Universiteit van Amsterdam

MSc System & Network Engineering

Automated analysis of AWS infrastructures

P.G.J. Bennink

Supervisor: Cedric van Bockhaven

July 10, 2018

Abstract

This project intends to find out what information can be gathered about an Amazon Web Services (AWS) infrastructure based on limited access gained through the infiltration of an EC2 instance and/or retrieved access keys. To this end we analyzed the AWS platform as a whole, the relationships that can exist between components, and what factors influence access to information about the infrastructure. We then built a tool that automatically gathers and parses information about the infrastructure via the AWS CLI to enumerate the components and the relationships between these components. It then imports this data into a graph database that allows for visualization of the data. This tool is able to combine the information multiple access keys can retrieve into one resulting database of the infrastructure, and can also include components that can be seen, but not yet accessed. We thus conclude that this type of AWS infrastructure analysis is useful for reconnaissance and security audits. Future work includes expanding the amount of supported services and commands.

Contents

1 Introduction **2**
1.1 Research question 3

2 Related work **4**

3 Methodology **4**

4 Analysis **5**
4.1 IAM 5
4.2 Information 8
4.3 Isolation 8

5 Development **9**
5.1 Crawling metadata 9
5.2 Brute forcing permissions 9
5.3 Data gathering and processing 10
5.4 Visualizing infrastructure 11

6 Conclusion **12**

7 Discussion **12**
7.1 Future work 13

1 Introduction

AWS is "a secure cloud services platform, offering compute power, database storage, content delivery and other functionality to help businesses scale and grow."[5] When someone successfully infiltrates a component of an AWS infrastructure, most of the time the next step is to try and expand this access to other parts of the infrastructure. However, this is difficult without any further knowledge of the infrastructure, and a lot of trial and error would be required. So before proceeding, reconnaissance of the infrastructure should be done to get a better idea of the components the infrastructure is made up of, the relationships between these components, and the best way to ultimately obtain access to valuable data within this infrastructure. Another important factor is finding out which of the found components can already be accessed with the found credentials.

Bloodhound, which relies on Microsoft's Active Directory to "easily identify highly complex attack paths that would otherwise be impossible to quickly identify within an Active Directory environment"[7], offers solutions to the problems specified above. However, the problem Bloodhound encounters is different, as Active Directory is generally one centralized directory containing all related entities. On AWS this is not the case, and data has to be sourced using a multitude of commands, sometimes executed multiple times on different components of the same type.

The goal of this project is thus to design a tool that crawls through an AWS infrastructure, identifies as many components as possible, and shows the relationships between the different components. In the analysis for and development of the tool we will pay attention to finding components of the infrastructure that we currently cannot yet control. This is specifically important in red team operations, where expanding access is one of the immediate goals. Another use case would be infrastructure administrators who are looking for misconfigurations, e.g. components that have more access than necessary. Based on the results we discuss what obstacles there still are in this type of analysis, and how they could be solved in the future.

AWS uses a service called Identity & Access Management (IAM) to define the privileges of access keys that entities (users, components or services) use to interact with the infrastructure. One of the problems this project aims to solve is the lack of knowledge about the privileges an access key has.

The tool built for this project relies on boto3, the Python AWS SDK[1], Neo4j, a graph database often used for the mapping and visualization of networks of any kind[11], and py2neo, a "a Python library and toolkit for working with Neo4j."[4] The tool has two functions. Firstly, for every AWS access key that it receives as input it will try to obtain information about the AWS infrastructure. It will then parse this information and put it in the database for the user to analyze. Secondly, the tool tries to access a predefined amount of AWS CLI commands, and for each command return whether that access key has the permission to execute it. The first function will only use a subset of commands that can be used to gather data, while this function will also try commands used to control the infrastructure and its components. This shows the user in what way they would be able to escalate privilege and/or expand access.

1.1 Research question

The research question will be defined as follows:

Given an infiltrated AWS component, what part of the related infrastructure would an automated tool be able to index?

The AWS component we chose is an EC2 instance, which is one of the services AWS offers and will be further explained in Section 4. The reason we assume an EC2 instance is the fact that EC2 instances are often the public-facing component within an infrastructure, which is thus a realistic component to have infiltrated first. Shell access to an EC2 instance also allows one to retrieve that instance's security credentials, which are the credentials this project uses to gather data about the infrastructure.

While for this project we assume an infiltrated AWS component, the results can be generalized to include any situation in which someone has obtained security credentials to an AWS infrastructure.

What should be made clear at this point is that the tool that was built during this project does not compromise any AWS components itself. This tool is meant for reconnaissance of infrastructures to which the user of the tool has gained access in a lawful manner (whether it be their own infrastructure or that of a consenting third party). Furthermore, all access to the AWS infrastructure happens via the AWS CLI.

2 Related work

While a certain amount of AWS analysis and/or visualization tools exist, most of these are meant to be used in combination with admin-credentials. In our situation, these would be the semi-metaphorical "keys to the kingdom", which we cannot assume we have access to for this project.

Nimbostratus is a tool developed by Andres Riancho, which he presented at Black Hat USA 2014.[6] This tool automates the retrieval of access keys and metadata from an EC2 instance as well as the exploitation of a few commands (if that EC2 instance has access to those). While the focus of Nimbostratus lies more with the exploitation of an EC2 instance than with the analysis of the greater infrastructure, this was a starting point for our analysis of AWS and possible weaknesses in (configurations of) IAM.

Another very similar tool is CloudMapper[8], developed by Duo Labs. The difference between this tool and the tool we are creating in this project is the situation in which it can be used. CloudMapper can be used when accessing a specific set of Actions is allowed. If access to even one of these is not possible, this tool will refuse to work. It is not so much a tool used for penetration testing as it is meant for analysis of an infrastructure to which a relatively high (and very specific) amount of access has already been established. This is not viable in the context of our project, as we do not know beforehand what access we have with each key, and in a lot of cases we will not be able to expand the access of an access key ourselves. Furthermore, Cloudmapper does not support inputting multiple keys and creating **one** mapping based on multiple keys with different permissions.

3 Methodology

The analysis of AWS was done by using the AWS platform (via the web interface, the CLI and the Python SDK), creating multiple infrastructures and reading the documentation. The development of the tool made for this project was done in Python 3. For the testing of the tool the aforementioned infrastructures were used.

The parts this project was made up of listed above were executed in an iterative manner due to the fact that the testing of the tool during development gave more insight into the working of AWS, and thus into what way the tool could be improved. We have nevertheless tried to separate the different parts in this report to make it clearer.

4 Analysis

AWS offers a vast amount of services. Defining what is part of 'the infrastructure' can be difficult, as a lot of the services AWS offers play a supporting (but nevertheless equally important) role in the infrastructure. Including all of these services in this project would not be possible, as each service requires a different method for interacting with it. Due to this we scoped this project to only include the following main services:

- Amazon Elastic Compute Cloud (Amazon EC2), which is "a web service that provides secure, resizable compute capacity in the cloud."[10] EC2 components (or instances, as they are called on AWS) are VMs ran on a hypervisor. These VMs can run any of the more than 70,000 Community AMIs (Amazon Machine Images), which are free, or one of the AMIs offered on the AWS Marketplace, which are paid. are the backbone of each infrastructure. They control all the other components in the infrastructure.
- Amazon S3, which is "object storage built to store and retrieve any amount of data from anywhere."[10]
- Amazon Relational Database Service (Amazon RDS), which "provides cost-efficient and resizable capacity while automating time-consuming administration tasks such as hardware provisioning, database setup, patching and backups."[10]

However, a setup with these services already makes use of other supporting services. Of these services one is the most important, namely IAM, as it is the security backbone that allows entities (users, components, services) to interact with one another.

4.1 IAM

In Amazon's words, IAM "enables you to manage access to AWS services and resources securely. Using IAM, you can create and manage AWS users and groups, and use permissions to allow and deny their access to AWS resources."[10] IAM uses a combination of an Access Key ID, a Secret Access Key and optionally a Token to allow users to interact with the infrastructure. Access Key IDs are public, and can thus, with the right privileges, be found using the AWS CLI. However, once created the Secret Access Key and Token cannot be retrieved from AWS again, meaning that if a user loses either of these keys they will have to generate new keys. Tokens are only used for temporary credentials.

IAM allows for the creating of groups, users, roles, and policies, the latter of which is the basis for the whole platform.

```
{
    "Version": "2012-10-17",
    "Statement": {
        "Effect": "Allow",
        "Action": "s3:ListBucket",
        "Resource": "arn:aws:s3:::example_bucket"
    }
}
```

Listing 1: An example of an AWS policy

Listing 1 show an example of a policy. A policy is made up of statements that either Allow or Deny a user or entity to access some command or resource. The important parts of these statements are `Effect`, `Action`, and `Resource`. `Effect` can either be Allow or Deny.

Figure 1: Decision flow chart of policy evaluation.

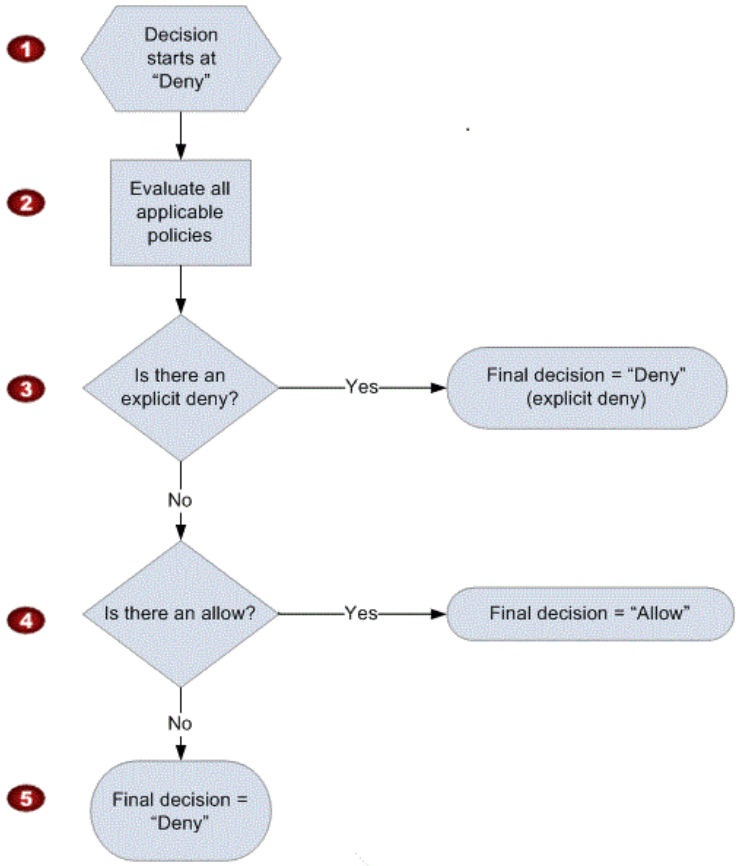

Source: AWS IAM Documentation

As can be seen in Figure 1, by default (if there is no policy that allows or denies a user or entity to do something) the access will be denied. If there is both an Allow and a Deny (in whatever order), access will also be denied. Only if the applicable policies solely explicitly allow an action will it be allowed.

`Action` specifies what actions the statement allows or denies. These actions are formatted as `Service:Action`, as multiple services use the same name for the same action. Wildcards are allowed, meaning that `*:List*` allows all List-actions to be used on all services.

In Listing `Resource` we can define on what specific resource an action can be used. In 1 the `S3:ListBucket` command can only be used on the bucket with the name `example_bucket`.

In terms of information necessary for the enumeration of components of the infrastructure, the `*:List*` and `*:Describe*` actions are the most important, as these are the commands that will statically retrieve information about the infrastructure without making modifications. What makes some of these commands even more interesting, is that their output cannot be controlled at a resource-level, meaning that if for instance `EC2:DescribeInstances` gets executed, it will either show all or none of the EC2 instances, which makes this a valuable source of information for this project. Amazon keeps a list with these types of commands, which can thus be used for this project.[3]

Groups contain users, and policies can either be attached to users directly or by creating a group of users and attaching policies to that group. Roles can be used to attach policies to services or instances.

Roles are basically groups for users, and can be attached to as many instances as necessary. If for example an EC2 instance needs to interact with an S3 bucket, a user with rights to make modifications in the IAM service can attach a role to an EC2 instance. They can then attach a policy to that role that allows all entities in that role to access all (or specific) S3 buckets in various ways (depending on what the instances need to do with the bucket).

These roles are interesting, as it could very well be that two instances that show an overlap in their required permissions get attached to the same role. This means that both instances have more access than they actually need, which can be a security risk (the severeness of which is of course dependent on the `Actions` and `Resources`.

As with all other IAM policies, the instances need access keys to make use of their permissions. These access keys get distributed via a metadata server, which each instance can access via `http://169.254.169.254`. This means that if one in whatever way has gained shell access to an EC2 instance, they also have access to very detailed metadata for that specific instance, which includes access keys with all permissions that that EC2 instance has. This is thus one way in which someone might be able to retrieve access keys from the AWS platform.

4.2 Information

All services and their actions can be controlled via either a CLI or one of the SDKs that Amazon offers. What becomes clear after analysis of the output of these commands is the fact that there is an overlap in information between these commands. For almost all types of components an `Action` exists that will enumerate all components of that type, `EC2:DescribeInstances` being an example of such a command. In most cases these types of commands also offers a lot of information about the surrounding components. In the case of `EC2:DescribeInstances` we are able to retrieve information about the availability zones, subnet IDs, VPC IDs, mounted volumes, security groups and owner ID of these instances. This is all information related to those instances, but also related to other components in the infrastructure.

Another example is that when access to `EC2:DescribeVpcs` (which lists the different VPCs with information about each) is prohibited, the existence of certain VPCs can still be confirmed based on the data about, for instance, an EC2 instance that resides in that VPC. As this project is about reconnaissance (with the intention of expanding access) this is important information that will be used in our resulting tool.

In a situation where access to for instance `EC2:DescribeSecurityGroups` is prohibited, a decent amount of information about the existence of these security groups can still be obtained via other Actions.

4.3 Isolation

As explained in Section 4.3, the information gained from lower-level components can be used to gain knowledge about the higher-level components. Conversely, one can also use the information of higher-level components to gain information about lower-level components.

AWS also allows users to place the components of their infrastructure in a Virtual Private Cloud (VPC). Among the services that support this are RDS and EC2. These VPCs contain security groups, which can be used to limit the amount, type, source and destination of traffic for components within that group.

These security measures are meant to make malicious use more difficult, as it just completely blocks all connections that do not abide by the rules of that security group. However, with access to `EC2:DescribeSecurityGroups`, which enumerates all security groups and the rules attached to them, making an assumption about what is inside the security group (based on properties like the group name and the allowed ports/IPs) would still be possible, and thus decide whether gaining access to this security group and its contents would be useful.

5 Development

The tool created for this project can be split up into three parts in terms of functionality, which correspond to the subsections below. Firstly, the crawling of an EC2 instance's metadata server for all metadata, including its IAM credentials. Secondly, the enumeration of permissions of all access keys that are given as input through bruteforcing, which may or may not include access keys obtained using the previous function. And thirdly, the gathering, processing and importing of data about the AWS infrastructure to Neo4j. The fourth section below describes the resulting output, namely the visualization of the database in Neo4j.

The first part is written in bash, and should be executed on an EC2 instance to which shell access has been gained, as these metadata servers can only be reached from or via the instance itself. While there are situations in which one can remotely obtain data from the metadata server (see Nimbostratus[6]), this is not within the scope of this project.

The second and third parts are written in Python 3, and can be executed from any computer with an internet connection that can run Neo4j, Python 3, and some Python packages. The access keys have at this point already been obtained, meaning there is no reason to run this on the infiltrated EC2 instance.

5.1 Crawling metadata

Once shell access to an EC2 instance has been achieved it is very easy to access the metadata server. To make it somewhat easier to crawl through the server we wrote a short bash script that crawls through the folders on the server and downloads all of the contents. It will then output the access keys found on the server to the user, who can then use those keys (among potential other found keys) as input for the second and third functions.

5.2 Brute forcing permissions

For all of the access keys the user has obtained this part of the tool tries to execute a predefined amount of commands to offer the user some information about the value of that key in the larger scheme of the infrastructure. For this project the tool ran through 25 different commands from the services EC2, RDS and S3. The EC2 service also includes commands about VPCs, the data volumes of the EC2 instances, and the security groups. We will discuss this further in section 7.1, but the reason we did not include more commands (both from the services specified above and other services) is mainly related to time. However, the technique used for this test can be extended to include other commands as well. Brute forcing of these permissions is done via boto3.

5.3 Data gathering and processing

As explained in Section 4, a lot of `*:Describe*` commands contain information about surrounding components. This is used in creating the relationships in the database. For instance, each security group is contained in some VPC. Which VPC this is can be found in the output for `EC2:DescribeSecurityGroups`. In the importing of the data a top-down approach is used: first `EC2:DescribeVpcs` is executed, if this returns data it will get imported in the database. After that `EC2:DescribeSecurityGroups` is executed. If any of the groups in the output of this command are located in a VPC that is not already in the database it will get added now, and a relationship between the security group and its VPC will be added. This same process is repeated for EC2 and RDS components, with relationships added between these components and their respective security groups, and S3 buckets (which do not have security groups). All of the commands on AWS are executed using boto3, and the importing of the data to the database is done using py2neo.

Now the access keys and their permissions need to be imported in the database. For this the results of the functionality described in Section 5.2 are used. For each of the commands used in that section we specified what type of component they are related to (e.g. `EC2:CreateDefaultVpc` is related to VPC). As in the previous example, the service to which the command belongs is not always that of the type of component, so this had to be added manually. If a command can be executed from an access key, a relationship between that access key and the related components will be made.

5.4 Visualizing infrastructure

Neo4j has its own query language (Cypher[2]) and a web interface which can be used to interact with the database (similar to PhpMyAdmin, but integrated). This web interface allows you to query the database and visualize the output, which is what we use for the visualization in this project.

Figure 2: Output of Neo4j query for example infrastructure.

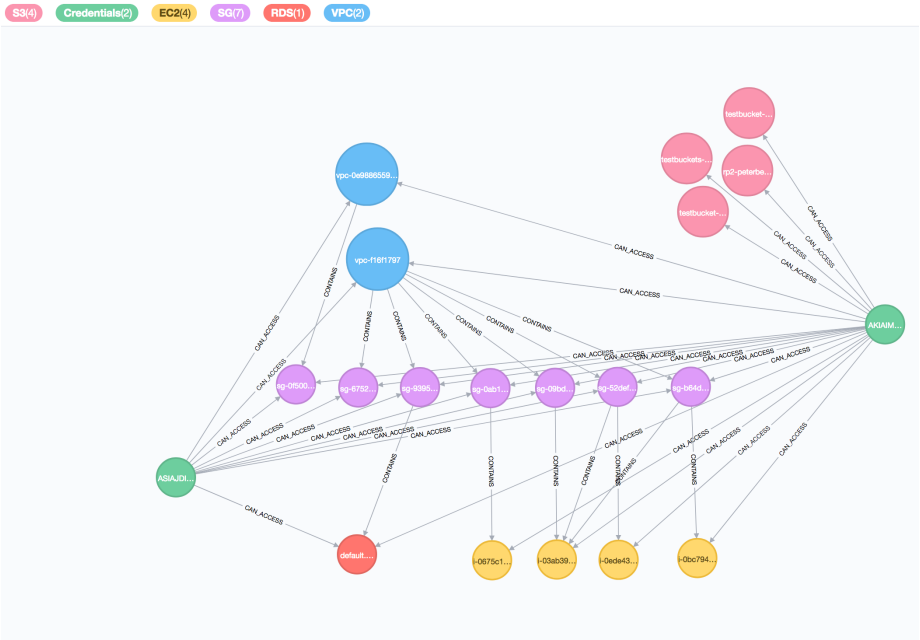

An example of this visualization is shown in Figure 2. The infrastructure shown here is a version of testing infrastructure used during this project. The directional edges between the component-nodes (all colors apart from the green nodes) indicate that the originating node contains the destination node, e.g. all security groups are contained in one of the two VPCs. The edges from the access keys (the green nodes) to the different components indicate that that access key can in some way access the destination node.

All nodes and edges have properties. The properties of the nodes offer more information about that component or access key. In the case of the components it will show the information that was retrieved about that component from the AWS CLI, in the case of the access keys it shows the Secret Access Key and potentially a token. The properties of the edges from the access keys indicate what commands can be executed on that component with that access key.

The figure shown above shows almost all types of components currently in the database. The volumes connected to the EC2 instances (yellow) are not shown. This is an example of what Neo4j can do. Just like with other query languages, you can query the database to only include certain types of nodes or relationships, based on label (e.g. 'EC2' or 'VPC') or properties (not shown in Figure, but an example would be 'VpcId' for the VPC nodes). This allows the user to query the database based on the components or relationships they are interested in.

6 Conclusion

Our research question was: **"Given an infiltrated AWS component, what part of the related infrastructure would an automated tool be able to index?"** It would be impossible to answer this question with something like a percentage or a calculation, as this all completely depends on the access key(s) acquired, what permissions they have, whether they overlap,

One of the main takeaways of this project is that AWS's default policies allow for a relatively large amount of enumeration. This is among other things due to the fact that a large amount of `*:Describe*` commands, which return information about all components of a specific type, cannot be restricted on a resource-level, meaning that if someone can access `EC2:DescribeInstances` they will be able to enumerate **all** EC2 instances, including a large amount of information that can then be used in the infiltration of those instances.

In terms of the access keys and their permissions, it seems that access keys do not need to be very privileged to be useful in terms of indexing the infrastructure. While higher level privileges might be useful from an infiltration perspective, the enumeration of components and their relationships does not benefit from this per se.

The tool created during this project is open source and available on Gitlab.[1]

7 Discussion

In general it seems that indexing an infrastructure is definitely possible in most cases, even though this does not hold for the whole infrastructure. While it is difficult to make a statement about this possibility for all infrastructures, access keys and/or components some conclusions can be drawn, as has been shown in the previous section. Based on these conclusions we are at least able to say that it would be worthwhile to pursue the development and research behind this tool further.

[1] AWS Infrastructure Analysis - Gitlab

7.1 Future work

The tool created during this project can be expanded, refined and improved upon in a vast number of ways. While this is mainly due to the vastness of AWS as a service, other factors have also contributed to this. The decisions made during this project were influenced by the limited time and the specific research question, and in some cases would have been different if this tool would be developed over a larger amount of time.

Path finding The most obvious improvement to this tool would be automated path finding, whereby you can input two nodes and get the shortest path between these two nodes as output. This is a feature in Linkurious[9], which allows users to interact with their graph-based data in an intuitive and graphical way. Bloodhound uses Linkurious.js for this purpose, which has been deprecated since 2016 and integrated in a new tool by the same organization called Ogma[2]. Integrating this would thus be a logical next step in the development. However, after inquiring with Linkurious it seems the minimal costs for using this library would be much more than would be afforable. Another solution with similar functionality will thus have to be found.

Further testing A second potential improvement would simply be further testing of this tool on larger infrastructures. Currently we have only tested this on our own testing infrastructure, which we kept expanding during the infrastructure, but further testing could reveal new ways in which the tool should be improved. While not in time for this project we are in talks with 3rd parties to test this tool on their infrastructure.

Penetration testing toolkit The tool created for this project is, in its current form, exclusively meant for reconnaissance. However, one way in which this tool could become part of a larger toolkit is by creating something comparable to Metasploit, but for AWS. Functionality in this toolkit could include scanning for and exploiting commands that can be used in privilege escalation, automatically scanning S3 buckets or EC2 volumes for high entropy contents (which could contain access keys or SSH private keys), or performing dictionary attacks to find usernames or bucket names.

Expansion An obvious improvement would be making the tool compatible with more AWS services. While the services in the scope of our project are among the most often used services[3], a lot of infrastructures will also use other services, and even our testing infrastructure uses more services than we analyzed for this project. An example of an interesting service to include is STS (Security Token Service) which among other things allows an entity with an access key to assume the role of another entity, which might have more or different privileges.

[2] Ogma - Linkurious

[3] Most popular AWS products of 2017 - globenewswire.com

Intelligent bruteforcing The permission bruteforcer created for this project checks a list of commands specified beforehand. It will check this whole list for each access key given. A more silent way to do this would be to prioritize the commands, and let the user of the tool choose how 'loud' the tool can be, much like the T-flag in Nmap allows its user to indicate the timing intervals between the different probes to stay under the radar of, for example, intrusion detection systems.[4] AWS allows an administrator to see which commands are executed using which access keys. If suddenly every possible command gets executed from multiple access keys this should raise some alarms. Prioritization could help prevent these alarms, making the reconnaissance more silent.

Resource-level permissions In its current form the bruteforcer only tries commands that do not support resource-level policies. Adding commands that can be restricted at resource-level would involve trying each command, for each access key, for each component that that command can be used on.

[4]Nmap manual: "Timing and Performance" - nmap.org

References

[1] AWS SDK for Python. https://aws.amazon.com/sdk-for-python/.

[2] Cypher Query Language Developer Guides & Tutorials. https://neo4j.com/developer/cypher/.

[3] Granting IAM Users Required Permissions for Amazon EC2 Resources. https://docs.aws.amazon.com/AWSEC2/latest/APIReference/ec2-api-permissions.html#ec2-api-unsupported-resource-permissions.

[4] The Py2neo v4 Handbook. http://py2neo.org/v4/.

[5] What is AWS? Amazon Web Services. https://aws.amazon.com/what-is-aws/.

[6] Andresriancho. andresriancho/nimbostratus. https://github.com/andresriancho/nimbostratus, Feb 2014.

[7] BloodHoundAD. BloodHound - Github. https://github.com/BloodHoundAD/Bloodhound/wiki.

[8] duo labs. Cloudmapper - Github. https://github.com/duo-labs/cloudmapper, Jun 2018.

[9] Linkurious. The graph intelligence platform. https://linkurio.us/.

[10] Jinesh Varia and Sajee Mathew. Overview of amazon web services. *Amazon Web Services*, 2014.

[11] Jim Webber. A programmatic introduction to neo4j. In *Proceedings of the 3rd annual conference on Systems, programming, and applications: software for humanity*, pages 217–218. ACM, 2012.