhanced Deployment

If an organization deploys any of the cloud models discussed previously (e.g., a private cloud for internal use or a public cloud for commoditization and profit), the organization must understand how to construct and administer the entire cloud stack. A very basic cloud service install which an administrator would have to create might resemble Figure 3, with installation occurring from left to right. Creating such a procedure is difficult and must be streamlined and instrumented with greater amounts of understanding/introspection into the install process. Not only will this enable greater cloud adoption, it will also give administrators the ability to quickly set up and tear-down cloud deployments for research and testing. We have created a set of installers for OpenStack that coalesce an immense amount of logging (from the system, network, and applications) to a central location. We have developed tools for automated install analysis as well as a platform to begin understanding how and why the system fails.

Our research points to the creation of a deployment process resembling object-oriented design patterns, in which the interactions and required functionality between each phase of the install are predefined. This way, a component in any step can be simply exchanged for another which provides the desired functionality. For instance, we have created automated installers for the major hypervisors that can be easily interchanged. In this manner, we can construct standard cloud configurations and take the uncertainties out of deployment. This allows for the creation of standard secure builds that can be vetted, tested, and guaranteed to produce repeatable results. Using this work, we can go from the bare metal to the fully operational application stack that can be re-provisioned in under an hour on tens of different hardware variants.

Cloud Analysis

While clouds are complex, one of the potential advantages of cloud-based computing is that it opens the possibility of understanding the entire infrastructure. This understanding comes down to intelligently gathering and processing a myriad of system packets and logs. Each component of the cloud stack, from the bare metal operating system, to the virtual machine manager, to the application being hosted by the virtual machine synthesizes logs that need to be analyzed. If this flood of information can be efficiently aggregated and correlated, this enables an administrator to understand the context of the applications, develop situational awareness, and leverage this awareness for detection and prevention. One potential avenue currently being explored for creating analyzable cloud infrastructure is by instrumenting logs in each level of the system and capturing them in a security information and event management (SIEM) solution such as Splunk. The SIEM solution will interrogate each service and aggregate the information, allowing for easy visualization of data and trends. Figure 4 is an example of the analysis framework we are investigating. With this framework, we have been able to quickly triage hardware and application failures as well as provide a record of the events on the system.

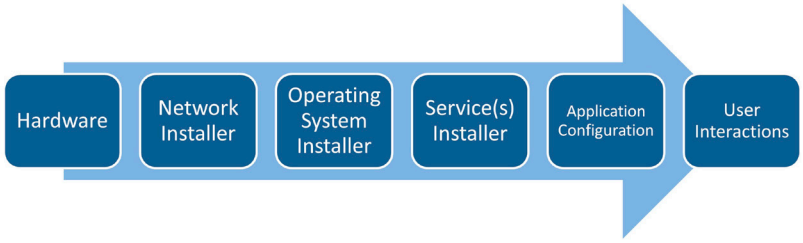

Figure 3 - Example cloud installation stack. As indicated by the arrow, the installation process flows from the hardware on the left to the final step of setting up the system for user interaction on the right.

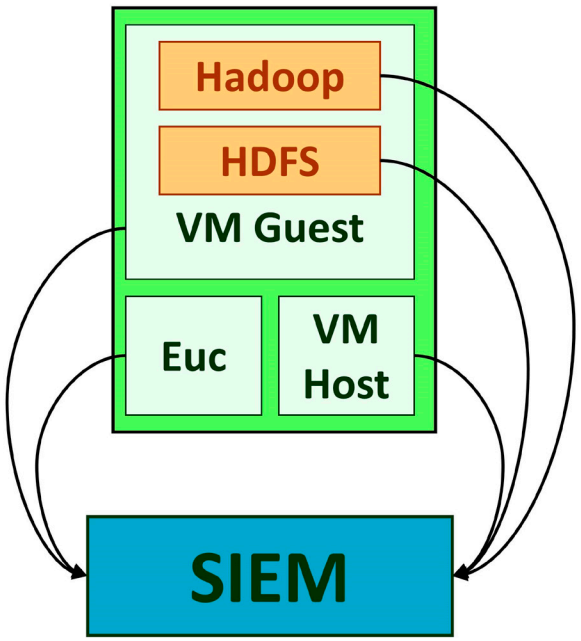

Figure 4 - System diagram of a full cloud analysis solution leveraging security information and event management solutions.

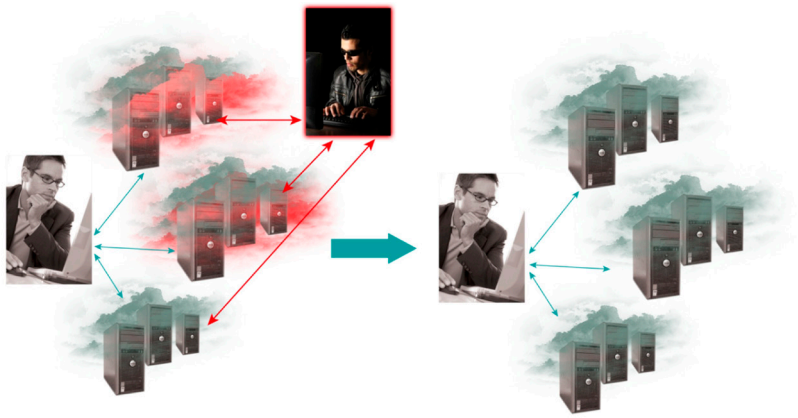

Figure 5 - The security of the cloud service should render the cloud usable to a legitimate user even when under attack.

While a common log area and visualizations are of great utility to an administrator, these are only a first step to better understand a cloud. Automated analysis techniques such as outlier identification and (un)supervised machine learning techniques can be built on top of this data, allowing for near real-time, less manually intensive identification of problems and security concerns. Also, much of this information would be useful to the end-user of the cloud. A critical research challenge for the future is enabling an end-user to leverage these logs in a secure manner.

Improved Cloud Storage

The third area we see as critical to the continued success of cloud computing is the continued development of improved data storage protocols. The concept of the cloud has been great for monetizing computing capabilities and a provider's return on investment has been tightly coupled with availability. This has typically left security as an afterthought or the burden has been placed on the user to ensure the confidentiality and integrity of their data. As seen in Figure 5, users need storage solutions that can seamlessly integrate multiple heterogeneous storage services (e.g., a local cloud storage service and Amazon S3) while providing the security the user needs and expects. Even with the system under attack, the user should be able to experience it as if the environment was benign.

One of the areas we are currently exploring is the use of wheat and chaff (W&C) storage. W&C uses multiple algebraic operations of linear subspaces to encode and replicate data. By encoding data in large finite fields, we create solutions which offer the end-user provable data confidentiality and integrity and provide lightweight checks on the user's data-related service level agreement. As part of this research, a completely non-preferential dynamic partitioning system was developed utilizing online codes [10] that allows for maximal robustness when splitting data between multiple cloud providers. This total system can provide the end-user with informed trade-offs between cost, performance, and security. This functionality is critical for the continued growth of hybrid solutions. For more information, see [17].

5. Looking Towards Continued Adoption

We believe the future of cloud computing rests in the opportunities and challenges present in hybrid cloud deployments. These allow organizations to have better resiliency to failure, establish data models for multiple types of data (i.e., increased privacy for data that remains in a private infrastructure), and optimize cost and resource usage by utilizing the appropriate cloud offerings. The solutions we present and continued work in the areas of creating automated cloud deployments, improved full cloud management, and secure storage will mitigate new cloud challenges before they become problematic.

Disclaimers:

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000

ABOUT THE AUTHORS

David Zage is a Principle Member of Technical Staff at Sandia National Laboratories. He received his Ph.D. in computer science from Purdue University. He has worked on various research topics including computer security, distributed systems, fault-tolerant protocols, routing, and wireless mesh networks.

E-mail: djzage@sandia.gov

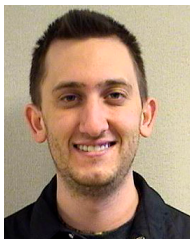

Dustin Franklin is a Graduate Student Intern at Sandia National Laboratories while he studies as a Ph.D. candidate in the computer science department at The University of New Mexico. His interests include distributed systems and adaptive intrusion response systems.

E-mail: drfrank@sandia.gov

Vincent Urias is a Principle Member of Technical Staff at Sandia National Laboratories, where he has spent the last ten years conducting cyber security research and development. His research areas include cyber testbedding, cyber modeling and simulation, as well as cyber analytics, cloud computing, and networking.

E-mail: veuria@sandia.gov

REFERENCES

1. Nelson Elhage. Virtunoid: Breaking out of kvm. Technical report, Black Hat USA, 2011.
2. Federal Computer Week. Cloud computing snapshot. Technical report, Federal Computer Week, 2010.
3. Frank Gens. New IDC IT cloud services survey: Top benefits and challenges. Technical report, IDC eXchange, 2009.
4. Nils Gruschka and Meiko Jensen. Attack surfaces: A taxonomy for attacks on cloud services. In Proceedings of IEEE CLOUD, 2010.
5. Jim Dempsey. <https://www.cdt.org/personnel/jim-dempsey> 2012.
6. Kaspersky Labs. Global it security risks: 2012. <http://www.kaspersky.com/downloads/pdf/kaspersky_global_it-security-risks-survey_report_eng_final.pdf> 2012.
7. Kostya Kortchinsky. Cloudburst: A vmware guest to host escape story. Technical report, Black Hat USA, 2009.
8. Mathew Schwartz. Sony hacked again, 1 million passwords exposed. <http://www.informationweek.com/security/attacks/sony-hacked-again-1-million-passwords-ex/229900111> 2011.
9. Doug Maughan et al. A roadmap for cybersecurity research. Technical report, Department of Homeland Security Science and Technology Directorate, 2009.
10. Petar Maymounkov. Online codes. Technical report, Technical report, New York University, 2002.
11. R McMillan. Sony cuts off sony online entertainment service after hack. Computer World, 2011.
12. Haroon Meer, Nicholas Arvanitis, and Marco Slaviero. Clobbering the cloud. In Black Hat USA, 2009.
13. David Molnar and Stuart Schechter. Self hosting vs. cloud hosting: Accounting for the security impact of hosting in the cloud. In Proceedings of WEIS, 2010.
14. Intel Research. Whatï¿½s holding back the cloud? Technical report, Intel, 2012.
15. Thomas Ristenpart, Eran Tromer, Hovav Shacham, and Stefan Savage. Hey, you, get off of my cloud: exploring information leakage in third-party compute clouds. In Proceedings of ACM CCS, pages 199–212. ACM, 2009.
16. David Mitchell Smith. Hype cycle for cloud computing. Technical report, Gartner's, 2011.
17. David Zage and James Obert. Utilizing linear subspaces to improve cloud security. In Proceedings of IEEE/IFIP DSN-W, 2012.
18. Yinqian Zhang, Ari Juels, Michael Reiter, and Thomas Ristenpart. Cross-vm side channels and their use to extract private keys. In Proceedings of ACM CCS, pages 305–316. ACM, 2012.