Is IaaS Platform Secure?

Nitin Kamble
INURTURE, Bangalore

Prince Pradhan
ADYPU, Pune

Mukesh Ghanchi
ADYPU, Pune

Ashutosh jaiswal
ADYPU, Pune

ABSTRACT
This paper put insight on how the security mechanisms enforced by cloud computing providers of their IaaS platforms and enabling their users to utilize the platforms securely. This paper provides the assessment of two popular IaaS platforms with respect to the following criteria: network security, authentication and API security, security attack protection and high availability, logging and monitoring, access right control, and security. The assessment is based on the information provided by cloud computing providers and usage of the IaaS platforms.

Keywords
Cloud computing, Infrastructure as a Service, security, Amazon Web Services, Google Cloud Platform.

1. INTRODUCTION
Why IaaS platforms? Cloud computing involves many other services like Platform as a Service (PaaS) or Software as a Service (SaaS) with their very own kinds of security issues. However, since these services are often more focused on certain tasks, they are not involved in this assessment.

In cloud computing security is major concern. It implies all the security threats that are affecting traditional server environments. However, due to the fact that customer data in IaaS platforms are stored and processed on shared hardware, cloud computing entails new security threats that do not exist in the traditional server environments. Let us take the management interfaces of IaaS platforms pose an additional access mechanism that can be a target for security attacks.

2. RELATED WORK
Cloud computing security is major concern and debated extensively. To mention few, like *Takabi et al.* [1] we address similar aspects of security considerations in this paper, but they featured only abstract concepts and did not assess any concrete IaaS platforms. Another researchers like *Bouayad et al.* [2] and *Khalil et al.* [3] concentrated on security challenges like instance isolation, resource visualization and identity management, but again without assessing any concrete IaaS platforms.

This paper try to fill a gap in this field, for that decided to assess two popular IaaS platforms like Amazon Web Services (AWS) Google Cloud Platform (GCP). This assessment is based on both the information provided by cloud computing providers and working experience with the IaaS platforms.

3. ASSESMENT
Amazon Web Services and Google Cloud Platform are the front-runners of scalable commercial platforms around the sphere. The reason behind their selection is their providers belong to the market leaders. Amazon provides IaaS like Amazon EC2 since 2006 while its challenger, Google Cloud Platform with Google Compute Engine, entered the market in 2012. Cloud computing is not a new market for Google because it has offered other kinds of cloud services like Google App Engine in terms of PaaS before. However, in the fast moving IT world, a six-year gap of experience in providing IaaS might be a vital factor.

3.1 Network Security
Network security is one of the key topics in any client-server environments. More or less every kind of service requires network communication with third parties. These third parties are depends on the service itself.

All the assessed platforms provide some means to regulate traffic by defining firewall rules. Amazon Web Services enable to regulate both inbound and outbound traffic [4] [5] [6], while Google Compute Engine enables to regulate only inbound traffic and recommends to use additional means like IP tables if outbound connections should be restricted [7]. All the platforms allow for the creation of virtual internal networks between groups of instances and VPN connections into these internal networks. In addition, Amazon Web Services isolate every single Amazon EC2 instance with their own firewall rules, which are called security groups. Google Compute Engine uses firewall rules for external access to an instance network and internal network communication [8].

Table 1. Network Security

	AWS	GCP
Inbound Rules	√	√
Outbound Rules	√	
VPN Gateway	√	√

Table 1 summarizes the assessment of Network Security with respect to cloud. The network security capabilities of other platforms are very similar, except for Google Compute Engine that does not enable to regulate outbound traffic. However, since this feature should not excessively degrade network security, no platform is considered to be better than the others in our evaluation.

3.2 API Security and Authentication
IaaS platforms have to ensure confidentially, integrity and availability of data on shared virtualized servers, and guarantee that user-defined restrictions such as network and management access are appropriately applied. Providers must implement suitable means to configure network access with firewalls or similar mechanisms, and enable smooth access control to cloud management interfaces such as web interface or an Application Programming Interface (API). Access to these interfaces has to be secured efficiently.

Multi-factor authentication (MFA) massively reduces the risk of security attacks based on a stolen password since access is still restricted without the additional authentication factors. As any service with elevated security requirements, IaaS platforms should provide multi-factor authentication or other methods that are considered more secure. Software-based implementations like time-based One-time Password (OTP) should provide a similar level of security as hardware tokens and out-of-band authentication if used correctly. The same consideration applies to certificate-based authentication methods. Amazon Web

Services support time based OTP and hardware OTP tokens by Gemalto as part of its web service called Identity and Access Management (IAM) [9]. Google Compute Engine supports the same additional authentication factors as all Google services, including time based OTP and out-of-band authentication via OAuth 2.0 [10].

Amazon Web Services verify the client and the client request with a signature over the request data, and prevent replay attacks with a timestamp in the request data [11]. OAuth 2.0 is used by Google Compute Engine for Application Programming Interface. OAuth2.0 depend on Transport Layer Security for server authentication as well as integrity of the data [12]. A man-in-the-middle, who is able to break into the TLS connection, could alter the API requests on Google Compute Engine has the same security issues when using its API, but provides a more secure certificate-based authentication with the CLI.

HTTP public key pinning (HPKP) is supported by Google Compute Engine for its API. Customers supporting HTTP public key pinning will not be vulnerable to man-in-the middle attacks with a valid but unauthorized certificate if the first connection to the server has been valid. The API of Amazon Web Services is not vulnerable to man-in-the-middle attacks on the TLS connection due to the additional request signature. Neither the web interface of Amazon Web Services nor their API are currently using HTTP public key pinning as an additional security feature.

Table 2. Authentication & API Security

	AWS	GCP
MFA	H, S	S, O
HPKP		√
Certificate Authentication	√	
Signed API Request	√	

H – Hardware OTP | S – Software OTP | O – Out of Band

In Table 2 summarization of assessment is illustrated. Due to the signed API requests and certificate-based authentication methods, Amazon Web Services provide the most secured API. Google Compute Engine completely rely on TLS for their APIs. This makes them the least secure in assessment.

3.3 High Availability

Distributed Denial of Service (DDoS) attacks can very easily target Instances. Due to their architecture as a distributed system, cloud computing environments tend to provide sophisticated means to weaken or even defend DDoS attacks. These might include load balancers, automated deployment of firewall rules and high availability instances distributed across availability zones. Also automatic notifications on high load situations can help administrators to quickly handle a DDoS attack. Cloud providers have enormous amounts of bandwidth and processing power shared across different data centers, thereby increasing the required resources for a successful DDoS attack. High availability not only is relevant to DDoS attack prevention but also part of the basic information security attributes likes confidentially, integrity and availability of data.

Amazon Web Services and Google Compute Platform infrastructure is a high availability environment with a guaranteed uptime covered by their Service Level Agreements (SLAs). Precisely, Amazon Web Services and Google Compute Engine assure a monthly uptime of at least 99.95% for their IaaS platforms [13] [14]. Google Compute Engine does a live migration of instances when hardware has to be maintained if not configured otherwise by users [15] in contrast, Amazon Web Services do not support live migration of instances so that during hardware maintenance, instances are shut down and restarted afterwards [16]. Users have to configure the instances to perform the desired actions automatically after a reboot to avoid the unwanted downtime of a high availability service in those cases. Amazon EC2 enables to check the health of instances and automatically reboot them when the health check fails [17]. In combination with a high availability service hosted in the instances, this feature can increase the availability of the whole system by keeping the individual instances running. A similar feature was not found in any other evaluated platform.

Cloud resources such as instances are often vended on a pay-per-use basis. This might add to unexpected costs during DDoS attacks, specifically in combination with automatic scaling of instances. Cloud providers should provide functionalities to set automatic scaling and cost limits. This will help users to get notified on other configurable triggers.

Auto-scaling rules for Amazon Web Services and Google Compute Engine empower to specify automatic scaling and cost limits while the former also features notifications on configurable triggers and billing alerts when the costs exceed specified limits [18]. This allows users to efficiently prevent cost exploitation attacks. In count of instances, many other resources like networks, storages and databases can cause unexpected and unwanted costs. These costs are not directly related to auto-scaling rules and can depend on many other factors. The billing alarms provided by Amazon Web Services are a very useful shield, unfortunately which is not found in Google Compute Engine.

Table 3. High Availability

	AWS	GCP
Notifications	√	
HA Platform	√	√
Automatic Restart	√	
Live Migration		√

In Table 3 summarization of assessment based on high availability is clarified. Amazon Web Services are supporting notifications on scaling actions or other configurable metrics, which allow users to quickly react to unusual or unwanted events in their cloud computing environs [19] [20]. It also supports, cost-based alerts and a watch dog service for instance health should enable users to easily deploy a high availability environment to Amazon EC2. The missing live migration support should be a minor issue when the environment is distributed across different availability zones. The other platforms also enable to deploy high available environments but none of them matches the features provided by Amazon Web Services. Therefore, Amazon Web Services are the best fit under this title.

3.4 Logging and Monitoring

Not only does security in cloud computing include the prevention of DDoS attacks, but also the analysis of security breaches caused by those attacks. Therefore, IaaS platforms should provide (at least) log files about access, executed commands and networks traffic to enable the investigation on the infrastructure level.

Currently Google Compute Engine logging is in beta phase which is not covered in any Service Level Agreement. Nevertheless, Google Compute Engine can log operations performed on instances, record metrics and store custom log entries. Furthermore, the system log files are stored directly on Google Cloud Platform [21], thus preventing a security breach in the instances from being covered up by modified log files. All log files and recorded metrics can be viewed using Google Cloud Logging [22].

Table 4. Logging and Monitoring

	AWS	GCP
Action Log	√	√
Instance Logs	√	√
Metric	√	√
User viewable action logs	√	√
User viewable request logs	√	√
User viewable metric	√	√

Amazon Web Services offer more versatile logging mechanisms that enable the logging of every API request, recording of metrics from all resources and storing of custom logs from instances. All log files and recorded metrics can be viewed from the web interface or accessed via an API.

In Table 4 summarization of assessment based on logging and monitoring is illustrated. The open source platforms provide much fewer logging and monitoring features than their commercial counterparts do. Amazon Web Services and Google Cloud Engine have the same abilities for a security beach analysis. Therefore, none of them can be considered better than the other one.

3.5 Access Control

For security reasons, IaaS platforms should provide fine-grained access control to cloud resources.

Amazon Web Services enabled to utilize the least required privilege principle for every user who wants to access or control parts of the IaaS platforms [24]. Amazon Web Services even enable to add more sophisticated conditions like the source IP network of a request or the time of the day to grant access on all or certain resources [25]. Google Compute Engine has a much simpler access management with only three different levels to a whole project: “can view”, “can edit” and “is owner” [26].

In cloud computing environments, access rights can easily be made less restrictive than intended. Some rules might accidentally override others and thereby grant unintended rights to unauthorized users or user groups. To avoid this issue, IaaS platforms should provide methods to get an overview of the over-granted access rights to cloud resources like instances.

Amazon Web Services provide the IAM policy simulator, which is an outstanding means to verify the access rights to cloud resources. As a result, all possible API operations on the IaaS platform can be simulated for every user and user group, even with the simulation of additional conditions used in an advanced rule set. No other evaluated platform has a remotely similar system to review the functionality of specified rules. Precisely, Google Compute Engine has a simple overview page where all users and user groups of a project are listed along with their access levels. Since only three different access levels can be granted on for a whole project, this seems feasible for a reviewing purpose.

Table 5. Access Control

	AWS	GCP
Command access rule	√	
Resource access rule	√	
Additional conditions	√	
User overview	√	√
Rule overview	√	

In Table 5 summarization of assessment based on access control is illustrated. All the evaluated platforms support a rule-based scaling of instances though this feature is currently labeled as beta in Google Compute Engine [27]. Since Google Cloud Platform does not have a command based rule system, a simple overview page fits the security needs. Among the IaaS platforms with a command-based rule system, Amazon Web Services provide the best control and review mechanisms.

Google Compute Engine has very limited access regulation mechanisms. This can lead to security issues if the project has different members. Limiting resource access to only required commands would be impossible then. Therefore, we consider Google Compute Engine to be the worst in our assessment.

3.6 Physical Security

Finally, IaaS platforms have to ensure security on physical and environmental level. This includes access control to data centers and server racks along with threat management like power failure, fire and natural disasters.

Amazon Web Services and Google Compute Engine make use different data centers which are globally sited. These data centers consist of different availability zones, which can operate independently. Availability zones having their own (emergency) power supplies are located in different fire protection zones and utilize their own connection to the Internet. Customers can use this redundancy in the data centers to spread their data across availability zones and thus, protect the data from a total system failure in one of the zones.

Table 6. Physical Security

	AWS	GCP
Redundant Locations	√	√
Availability zones	√	√
Threat management	√	√
Physical server access	√	√

In Table 5 summarization of assessment based on physical security is illustrated. As can be seen, the commercial IaaS platforms are more advanced than their commercial counterparts. In precise, physically restricted access to a server and natural disaster protection are not the responsibility of open source IaaS platforms and to achieve these features we have to take help of external means.

4. FUTURE WORK

Cloud computing providers and customers of cloud services are sharing security responsibilities. The type of an IaaS platform defines who is responsible for which part of the system. For example, in case of commercial platforms like Amazon Web

Services and Google Cloud Platform, the providers are responsible for the physical infrastructure, networks and isolations via hypervisors [28] [29]. Customers have to take care of the utilized operating system and firewall configuration, and access management in the IaaS platforms.

Cloud computing environments are usually running multiple virtualized instances on shared host servers. It is crucial that these instances cannot affect each other or even access resources from other instances, especially since they might belong to different tenants. Typically memory, disk and network access are isolated from other instances by a hypervisor. Like all other software systems, hypervisors might include unknown or undisclosed vulnerabilities that break the instance isolation. For example, the Venom vulnerability in the QEMU virtual Floppy Disk Controller has recently threatened the commonly used hypervisors QEMU, KVM and Xen [30]. Moreover, highly sophisticated attacks on hardware properties like the row hammer attack are possible security threats on shared systems [23]. These basic risks affecting both commercial and open source platforms are not evaluated in this paper, but they have to be considered in the future. Mandatory

Access Control (MAC) systems like SELinux and AppArmor can provide an additional security layer and minimize the impact of vulnerabilities in the hypervisor, but they are also just another software security layer.

Instance isolation is the keystone in cloud security. Users of both public and private clouds have to believe that this lowest security level is working correctly. Commercial providers claim to continuously and thoroughly scan their infrastructure for vulnerabilities. Host systems on both evaluated commercial platforms are equipped with Trusted Platform Modules or similar systems that are validating the integrity of the host software [4] [7]. This level of defense is not part of any evaluated open source platform. Rather, it has to be added manually with intrusion detection systems and vulnerability scanners to reach a similar security level.

Commercial platforms may have data centers all across the world and thus, it may be not always transparent to the customers where their data are stored and processed. From another point of view, privacy laws can require that customer data are stored and processed in a specific region or country. Therefore, in the future, we are also going to evaluate if the cloud computing providers declare where customer data are stored and processed and if the customers can choose the location.

5. CONCLUSION

Amazon Web Services and Google Cloud are assessed in this paper with reference to IaaS Platform. The assessment has showed that Amazon Web Services offer the most useful and advanced tools for customers by supporting a secure configuration of their IaaS platform. Regulating access to every possible API command even with additionally required conditions and the possibilities to inspect these access rules by simulator of the IAM policies are incomparable by any other assessed platform. The provision of setting scaling rules, monitoring nearly every aspect of the cloud computing environment and creating notifications on a lot of different triggers surpass the competencies of competitors. In addition, Amazon Web Services are the only evaluated platform that does not completely rely on the security of TLS for their API and as an alternative uses additional signature for the requests.

6. REFERENCES

[1] H. Takabi, J. Joshi, and G.-J. Ahn, "Security and privacy challenges in cloud computing environments," Security Privacy, IEEE, vol. 8, no. 6, pp. 24–31, Nov 2010.

[2] A. Bouayad, A. Blilat, N. El Houda Mejhed, and M. El Ghazi, "Cloud computing: Security challenges," in Information Science and Technology (CIST), 2012 Colloquium in, Oct 2012, pp. 26– 31.

[3] I. Khalil, A. Khreishah, S. Bouktif, and A. Ahmad, "Security concerns in cloud computing," in Information Technology: New Generations (ITNG), 2013 Tenth International Conference on, April 2013, pp. 411–416.

[4] "Cloud Security - AWS," http://aws.amazon.com/security, 2015, [Online; accessed 06.11.2015].

[5] "Integrity life-cycle - OpenStack Security Guide," http://docs.openstack.org/security-guide/content/integrity-life-cycle.html, [Online; accessed 06.11.2015].

[6] "Authentication Overview - OpenNebula 4.12.1 documentation," http://docs.opennebula.org/4.12/administration/authentication/external_auth.html, 2015, [Online; accessed 06.11.2015].

[7] "Google Security Whitepaper – Google Cloud Platform," https://cloud.google.com/security/whitepaper#custom_server_hardware_and_software, 2015, [Online; accessed 06.11.2015].

[8] "Networking and Firewalls - Compute Engine – Google Cloud Platform", https://cloud.google.com/compute/docs/networking,2015, [Online; accessed 06.11.2015].

[9] [9] "AWS - Multi-Factor Authentication," https://aws.amazon. com/iam/details/mfa/, [Online; accessed 06.11.2015].

[10] "Authentication - Compute Engine – Google Cloud Platform," https://cloud.google.com/storage/docs/authentication, 2015, [Online; accessed 06.11.2015].

[11] "Signing AWS API Requests - Amazon Web Services," http://docs.aws.amazon.com/general/latest/gr/signing_aws _ api_requests.html, [Online; accessed 06.11.2015].

[12] "RFC 6749 - The OAuth 2.0 Authorization Framework."

[13] "Google Compute Engine Service Level Agreement (SLA)," https://cloud.google.com/compute/sla, 2015, [Online; accessed 06.11.2015].

[14] "Amazon EC2 SLA," http://aws.amazon.com/ec2/sla/, 2015, [Online; accessed 06.11.2015].

[15] "Instances - Compute Engine – Google Cloud Platform," https://cloud.google.com/compute/docs/instances/#onhostmaintenance, 2015, [Online; accessed 06.11.2015].

[16] "Amazon EC2 Maintenance Help Page," http://aws.amazon.com/maintenance-help/, 2015, [Online; accessed 6.11.2015].

[17] "New – Auto Recovery for Amazon EC2 | AWS Official Blog," https://aws.amazon.com/blogs/aws/ new-auto-recovery-for-amazon-ec2/,2015, [Online; accessed 06.11.2015].

[18] "Monitor Your AWS Charges with Billing Alerts Using Amazon CloudWatch," http://aws.amazon.com/about-aws/whatsnew/2012/05/10/announcing-aws-billing-alerts/, 2015, [Online; accessed 06.11.2015].

[19] "Getting Notifications When Your Auto Scaling Group Changes – Auto Scaling," http://docs.aws.amazon.com/AutoScaling/latest/DeveloperGuide/ASGettingNotifications.html, 2015, [Online; accessed 06.11.2015].

[20] "Ceilometer/Alerting - OpenStack," https://wiki.openstack.org/wiki/Ceilometer/Alerting, [Online; accessed 06.11.2015].

[21] "List of Log Types - Cloud Logging – Google Cloud Platform," https://cloud.google.com/logging/docs/view/logs_ index, 2015, [Online; accessed 06.11.2015].

[22] "Pricing - Cloud Logging – Google Cloud Platform," https://cloud.google.com/logging/pricing, 2015, [Online; accessed 06.11.2015].

[23] M. Seaborn and T. Dullien, "Exploiting the DRAM rowhammer bug to gain kernel privileges," http://googleprojectzero.blogspot.de/2015/03/exploiting-dramrowhammer-bug-to-gain.html, 2015, [Online; accessed 06.11.2015].

[24] "Chapter 9. Managing Projects and Users - OpenStack Operations Guide," http://docs.openstack.org/openstack-ops/content/projects_users.html, 2015, [Online; accessed 06.11.2015].

[25] "AWS - Manage Permissions and Policies," http://aws.amazon.com/de/iam/details/manage-permissions/, 2015, [Online; accessed 06.11.2015].

[26] "Configuring permissions on Google Cloud Platform," https://cloud.google.com/docs/permissions-overview, 2015, [Online; accessed 06.11.2015].

[27] "Autoscaler - Compute Engine – Google Cloud Platform," https://cloud.google.com/compute/docs/autoscaler/, 2015, [Online; accessed 06.11.2015].

[28] "AWS Shared Responsibility Model," http://aws.amazon.com/compliance/shared-responsibility-model/, 2015, [Online; accessed 06.11.2015].

[29] "Security and Compliance on the Google Cloud Platform – Google Cloud Platform," https://cloud.google.com/security/, 2015, [Online; accessed 06.11.2015].

[30] J. Geffner, "Venom vulnerability," http://venom.crowdstrike.com/, 2015, [Online; accessed 06.11.2015].