Cloud Computing Security Issues

V.I. Guzhov, K.O. Bazhenov, S.P. Ilinykh, A.R. Vagizov
Novosibirsk State Technical University

***Annotation*: Cloud computing is a fundamental shift from traditional client/server or tier architecture which lays emphasis on effective utilization of IT infrastructure, reduction in operational cost and optimum customer satisfaction thus helping the enterprise in increased profits with satisfied customers. There are many security issues than the traditional "authentication & authorization, integrity, consistency, backup & recovery" which might turn out to be abuse of technology, if not considered.**

***Keywords*: Cloud computing, security, privacy.**

I. Introduction

Cloud computing has been defined by NIST as 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 cloud provider interaction [1]. Cloud computing technologies can be implemented in a wide variety of architectures, under different service and deployment models, and can coexist with other technologies and software design approaches. The security challenges cloud computing presents, however, are formidable, especially for public clouds whose infrastructure and computational resources are owned by an outside party that sells those services to the general public.

II. About Cloud Computing

In addition to the usual challenges of developing secure IT systems, cloud computing presents an added level of risk because essential services are often outsourced to a third party. The externalized aspect of outsourcing makes it harder to maintain data integrity and privacy, support data and service availability, and demonstrate compliance.

In effect, cloud computing shifts much of the control over data and operations from the client organization to their cloud providers, much in the same way organizations entrust part of their IT operations to outsourcing companies. Even basic tasks, such as applying patches and configuring firewalls, can become the responsibility of the cloud service provider, not the user. This means that clients must establish trust relationships with their providers and understand the risk in terms of how these providers implement, deploy, and manage security on their behalf. This trust but verify relationship between cloud service providers and consumers is critical because the cloud service consumer is still ultimately responsible for compliance and protection of their critical data, even if that workload had moved to the cloud. In fact, some organizations choose private or hybrid models over public clouds because of the risks associated with outsourcing services.

Other aspects about cloud computing also require a major reassessment of security and risk. Inside the cloud, it is difficult to physically locate where data is stored. Security processes that were once visible are now hidden behind layers of abstraction. This lack of visibility can create a number of security and compliance issues.

In addition, the massive sharing of infrastructure with cloud computing creates a significant difference between cloud security and security in more traditional IT environments. Users spanning different corporations and trust levels often interact with the same set of computing resources. At the same time, workload balancing, changing service level agreements, and other aspects of today's dynamic IT environments create even more opportunities for misconfiguration, data compromise, and malicious conduct.

Infrastructure sharing calls for a high degree of standardized and process automation, which can help improve security by eliminating the risk of operator error and oversight. However, the risks inherent with a massively shared infrastructure mean that cloud computing models must still place a strong emphasis on isolation, identity, and compliance.

Cloud computing is available in several service models (and hybrids of these models). Each presents different levels of responsibility for security management.

- **Software-as-a-Service**. Software-as-a-Service (SaaS) is a model of software deployment where by one or more applications and the computational resources to run them are provided for use on demand as a turnkey service. Its main purpose is to reduce the total cost of hardware and software development, maintenance, and operations. Security provisions are carried out mainly by the cloud provider. The cloud subscriber does not manage or control the underlying cloud

infrastructure or individual applications, except for preference selections and limited administrative application settings.

- **Platform-as-a-Service**. Platform-as-a-Service (PaaS) is a model of software deployment whereby the computing platform is provided as an on-demand service upon which applications can be developed and deployed. Its main purpose is to reduce the cost and complexity of buying, housing, and managing the underlying hardware and software components of the platform, including any needed program and database development tools. The development environment is typically special purpose, determined by the cloud provider and tailored to the design and architecture of its platform. The cloud subscriber has control over applications and application environment settings of the platform. Security provisions are split between the cloud provider and the cloud subscriber.
- **Infrastructure-as-a-Service**. Infrastructure-as-a-Service (IaaS) is a model of software deployment whereby the basic computing infrastructure of servers, software, and network equipment is provided as an on-demand service upon which a platform to develop and execute applications can be established. Its main purpose is to avoid purchasing, housing, and managing the basic hardware and software infrastructure components, and instead obtain those resources as virtualized objects controllable via a service interface. The cloud subscriber generally has broad freedom to choose the operating system and development environment to be hosted. Security provisions beyond the basic infrastructure are carried out mainly by the cloud subscriber.

Figure 1 illustrates the differences in scope and control between the cloud subscriber and cloud provider, for each of the service models discussed above. Five conceptual layers of a generalized cloud environment are identified in the center diagram and apply to public clouds, as well as each of the other deployment models. The arrows at the left and right of the diagram denote the approximate range of the cloud provider's and user's scope and control over the cloud environment for each service model. In general, the higher the level of support available from a cloud provider, the more narrow the scope and control the cloud subscriber has over the system.

The two lowest layers shown denote the physical elements of a cloud environment, which are under the full control of the cloud provider, regardless of the service model. Heating, ventilation, air conditioning (HVAC), power, communications, and other aspects of the physical plant comprise the lowest layer, the facility layer, while computers, network and storage components, and other physical computing infrastructure elements comprise the hardware layer.

The remaining layers denote the logical elements of a cloud environment. The virtualized infrastructure layer entails software elements, such as hypervisors, virtual machines, virtual data storage, and supporting middleware components used to realize the infrastructure upon which a computing platform can be established. While virtual machine technology is commonly used at this layer, other means of providing the necessary software abstractions are not precluded. Similarly, the platform architecture layer entails compilers, libraries, utilities, and other software tools and development environments needed to implement applications. The application layer represents deployed software applications targeted towards end-user software clients or other programs, and made available via the cloud.

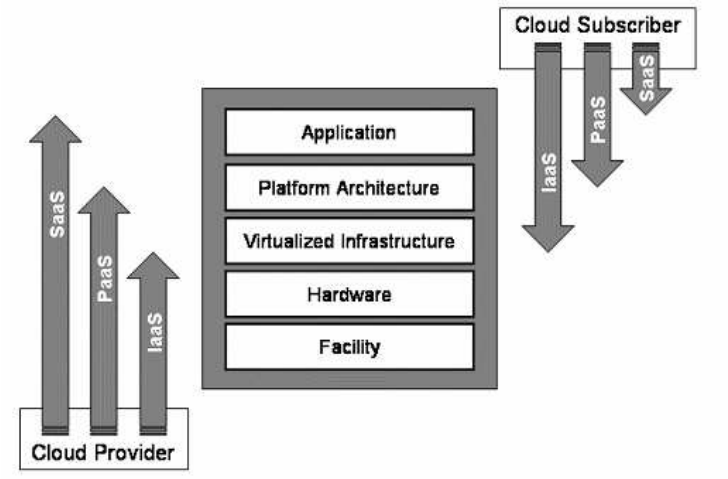

Figure 1. Differences in Scope and Control among Cloud Service Models

III. Cloud computing issues

What are the "security" concerns that are preventing companies from taking advantage of the cloud? Numerous studies, for example IDC's 2008 Cloud Services User Survey [2] of IT executives, cite security as the number one challenge for cloud users.

The Cloud Security Alliance's initial report [3] contains a different sort of taxonomy based on 15 different security domains and the processes that need to be followed in an overall cloud deployment. We categorize the security concerns as:

- Traditional security
- Availability
- Third-party data control

Traditional Security

These concerns involve computer and network intrusions or attacks that will be made possible or at least easier by moving to the cloud. Cloud providers respond to these concerns by arguing that their security measures and processes are more mature and tested than those of the average company.

Concerns in this category include:

VM-level attacks. Potential vulnerabilities in the hypervisor or VM technology used by cloud vendors are a potential problem in multi-tenant architectures. Vulnerabilities have appeared in VMWare, Xen, and Microsoft's Virtual PC and Virtual Server. Vendors such as Third Brigade mitigate potential VM-level vulnerabilities through monitoring and firewalls.

Cloud provider vulnerabilities. These could be platform-level, such as an SQL-injection or cross-site scripting vulnerability in salesforce.com. For instance, there have been a couple of recent Google Docs vulnerabilities [4] and [5]. There is nothing new in the nature of these vulnerabilities; only their setting is novel. In fact, IBM has repositioned its Rational AppScan tool, which scans for vulnerabilities in web services as a cloud security.

Phishing cloud provider. Phishers and other social engineers have a new attack vector, as the Salesforce phishing incident [6] shows.

Expanded network attack surface. The cloud user must protect the infrastructure used to connect and interact with the cloud, a task complicated by the cloud being outside the firewall in many cases. For instance, [7] shows an example of how the cloud might attack the machine connecting to it.

Authentication and Authorization. The enterprise authentication and authorization framework does not naturally extend into the cloud. How does a company meld its existing framework to include cloud resources? Furthermore, how does an enterprise merge cloud security data (if even available) with its own security metrics and policies?

Availability

These concerns center on critical applications and data being available. Well-publicized incidents of cloud outages include Gmail (one-day outage in mid-October 2008 [8]), Amazon S3 (over seven-hour downtime on July 20, 2008 [9]), and FlexiScale (18-hour outage on October 31, 2008 [10]).

1. Uptime. As with the Traditional Security concerns, cloud providers argue that their server uptime compares well with the availability of the cloud user's own data centers.Besides just services and applications being down, this includes the concern that a third-party cloud would not scale well enough to handle certain applications. SAP's CEO, Leo Apotheker said: "There are certain things that you cannot run in the cloud because the cloud would collapse...Don't believe that any utility company is going to run its billing for 50 million consumers in the cloud." (11/24/08, searchSAP.com)
2. Single point of failure. Cloud services are thought of as providing more availability, but perhaps not - there are more single points of failure and attack.
3. Assurance of computational integrity. Can an enterprise be assured that a cloud provider is faithfully running a hosted application and giving valid results? For example, Stanford's Folding@Home project gives the same task to multiple clients to reach a consensus on the correct result.

Third-party data control

The legal implications of data and applications being held by a third party are complex and not well understood. There is also a potential lack of control and transparency when a third party holds the data. Part of the hype of cloud computing is that the cloud can be implementation independent, but in reality regulatory compliance requires transparency into the cloud.

All this is prompting some companies to build private clouds to avoid these issues and yet retain some of the advantages of cloud computing. For example, Benjamin Linder, Scalent System's CEO, says [11]: "What I find as CEO of a software company in this space, Scalent Systems, is that most enterprises have a hard time trusting external clouds for their proprietary and high-availability systems. They are instead building internal "clouds", or "utilities" to serve their internal customers in a more controlled way."

1 Due diligence. If served a subpoena or other legal action, can a cloud user compel the cloud provider to respond in the required time-frame? A related question is the provability of deletion, relevant to an enterprise's retention policy: How can a cloud user be guaranteed that data has been deleted by the cloud provider?

2 Auditability. Audit difficulty is another side effect of the lack of control in the cloud. Is there

sufficient transparency in the operations of the cloud provider for auditing purposes? Currently, this transparency is provided by documentation and manual audits. Information Security A related concern is proper governance of cloud-related activity. It's easy, perhaps too easy, to start using a cloud service.One popular auditing guideline is the SAS 70, which defines guidelines for auditors to assess internal controls, for instance controls over the processing of sensitive information. SOX and HIPAA are other well-known regulations. US government agencies generally need to follow guidelines from FISMA, NIST, and FIPS.Certain regulations require data and operations to remain in certain geographic locations. Cloud providers are beginning to respond with geo-targeted offerings [12].

3 Contractual obligations. One problem with using another company's infrastructure besides the uncertain alignment of interests is that there might be surprising legal implications. For instance, here is a passage from Amazon's terms of use [13]: Non-Assertion. During and after the term of the Agreement, with respect to any of the Services that you elect to use, you will not assert, nor will you authorize, assist, or encourage any third party to assert, against us or any of our customers, end users, vendors, business partners (including third party sellers on websites operated by or on behalf of us), licensors, sublicensees or transferees, any patent infringement or other intellectual property infringement claim with respect to such Services.This could be interpreted as implying that after you use EC2, you cannot file infringement claims against Amazon or its customers suggesting that EC2 itself violates any of your patents. It's not clear whether this non-assert would be upheld by the courts, but any uncertainty is bad for business.

4 Cloud Provider Espionage. This is the worry of theft of company proprietary information by the cloud provider. For example, Google Gmail and Google Apps are examples of services supported by a private cloud infrastructure. Corporate users of these services are concerned about confidentiality and availability of their data.

5 Data Lock-in. How does a cloud user avoid lock-in to a particular cloud-computing vendor? The data might itself be locked in a proprietary format, and there are also issues with training and processes. There is also the problem of the cloud user having no control over frequent changes in cloud-based services.

6 Transitive nature. Another possible concern is that the contracted cloud provider might itself use subcontractors, over whom the cloud user has even less control, and who also must be trusted. One example is the online storage service called The Linkup, which in turn used an online storage company called Nirvanix. The Linkup shutdown after losing sizeable amounts of customer data, which some say was the fault of Nirvanix [14]. Another example is Carbonite [15], who is suing its hardware providers for faulty equipment causing loss of customer data.

IV. NEW PROBLEMS

In this section we outline new problem areas in security that arise from cloud computing. These problems may only become apparent after the maturation and more widespread adoption of cloud computing as a technology.

Cheap data and data analysis. The rise of cloud computing has created enormous data sets that can be monetized by applications such as advertising. Google, for instance, leverages its cloud infrastructure to collect and analyze consumer data for its advertising network. Collection and analysis of data is nowpossible cheaply, even for companies lacking Google's resources. What is the impact on privacy of abundant data and cheap data-mining? Because of the cloud, attackers potentially have massive, centralized databases available for analysis and also the raw computing power to mine these databases. For example, Google is essentially doing cheap data mining when it returns search results. How much more privacy did one have before one could be Googled?

Because of privacy concerns, enterprises running clouds collecting data have felt increasing pressure to anonymize their data. EPIC has called for Gmail, Google Docs, Google Calendar, and the company's other Web applications to be shut down until appropriate privacy guards are in place [16]. Google and Yahoo!, because of pressure from privacy advocates, now have an 18 month retention policy for their search data, after which it will be anonymized. This means that some identifying data will be removed such as IP addresses and cookie information. The anonymized data is retained though, to support the continual testing of their algorithms. Another reason to anonymize data is to share data with other parties. We note that anonymizing data is a difficult problem. For example, in [17] the Netflix data set was partially de-anonymized, and in [18] the then-Governor of Massachusetts was identified as a patient of Massachusetts General Hospital from an anonymized list of discharged patients. Tools are needed for effective anonymization, which will increase in importance as clouds proliferate and more data is collected that needs to be analyzed safely or shared.

Cost-effective defense of availability. Availability also needs to be considered in the context of an adversary whose goals are simply to sabotage activities. Increasingly, such adversaries are becoming realistic as political conflict is taken onto the web, and as the recent cyber attacks on Lithuania confirm [21]. The damages are not only related to the losses of productivity, but extend to losses due to the degraded trust in the infrastructure, and potentially costly backup measures. The cloud computing model encourages single points of failure. It is therefore important to develop methods for sustained availability (in the context of attack), and for recovery from attack. The latter could operate on the basis of minimization of losses, required service levels, or

similar measures.

Increased authentication demands. The development of cloud computing may, in the extreme, allow the use of thin clients on the client side. Rather than a license purchased and software installation on the client side, users will authenticate in order to be able to use a cloud application. There are some advantages in such a model, such as making software piracy more difficult and giving the ability to centralize monitoring. It also may help prevent the spread of sensitive data on untrustworthy clients.

Thin clients result in a number of opportunities related to security, including the paradigm in which typical users do not have to worry about the risks of any actions - their security is managed by the cloud, which maintains the software they run. This architecture stimulates mobility of users, but increases the need to address authentication in a secure manner. In addition, the movement towards increased hosting of data and applications in the cloud and lesser reliance on specific user machines is likely to increase the threat of phishing and other abusive technologies aimed at stealing access credentials, or otherwise derive them, e.g., by brute force methods.

Mash-up authorization. As adoption of cloud computing grows, we are likely to see more and more services performing mash-ups of data. This development has potential security implications, both in terms of data leaks, and in terms of the number of sources of data a user may have to pull data from - this, in turn, places requirements on how access is authorized for reasons of usability. While centralized access control may solve many of these problems, that may not be possible - or even desirable.

One example in this area is provided by Facebook. Facebook users upload both sensitive and non-sensitive data. This data is both utilized by Facebook to present the data to other users, and also utilized by third party applications that are run by the platform. These applications are typically not verified by Facebook. Hence, there is a drive to create malicious applications that run in Facebook's cloud to steal sensitive data, e.g., see [20].

V. Conclusion

Cloud computing is the most popular notion in IT today; even an academic report [21] from UC Berkeley says "Cloud Computing is likely to have the same impact on software that foundries have had on the hardware industry." They go on to recommend that "developers would be wise to design their next generation of systems to be deployed into Cloud Computing". While many of the predictions may be cloud hype, we believe the new IT procurement model offered by cloud computing is here to stay. Whether adoption becomes as prevalent and deep as some forecast will depend largely on overcoming fears of the cloud.

Cloud fears largely stem from the perceived loss of control of sensitive data. Current control measures do not adequately address cloud computing's third-party data storage and processing needs. In our vision, we propose to extend control measures from the enterprise into the cloud through the use of Trusted Computing and applied cryptographic techniques. These measures should alleviate much of today's fear of cloud computing, and, we believe, have the potential to provide demonstrable business intelligence advantages to cloud participation.

Our vision also relates to likely problems and abuses arising from a greater reliance on cloud computing, and how to maintain security in the face of such attacks. Namely, the new threats require new constructions to maintain and improve security. Among these are tools to control and understand privacy leaks, perform authentication, and guarantee availability in the face of cloud denial-of-service attacks.

REFERENCES

[1] Peter Mell, Tim Grance, The NIST Definition of Cloud Computing, Version 15, October, 7. 2009. URL: http://csrc.nist.gov/groups/SNS/cloud-computing

[2] IT Cloud Services User Survey, pt.2: Top Benefits & Challenges, URL: http: //blogs.idc.com/ie/?p=210.

[3] Security Guidance for Critical Areas of Focus in Cloud Computing. URL: http://www.cloudsecurityalliance.org/guidance/csaguide.pdf.

[4] Google Docs Glitch Exposes Private Files. http://www.pcworld.com/article/160927/google_docs_glitch_exposes_private_files.html.

[5] Security issues with Google Docs. URL: http://peekay.org/2009/03/26/security-issues-with-google-docs/ .

[6] Salesforce.com Warns Customers of Phishing Scam. URL: http://www.pcworld.com/businesscenter/article/139353/salesforcecom_warns_customers_of_phishing_scam.html .

[7] Security Evaluation of Grid Environments. https://hpcrd.lbl.gov/HEPCybersecurity/HEP-Sec-Miller-Mar2005.ppt.

[8] Extended Gmail outage hits Apps admins. http://www.computerworld.com/action/article.do?command =viewArticleBasic&articleId=9117322.

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

[10] FlexiScale Suffers 18-Hour Outage. URL: http://www.thewhir.com/web-hosting-news/103108_FlexiScale_Suffers_18_Hour_Outage.

[11] Disaster-Proofing The Cloud. URL: http://www.forbes.com/2008/11/24/cio-cloud-disaster-tech-cio-cx_dw_1125cloud.html.

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

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

[14] Loss of customer data spurs closure of online storage service 'The Linkup'. URL: http://www.networkworld.com/news/2008/081108-linkup-failure.html?page=1.

[15] Latest cloud storage hiccups prompts data security questions. URL: http://www.computerworld.com/action/article.do?command=viewArticleBasic&articleId=9130682&source=NLT_PM.

[16] FTC questions cloud-computing security. URL:

http://news.cnet.com/8301-13578_3-10198577-38.html?part=rss&subj=news&tag=2547-1_3-0-20.
[17] Narayanan, A. and Shmatikov, V. Robust De-anonymization of Large Sparse Datasets. In IEEE Symposium on Security and Privacy. IEEE Computer Society, 2008.
[18] Sweeney, L. Weaving technology and policy together. J. of Law, Medicine and Ethics, 25, 2-3 (1997).
[19] Lithuania Weathers Cyber Attack, Braces for Round 2. URL: http://blog.washingtonpost.com/securityfix/2008/07/lithuani a_weathers_cyber_attac_1.html.
[20] Facebook users suffer viral surge. URL: http://news.bbc.co.uk/2/hi/technology/7918839.stm.
[21] 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.

Vladimir Guzhov is dean of the faculty of Automatics and Computer Engineering in Novosibirsk State Technical University, professor, doctor of technical sciences. He is the author of 120 science papers including 4 patents. The science interests and competence field is program systems, high accuracy measurements.
E-mail: vig@edu.nstu.ru

Sergey Ilinykh is assistant-professor inNovosibirsk State Technical University, professor, PhD. He is the author of 50 science papers including 4 patents and 1 high school-book. The science interests and competence field is, program systems, laser systems.
E-mail: i sp51@yandex.ru
Alexander Vagizov is undergraduate student in Novosibirsk State Technical University. The science interests and competence field is optical measurement systems, high accuracy measurements. E-mail: famas-yandex.ru

Konstantin Bazhenov is a post graduate student in the Novosibirsk State Technical University. He graduated from the NSTU in 2009 from the Faculty of Automation and Computer Engineering with the diploma in engineering. He is the author of 3 publications.
The science interests and competence field is high accuracy measurements, program systems and cloud computing.
E-mail: kos_4ll@mail.ru