Cloud Checkr

whitepaper

THE AUDITOR'S GUIDE TO **AMAZON WEB SERVICES**

Amazon Web Services (AWS) is a leader in the Infrastructure as a Service (IaaS) space. AWS features have matured quickly, and adoption is gaining critical mass. If AWS, and the cloud as a whole, is to be accepted by the enterprise, companies must learn how to utilize the cloud while meeting the compliance and governance standards which already exist in the datacenter. This will require a combination of knowledge and access to the proper tools.

Concerns about security present the biggest objections to cloud adoption. Why does security present an obstacle? First, the cloud is new and there is a general unease surrounding the new and unknown. Second, the initial unease is compounded by the current lack of subject matter experts regarding cloud security. Many of the tool vendors have re-marketed their current offerings as “cloudenabled”, but upon deeper examination they fail to work in a cloud environment. This adds to the uncertainty surrounding cloud security.

Why is cloud security so tricky?

The architecture of a cloud is fundamentally different than that of a traditional datacenter. Cloud deployments aren't necessarily more difficult to secure, but they do require a new approach to security. Clouds are elastic.

This means new servers are ramped up and down in a dynamic and ongoing fashion. Attempting to run a network scan on the cloud does not work for both legal (Amazon terms of service prohibit this) and technological reasons (IP addresses are assigned very differently in the cloud).

Cloud storage is not done on local hard drives or SANs. Instead, data is stored on cloud storage systems such as Elastic Block Storage (EBS) or Simple Storage Service (S3). All these fundamental architecture differences make it imperative that a company develops the technical expertise and finds the right tools designed to operate in this environment.

Physical security questions?

Many people assume that the biggest concern with cloud security is with the cloud service provider itself. Because cloud architectures are multi-tenant, a major concern is that the walls of a shared resource could be breached. Unlike a traditional datacenter, a cloud application is not housed inside a clearly defined DMZ, in a single physical location, or behind a firewall.

In a cloud environment, specifically Amazon Web Services, physical security is handled by the service provider. Amazon, for example, has extensive security procedures in place, and has completed numerous certifications that can be reviewed at http://aws.amazon.com/security/.

Here is the list of certifications achieved at the time this white paper was published:

- SOC 1/SSAE 16/ISAE 3402
- FISMA Moderate
- PCI DSS Level 1
- ISO 27001
- International Traffic In Arms Compliance
- FIPS 140-2
- HIPAA

Physical security concerns are unfounded.

In Fact, Amazon has more stringent physical security policies than even the most security conscious company. In addition to their immense resources, Amazon employs some of the most talented security experts into the world. There are few scenarios in which physical security is not improved by moving to AWS.

A much more important concern when moving to the cloud is network, host, and application level security. These are the responsibility of the application owner, not the cloud provider, and should be implemented based on the best practices of the business.

How data is stored and used on the cloud is outside of the responsibility of the cloud service provider. Are you setting the proper permissions on files stored on S3? Have you configured the Network ACLs correctly? Do you have the appropriate password policies in place? These are all dependent on how you configure, setup, and utilize your AWS account.

Steps for an auditor

So you have been tasked with auditing an AWS account and need to figure out what steps you need to perform. From the high level, there is little that has changed from a traditional IT audit. However, at the task level, the steps can be very different. Here is a guide:

Step 1

Survey the landscape. Get a complete inventory of what's running in the cloud and how it's setup. Remember, if you can't get complete visibility into what's running on the cloud, you can't possibly assess the risk.

This can be very different than surveying the landscape in the data center. In the data center, you might perform a network scan looking for what's running on a specific subnet. In the cloud, you could not just scan a range of IP addresses because they are shared with other AWS users.

Instead you need to use the AWS Management Console or AWS API to get an inventory of what is running under an account.

Step 2

Review best practices, security standards, and policy compliance. Once you've inventoried what is running in the cloud, you need to make sure that there are no problems

There are many different standards to base your assessment on and selecting a standard depends on the industry, the data sitting in the cloud, and who can access the system (are anonymous users allowed access, employees only, or are 3rd-party vendor allowed access).

Step 3

Assess where your highest risks exist, evaluate necessary exceptions, and remediate what needs to be addressed. Many systems will have exceptions to policies that must be reviewed to determine what risk they actually pose. For example, a corporate security policy may dictate that all files should require authentication to be accessed.

In a cloud application, some files may be informational and meant to be accessed by anonymous potential customers. In these cases, it should be acknowledged that these files are meant to have open access.

An auditor should spend a significant amount of time evaluating risks and providing intelligent recommendations based on the inventory and its review.

Step 4

Review again. Once a complete review is performed, analysis is conducted, exceptions are made, and remediation is performed, a follow up review should be completed. This ensures no new issues have come to light, and that the remediation has been done properly.

Conclusion

These four steps should be repeated on an on-going basis. Compliance and security should be setup as a repeatable process.

Cloud deployments are inherently dynamic and elastic. This fluidity requires greater monitoring than a static environment. Consequently, audits should be performed on at least a quarterly basis (ideally, they would be performed more regularly) to ensure compliance is maintained.

Classifying Data

An important part of any audit is identifying and classifying data, particularly sensitive data. The number of objects, files, folders, and buckets sitting on S3 under an account can quickly grow. Getting this content under control is crucial to ensuring you know where to look for inadvertently exposed sensitive information.

There are so many powerful features of S3 for publishing and permissioning objects and buckets that it takes very little to inadvertently expose a file. A common example of this is a feature of S3 that allows the content of a bucket to be published directly to the web. This can be done by setting the property 'Website Endpoint' for the bucket within the AWS Management Console.

See the following description in the Management Console:

"You can host your static websites entirely out of Amazon S3. Once your bucket has been configured as a website, you can access all your content via the Amazon S3 website endpoint for your bucket."

This can create issues if not properly locked down. You can find examples of this by doing a quick search on Google for people that have inadvertently exposed sensitive files this way. Running the following search term on Google.com shows a list of Excel spreadsheets containing the word "salary" exposed to the public internet:

```
salary site:amazonaws.com filetype:xls
```

This query returned 594 results. As you can see, results from Amazon S3 can end up being indexed by Google accidentally. Taking it a step further, you can narrow the Google search to files from your buckets using the following query.

```
site:yourbucketname.amazonaws.com
```

The lesson is to carefully review what you have published and ensure that you aren't inadvertently exposing confidential or sensitive data. Even if you remove the content from your website, you may still need to clear it from the Google cache.

Classifying data requires that you first correctly determine its content type. Then, after the determination, apply the appropriate retention policies, encryption standards, and permissions to the content.

Verify Where Your Data Sits.

Some of the advantages of the cloud can present new complications for compliance. An example of this is the fact that data in the cloud is not necessarily tied to a single physical location. As such, there must be greater understanding of where your data physically sits to ensure you are maintaining compliance with local regulatory rules and laws.

Amazon Web Services uses multiple location layers; the first of which are Regions. Regions are composed of multiple Availability Zones (AZ) which are geographically dispersed. At the time this paper was written, Amazon utilized eight regions: US East (Northern Virginia), US West (Oregon), US West (Northern California), EU (Ireland), Asia Pacific (Singapore), Asia Pacific (Tokyo), South America (Sao Paulo) and the AWS GovCloud.

Within a region, Amazon uses multiple availability zones that are "physically distinct, independent infrastructure, and are engineered to be highly reliable". The idea is that if you design your application across multiple AZs in a single region, it should survive anything short of a catastrophic natural disaster.

As noted above, Amazon offers a special region known as GovCloud. Some sensitive US government applications are subject to compliance regulations such as International Traffic in Arms Regulations (ITAR). ITAR stipulates that data be accessible by US persons only. This requirement is something that would be impossible to manage across all US regions let alone international regions. To meet these standards, Amazon has built a region specifically with US staff only.

Know your region's particular regulations.

Even within countries, specific regulations apply for different states or territories. For example, if you select the US West (Northern California) region, you should be aware of California Data Breach Notification Bill (SB 24) which requires you to notify consumers if their personal data has been breached. This would apply even if your company did not do business in California and if none of your customers were located in California.

By housing your data in a specific region, you must understand the specific responsibilities that decision carries.

This is an important concept in the cloud. Think about how a foreign government could seize data sitting in an Amazon data center.

Recall back to an incident in Turkey in 2007 in which a Turkish court ordered all WordPress blogs blocked because some of the blogs hosted on the platform were declared defamatory by the Turkish court (http://en.blog.wordpress.com/2007/08/19/why-were-blocked-in-turkey/).

The Amazon CloudFront Edge Network (http://aws.amazon.com/cloudfront/) uses a network of edge locations around the world including the United States, Europe, Asia, and South America to provide low-latency access to data.

When you create a distribution, you need to understand completely what content is being distributed, what the risks of it being seized are, and the likelihood of that occurring. You should carefully consider any other potential risks geographic and political situations may cause.

Verify Best Practices.

Best practices provide a standard to use in implementing a business process and are the result of experience and thoughtful consideration. Best practices are usually not "quick and dirty". They are, however, the best way to get things done in the long term.

Clouds are typically setup with less formality and rigorousness. This often leads to best practices being de-prioritized. It's an auditor's job to put these best practices back into play.

Best Practices

Below is a list of best practices that should be considered for your AWS accounts:

#1 – Use of the AWS account should be limited to creating the account and billing management. Administration of computing and storage resources should be delegated and executed through users created within the Identity and Access Management (IAM) service.

"Amazon recommends that the AWS account not be interacted with at the AWS account level. Instead users should be created and used to interact with the system."

#2 - Require use of Multi-factor Authentication for the AWS account and other users with full administrative privileges. MFA is simple to configure in AWS and an MFA device can be purchased for as little as $15, making it a good option for increased security.

#3 - Use groups for all access control. Permissions should be granted to groups only, not directly to users. This allows permissions to be revoked and granted in a repeatable fashion so that mistakes are not made resulting in improper security grants.

#4 - Require strong password policies to be in place. Under the IAM service, a password policy should enabled.

#5 – Under the IAM password policy, require a minimum password length of 6 or more characters.

#6 - Under the IAM password policy, require at least one uppercase letter to be used in every password.

#7 - Under the IAM password policy, require at least one lowercase letter to be used in every password.

#8 - Under the IAM password policy, require at least one number to be used in every password.

#9 - Under the IAM password policy, require at least one non-alphanumeric character.

#10 - Under the IAM password policy, allow users to change their own password.

#11 – Create an administrators group to use for controlling the list of users with administrative privileges. From the IAM Getting Started Guide

(http://docs.amazonwebservices.com/IAM/latest/GettingStartedGuide/SetUpAdminsGroup.html):

"Having an administrators group for your AWS account isn't required, but we strongly recommend it."

#12 – Consider converting large files to Reduced Redundancy Storage. Review each file greater than 100 GB to verify how critical the file is. Suggest that very large files which are non-critical and reproducible be converted to Reduced Redundancy Storage. Reduced Redundancy costs are between 25% and 33% less than Standard Storage.

For instance, the first 1 terabyte of disk space on S3 is $0.125 per GB per month. For Reduced Redundancy Storage the analogous cost is $0.093 per GB per month.

While the price is lower, the durability is reduced. Standard Storage provides 99.999999999% durability while Reduced Redundancy Storage provides 99.99% durability.

#13 - Check for S3 buckets, folders, or objects granted the permission "Open/Download" to Everyone. This permission allows anonymous users to read a file but does not allow them to update the file.

Review all cases in which this permission is granted to verify that it is appropriate.

#14 - Check for S3 buckets, folders, or objects granted "View Permissions" and "Edit permissions" to Everyone. This permission allows anonymous users to edit or delete a file.

Review all cases in which this permission is granted to verify that it is appropriate.

#15 - Check that logging is enabled for S3 buckets. This creates a log file of all accesses to the objects and bucket. Without this option, you will be unable to track and identify who has accessed a file/object in S3.

#16 - Check that notifications have been enabled for S3 buckets. There is a property for each S3 bucket:

"Enabling notifications causes a message to be published to an Amazon Simple Notification Service (SNS) Topic when Amazon S3 detects that a Reduced Redundancy Storage object stored in this bucket is lost."

If you do not enable this option, you will not be notified when an S3 object is corrupt or lost.

#17 - Check for S3 buckets that have a website endpoint enabled. Verify the permissions and content within the directory to verify that no sensitive data has been exposed.

#18 - Verify automatic backups are enabled for Amazon RDS. In March of 2012, the maximum retention period for automated backups was increased from eight days to thirty five days.

It is recommended you change the backup retention period to 30 days.

#19 - Use the InnoDB Engine with RDS MySQL. From the RDS FAQ

(http://aws.amazon.com/en/rds/faqs/):

"Amazon RDS automated backups and DB Snapshots are currently supported for the InnoDB engine only. Use of these features with other MySQL engines, including MyISAM, may lead to unreliable behavior while restoring from backups. Specifically, since storage engines like MyISAM do not support reliable crash recovery, your tables can be corrupted in the event of a crash. For this reason, we encourage you to use the InnoDB storage engine."

#20 - In Amazon RDS, verify that the master username is not the default "awsuser".

#21 - In Amazon RDS, verify that master password is not the default "mypassword".

#22 - In Amazon RDS, verify that there are no orphaned Read Replica instances. From the RDS FAQ

(http://aws.amazon.com/en/rds/faqs/):

"A Read Replica will stay active and continue accepting read traffic even after its corresponding source DB Instance has been deleted. You must explicitly delete the Read Replica DB Instance."

#23 - In Amazon RDS, verify that the DB Security Groups are not set to broad IP ranges. From the Amazon RDS user guide (http://awsdocs.s3.amazonaws.com/RDS/latest/rds-ug.pdf):

"Ensure you authorize only specific IP ranges or EC2 security groups. We highly discourage authorizing broad IP ranges (for example, 0.0.0.0/0)."

#24 - In Amazon Virtual Private Cloud (VPC), verify that the security groups are allowing only HTTP/HTTPS traffic from anywhere and only inbound SSH and RDP.

Additional network access is seldom required. See the Amazon VPC Getting Started Guide:

http://docs.amazonwebservices.com/AmazonVPC/latest/GettingStartedGuide/SecurityGroup.html

#25 - Check that all Virtual Private Gateways have both tunnels enabled. See the Amazon VPC Network Administrator Guide:

"Note that the VPN connection consists of two separate tunnels: Tunnel #1 and Tunnel #2. Two redundant tunnels provide an increased availability in the case of a device failure."

Summary

As you can see, with just a small amount of digging there are plenty of best practices to review and follow. Each of the Amazon Web Services has its own set of configurations and setup. If a specific service in AWS is being used it should be carefully reviewed to verify that it is following these as well as all other generally accepted best practices.

The cloud is here.

It offers compelling cost and efficiency benefits. Startups have already embraced the cloud. We hope that the steps in this white paper will start you on the right path to providing the knowledge, processes and procedures to insure that your Enterprise clients also enter the cloud successfully.

About the Author

Aaron C. Newman, a founder of CloudCheckr Inc., is a world renowned security expert. He authored The Oracle Security Handbook (published by Oracle Press) and has spent over 20 years working in the security arena. He has created viable security solutions for both the database and operating system levels.

About the Company

CloudCheckr Inc. is a Rochester, NY software company that specializes in bringing clarity to cloud deployments for enterprise cloud users. CloudCheckr Inc. leveraged its team's expertise and knowledge to develop CloudCheckr (www.cloudcheckr.com).

CloudCheckr uses read only credentials to fully inventory your cloud deployment. It then formats this information into actionable knowledge, cost, and best practice reports and checks.