

James Witts - S14139866
Ahmed Ismail - S15120305
Cheng Jiacheng - S16153760

CMP6203 - Cloud Computing: AWS Report

This report describes **six** main stages in a project using *Amazon Web Services (AWS)* from initial development on a local machine through deployment, testing and concluding with some remarks for further development

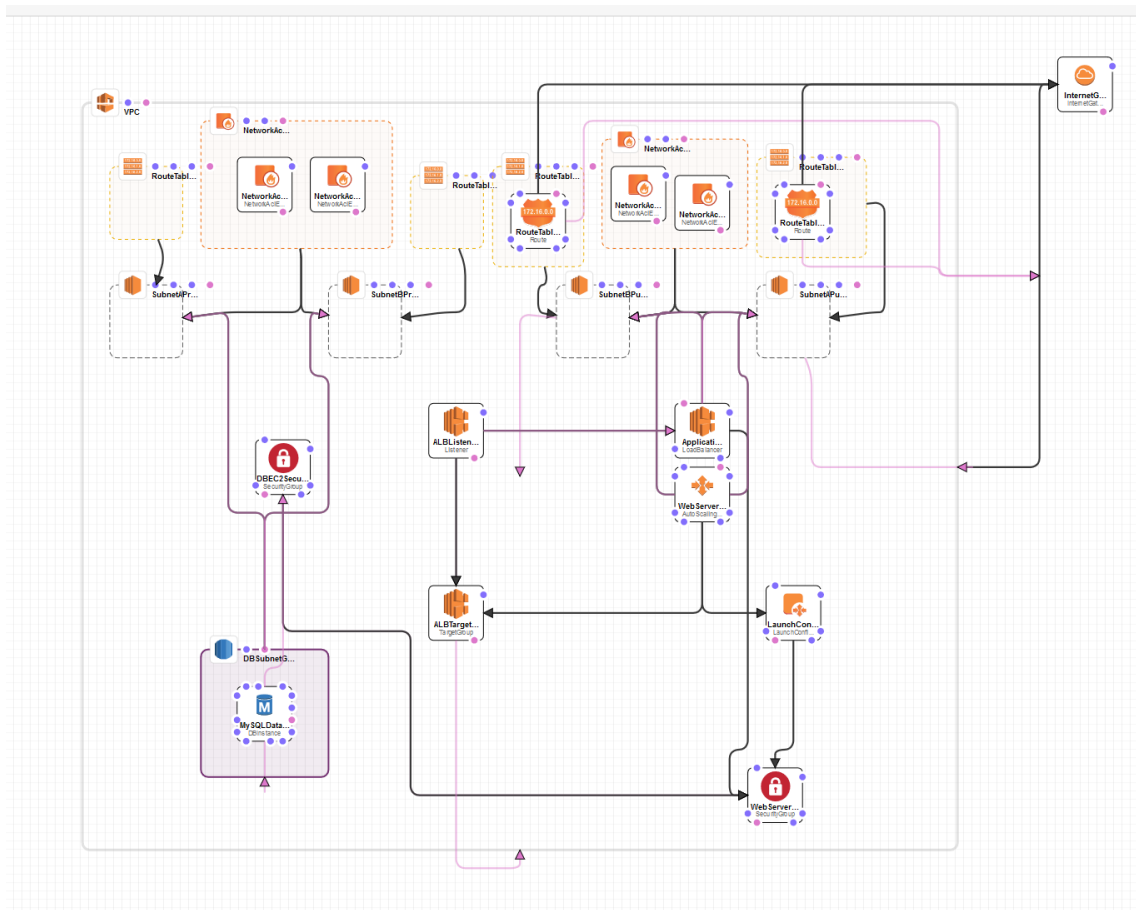
STAGE ONE: Local development

To start with we developed our application on a windows 8 laptop using XAMPP Lamp stack from Bitnami that uses a MySQL database running on port 3306. The application is a php web app that allows a user to insert data into one of three different tables; users, bookmarks or tags. The default engine for the MySQL database is MyISAM. In the future it might be better to changes this to Innodb since the “tags” table we created uses a foreign key dependent on user data, which will be replaced in Innodb by a unique identifier when running from an Amazon RDS MySQL database.

STAGE TWO: Deployment - North California region

When moving the XAMPP stack to Amazon we first prepared all the resources we needed by using a CloudFormation stack template which we made in CloudFormation Designer by merging two existing templates together from two different sources. One came from a git repository and the other a sample template provided by CloudFormation Designer. When merging the two together the workflow was mostly trial and error and we eventually got a stack that created most of the resources we needed.

A schematic overview of the full design is given in Appendix I.



The following diagram shows the output from creating a stack using the *Cloudformation* tool:

The screenshot displays the AWS CloudFormation console interface. At the top, there are navigation menus for 'Services' and 'Resource Groups'. The main header shows 'CloudFormation' and 'Stacks'. Below this, there are buttons for 'Create Stack', 'Actions', and 'Design template'. A filter section is set to 'Active' and 'By Stack Name'. The main content area shows a table with columns for 'Stack Name', 'Created Time', 'Status', and 'Description'. The stack 'AwsCloudCompStackNorthCv2' is listed with a status of 'CREATE_COMPLETE' and a description: 'VPC; public and private subnets in two availability zones, a cloudonaut.io template modified by James Ahmed Jiacheng'. Below the table, there are tabs for 'Overview', 'Outputs', 'Resources', 'Events', 'Template', 'Parameters', 'Tags', 'Stack Policy', 'Change Sets', and 'Rollback Triggers'. The 'Events' tab is selected, showing a list of events with columns for 'Status', 'Type', 'Logical ID', and 'Status Reason'. The events are filtered by 'Status' and show a sequence of 'CREATE_COMPLETE' and 'CREATE_IN_PROGRESS' events for various AWS resources like RDS instances, Elastic Load Balancing listeners, and subnets.

Stack Name	Created Time	Status	Description
AwsCloudCompStackNorthCv2	2018-05-16 01:00:49 UTC+0100	CREATE_COMPLETE	VPC; public and private subnets in two availability zones, a cloudonaut.io template modified by James Ahmed Jiacheng

Time	Status	Type	Logical ID	Status Reason
01:17:39 UTC+0100	CREATE_COMPLETE	AWS::CloudFormation::Stack	AwsCloudCompStackNorthCv2	
01:17:32 UTC+0100	CREATE_COMPLETE	AWS::RDS::DBInstance	MySQLDatabase	
01:03:44 UTC+0100	CREATE_COMPLETE	AWS::ElasticLoadBalancingV2::Listener	ALBListener	
01:03:43 UTC+0100	CREATE_IN_PROGRESS	AWS::ElasticLoadBalancingV2::Listener	ALBListener	Resource creation initiated
01:03:43 UTC+0100	CREATE_IN_PROGRESS	AWS::ElasticLoadBalancingV2::Listener	ALBListener	Resource creation initiated
01:03:39 UTC+0100	CREATE_COMPLETE	AWS::ElasticLoadBalancingV2::LoadBalancer	ApplicationLoadBalancer	
01:01:54 UTC+0100	CREATE_COMPLETE	AWS::EC2::SubnetRouteTableAssociation	RouteTableAssociationAPublic	
01:01:53 UTC+0100	CREATE_COMPLETE	AWS::EC2::SubnetNetworkAcAssociation	SubnetNetworkAcAssociationBPublic	
01:01:53 UTC+0100	CREATE_COMPLETE	AWS::EC2::SubnetNetworkAcAssociation	SubnetNetworkAcAssociationAPrivate	
01:01:53 UTC+0100	CREATE_COMPLETE	AWS::EC2::SubnetRouteTableAssociation	RouteTableAssociationBPublic	
01:01:52 UTC+0100	CREATE_COMPLETE	AWS::EC2::SubnetRouteTableAssociation	RouteTableAssociationBPrivate	
01:01:52 UTC+0100	CREATE_COMPLETE	AWS::EC2::SubnetRouteTableAssociation	RouteTableAssociationAPrivate	
01:01:52 UTC+0100	CREATE_COMPLETE	AWS::EC2::SubnetNetworkAcAssociation	SubnetNetworkAcAssociationBPrivate	
01:01:51 UTC+0100	CREATE_COMPLETE	AWS::EC2::SubnetNetworkAcAssociation	SubnetNetworkAcAssociationAPublic	
01:01:50 UTC+0100	CREATE_COMPLETE	AWS::EC2::Route	RouteTablePublicInternetRoute	
01:01:50 UTC+0100	CREATE_COMPLETE	AWS::EC2::Route	RouteTablePublicBInternetRoute	
01:01:43 UTC+0100	CREATE_IN_PROGRESS	AWS::RDS::DBInstance	MySQLDatabase	Resource creation initiated
01:01:42 UTC+0100	CREATE_IN_PROGRESS	AWS::RDS::DBInstance	MySQLDatabase	Resource creation initiated
01:01:39 UTC+0100	CREATE_COMPLETE	AWS::AutoScaling::AutoScalingGroup	WebServerGroup	
01:01:39 UTC+0100	CREATE_IN_PROGRESS	AWS::EC2::SubnetRouteTableAssociation	RouteTableAssociationAPublic	Resource creation initiated
01:01:38 UTC+0100	CREATE_IN_PROGRESS	AWS::AutoScaling::AutoScalingGroup	WebServerGroup	Resource creation initiated
01:01:38 UTC+0100	CREATE_IN_PROGRESS	AWS::EC2::SubnetRouteTableAssociation	RouteTableAssociationAPublic	Resource creation initiated
01:01:38 UTC+0100	CREATE_IN_PROGRESS	AWS::EC2::SubnetNetworkAcAssociation	SubnetNetworkAcAssociationBPublic	Resource creation initiated
01:01:38 UTC+0100	CREATE_COMPLETE	AWS::RDS::DBSubnetGroup	DBSubnetGroup	
01:01:37 UTC+0100	CREATE_IN_PROGRESS	AWS::ElasticLoadBalancingV2::LoadBalancer	ApplicationLoadBalancer	Resource creation initiated
01:01:37 UTC+0100	CREATE_IN_PROGRESS	AWS::EC2::SubnetRouteTableAssociation	RouteTableAssociationBPublic	Resource creation initiated
01:01:37 UTC+0100	CREATE_IN_PROGRESS	AWS::EC2::SubnetNetworkAcAssociation	SubnetNetworkAcAssociationBPublic	Resource creation initiated
01:01:37 UTC+0100	CREATE_IN_PROGRESS	AWS::AutoScaling::AutoScalingGroup	WebServerGroup	Resource creation initiated
01:01:37 UTC+0100	CREATE_IN_PROGRESS	AWS::RDS::DBSubnetGroup	DBSubnetGroup	Resource creation initiated
01:01:37 UTC+0100	CREATE_IN_PROGRESS	AWS::EC2::SubnetNetworkAcAssociation	SubnetNetworkAcAssociationAPrivate	Resource creation initiated
01:01:37 UTC+0100	CREATE_IN_PROGRESS	AWS::EC2::SubnetRouteTableAssociation	RouteTableAssociationBPrivate	Resource creation initiated
01:01:37 UTC+0100	CREATE_IN_PROGRESS	AWS::ElasticLoadBalancingV2::LoadBalancer	ApplicationLoadBalancer	Resource creation initiated

The following modifications were made after cloudformation stack completion from the *AWS console*:

Security

- Added the inbound rule on web servers from port 3306 for the DBEC2SecurityGroup
- Modified DBInstance to the DBEC2SecurityGroup and removed the default one.
- DBEC2SecurityGroup ingress rule to allow all traffic from web servers. (to get it to work)

S3 storage

- Created an s3 bucket in the current region
- Moved the "bcu.jpg" into the S3 bucket.
- Gave read-only permissions for the public for the "bcu.jpg" image.

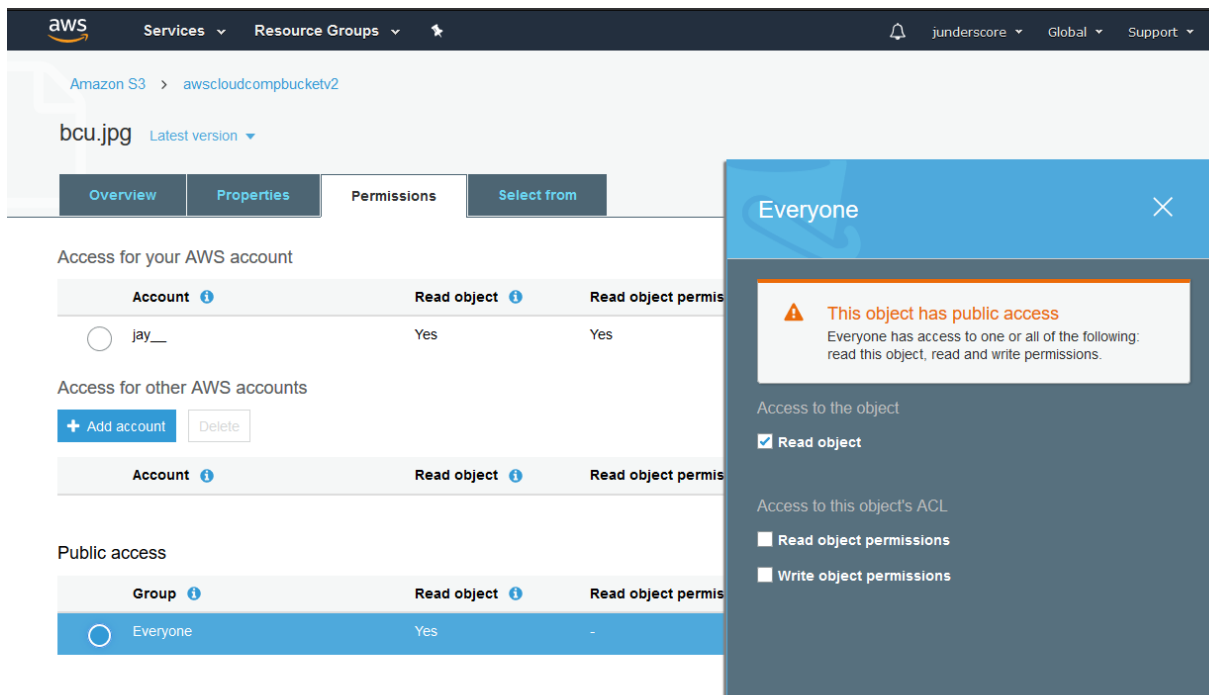


Illustration 1: Screenshot of S3 Storage Interface

MySQL scripting using Workbench

- Created a connection to the database
- Created the tables with the "cake_bookmarks_createTables.sql" file.

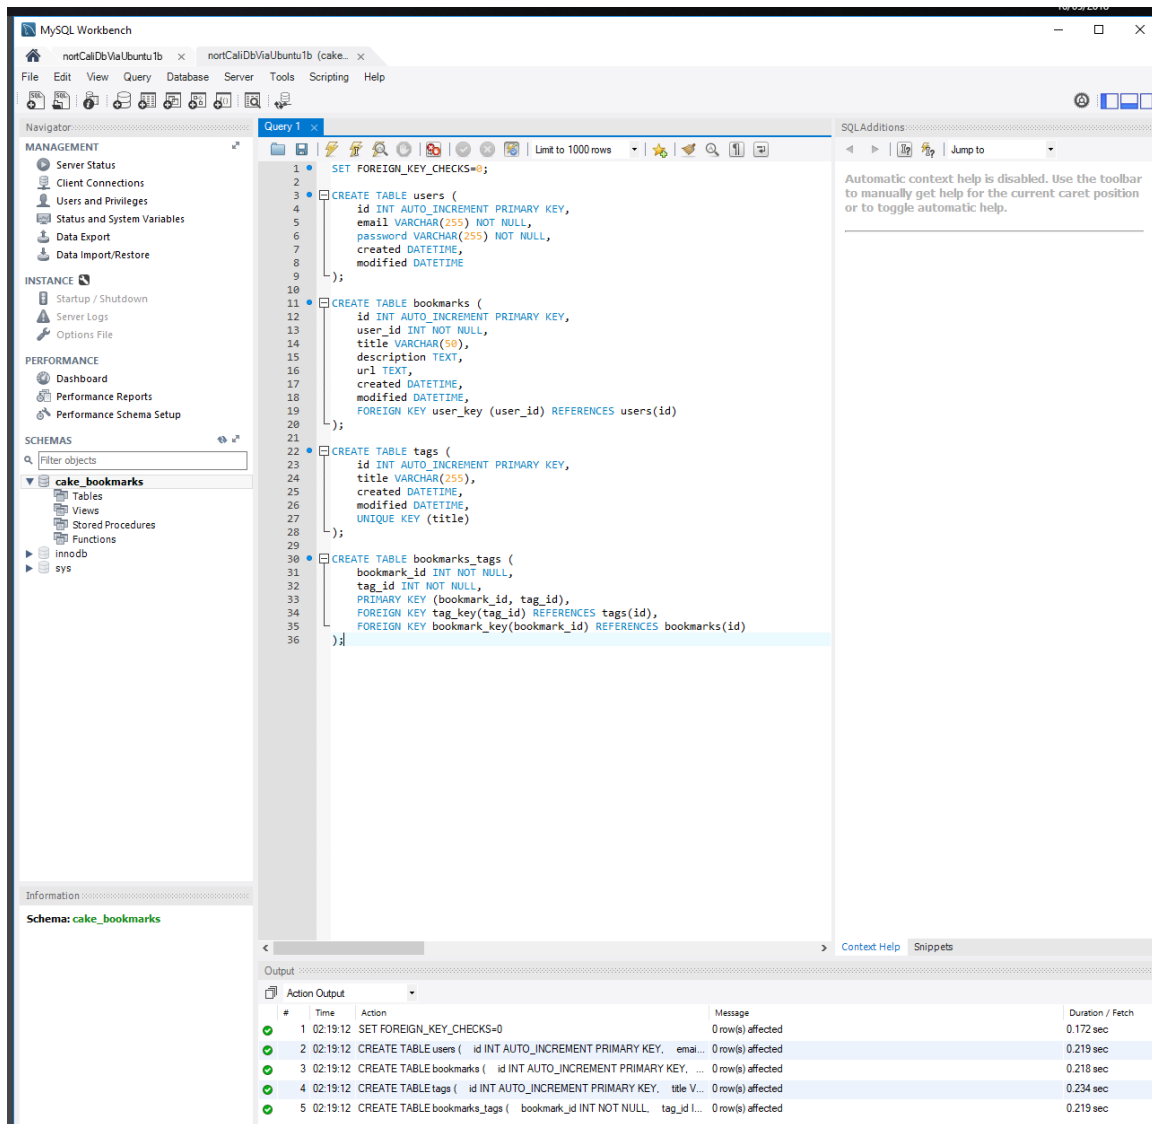


Illustration 2: Screenshot of MySQL Workbench Interface

Configured Webserver using WinSCP and Putty

This procedure was carried out for both ubuntu EC2 instances that were created:

- Connect via ssh client and execute the following

```
sudo su
apt-get update && apt-get upgrade -y
usermod -a -G root ubuntu
chown -R ubuntu:ubuntu /var/www/html
chmod -R 755 /var/www/html
a2enmod rewrite
sudo nano /etc/apache2/sites-available/000-default.conf
```
- Added to webserver configuration file:

```
<VirtualHost *:80>
    <Directory /var/www/html>
        Options Indexes FollowSymLinks MultiViews
        AllowOverride All
        Require all granted
    </Directory>
    . . .
</VirtualHost>
```
- Modified *document root*

```
/var/www/html/bcu_bookmarker
<<Ctrl+o>> then <<Ctrl+x>> to save the changes.
```
- `sudo nano /var/www/html/.htaccess`

```
RewriteEngine on
<<Ctrl+o>> then <<Ctrl+x>> to save the changes.
systemctl restart apache2
```
- Dragged and dropped "bcu_bookmarker" client application into /var/www/html folder. (Long wait)

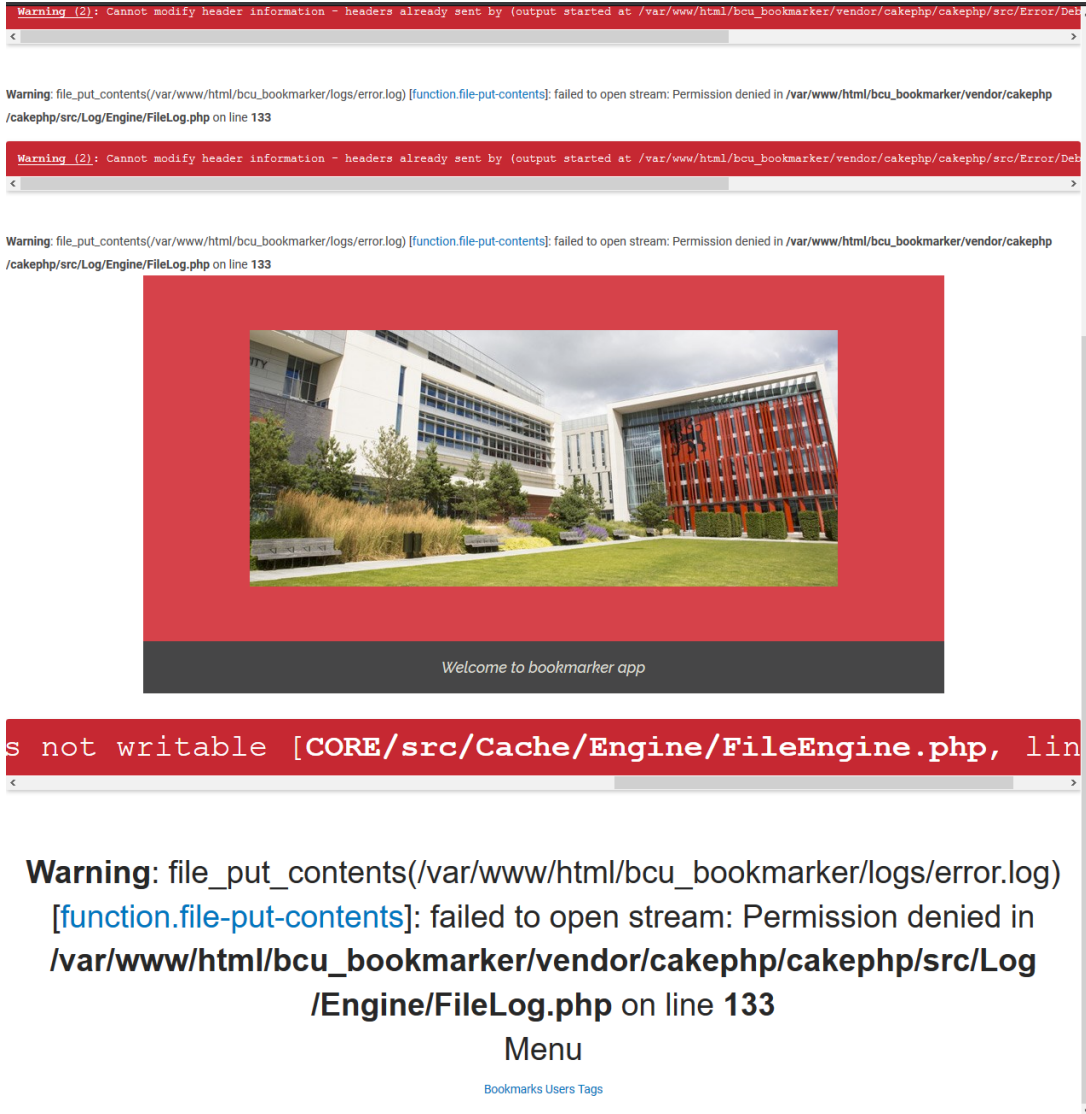


Illustration 3: Screenshot of Web App in error reporting mode

```

Warning (512): /var/www/html/bcu_bookmarker/tmp/cache/ is not writable [CORE/src/Cache/Engine/FileEngine.php, line 437]

Warning: file_put_contents(/var/www/html/bcu_bookmarker/logs/error.log) [function.file-put-contents]: failed to open stream: Permission denied in /var/www/html/bcu_bookmarker/vendor/cakephp/cakephp/src/Log/Engine/FileLog.php on line 133

Warning (512): _cake_routes_cache was unable to write 'routeCollection' to Cake\Cache\Engine\FileEngine cache [CORE/src/Cache/Cache.php, line 290]

Warning: file_put_contents(/var/www/html/bcu_bookmarker/logs/error.log) [function.file-put-contents]: failed to open stream: Permission denied in /var/www/html/bcu_bookmarker/vendor/cakephp/cakephp/src/Log/Engine/FileLog.php on line 133

Warning (2): file_put_contents(/var/www/html/bcu_bookmarker/logs/error.log) [<a href='https://secure.php.net/function.file-put-contents'>function.file-put-contents</a>]: failed to open stream: Permission denied [CORE/src/Log/Engine/FileLog.php, line 133]

Warning: file_put_contents(/var/www/html/bcu_bookmarker/logs/error.log) [function.file-put-contents]: failed to open stream: Permission denied in /var/www/html/bcu_bookmarker/vendor/cakephp/cakephp/src/Log/Engine/FileLog.php on line 133

Warning (512): Unable to emit headers. Headers sent in file=/var/www/html/bcu_bookmarker/vendor/cakephp/cakephp/src/Error/Debugger.php line=853 [CORE/src/Http/ResponseEmitter.php, line 48]

Warning: file_put_contents(/var/www/html/bcu_bookmarker/logs/error.log) [function.file-put-contents]: failed to open stream: Permission denied in /var/www/html/bcu_bookmarker/vendor/cakephp/cakephp/src/Log/Engine/FileLog.php on line 133

Warning (2): Cannot modify header information - headers already sent by (output started at /var/www/html/bcu_bookmarker/vendor/cakephp/cakephp/src/Error/Debugger.php:853) [CORE/src/Http/ResponseEmitter.php, line 148]

Warning: file_put_contents(/var/www/html/bcu_bookmarker/logs/error.log) [function.file-put-contents]: failed to open stream: Permission denied in /var/www/html/bcu_bookmarker/vendor/cakephp/cakephp/src/Log/Engine/FileLog.php on line 133

Warning (2): Cannot modify header information - headers already sent by (output started at /var/www/html/bcu_bookmarker/vendor/cakephp/cakephp/src/Error/Debugger.php:853) [CORE/src/Http/ResponseEmitter.php, line 181]

Warning: file_put_contents(/var/www/html/bcu_bookmarker/logs/error.log) [function.file-put-contents]: failed to open stream: Permission denied in /var/www/html/bcu_bookmarker/vendor/cakephp/cakephp/src/Log/Engine/FileLog.php on line 133

```

Database Error
PDOException

toggle vendor stack frames

```

) PDO->__construct
CORE/src/Database/Driver.php, line 92

) Cake\Database\Driver->connect
CORE/src/Database/Driver/MySQL.php, line 103

) Cake\Database\Driver\MySQL->connect
CORE/src/Database/Schema/BaseSchema.php, line 45

) Cake\Database\Schema\BaseSchema->__construct
CORE/src/Database/Dialect/MySQLDialectTrait.php, line 63

) Cake\Database\Driver\MySQL->schemaDialect
CORE/src/Database/Schema/Collection.php, line 52

) Cake\Database\Schema\Collection->__construct
CORE/src/Database/Schema/CachedCollection.php, line 42

```

Error: SQLSTATE[HY000] [2002] Connection timed out

If you are using SQL keywords as table column names, you can enable identifier quoting for your database connection in config/app.php.

If you want to customize this error message, create src/Template/Error/pdo_error.ctp

Illustration 4: More errors reported during debugging

In app config

- We set the debug variable to false to hide the errors
- Entered the db host as the rds endpoint.
 - `amtzcbah63q21n.ccot9a10nlqj.us-west-1.rds.amazonaws.com`
- Modified security salt variable
- In the home.ctp we linked to s3 bucket in home.ctp img tag src:
 -

<https://s3-us-west-1.amazonaws.com/awsclooudcompbucketv2/bcu.jpg>

This completed the initial deployment of the app.

The next step was to **populate database with test data:**

Users Home

i The user has been saved.

ACTIONS

New User

List Bookmarks

New Bookmark

Users

Id	Email	Password	Created	Modified	Actions
1	testuser1@hotmail.com	genericPass	5/16/18, 3:09 AM	5/16/18, 3:09 AM	View Edit Delete
2	testuser2@gmail.com	asecondpassword	5/16/18, 3:24 AM	5/16/18, 3:24 AM	View Edit Delete
3	jfdvjvfjij@aol.com	hahahahah	5/16/18, 5:23 AM	5/16/18, 5:23 AM	View Edit Delete

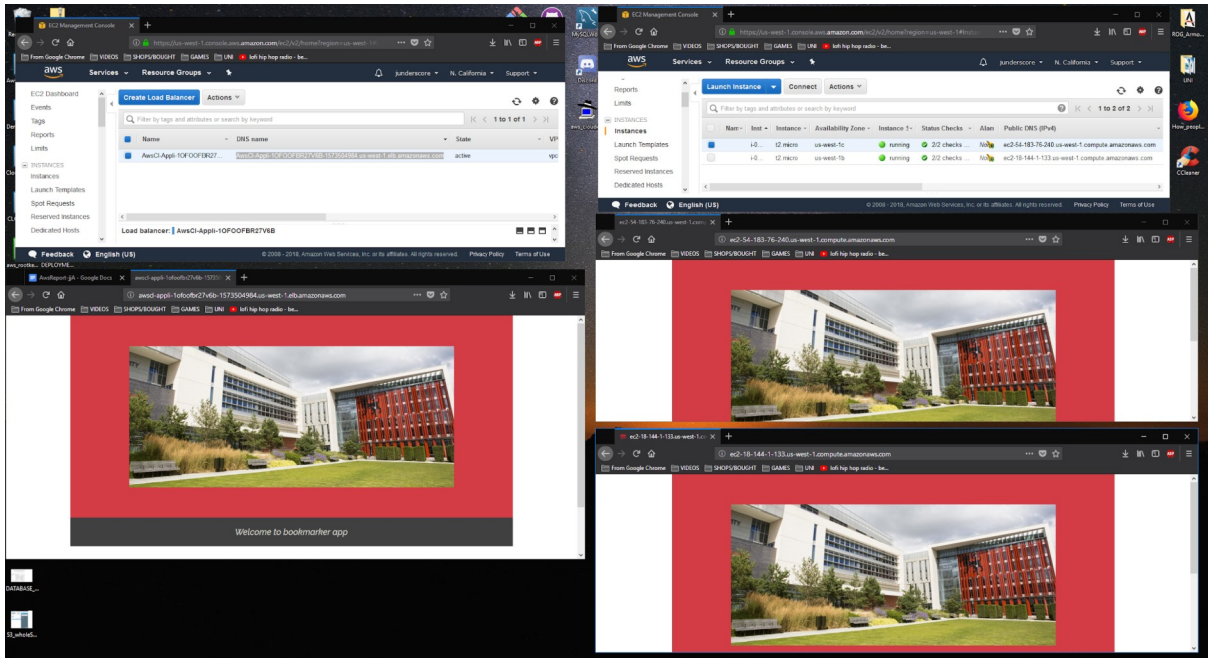
< previous next >

Page 1 of 1, showing 3 record(s) out of 3 total

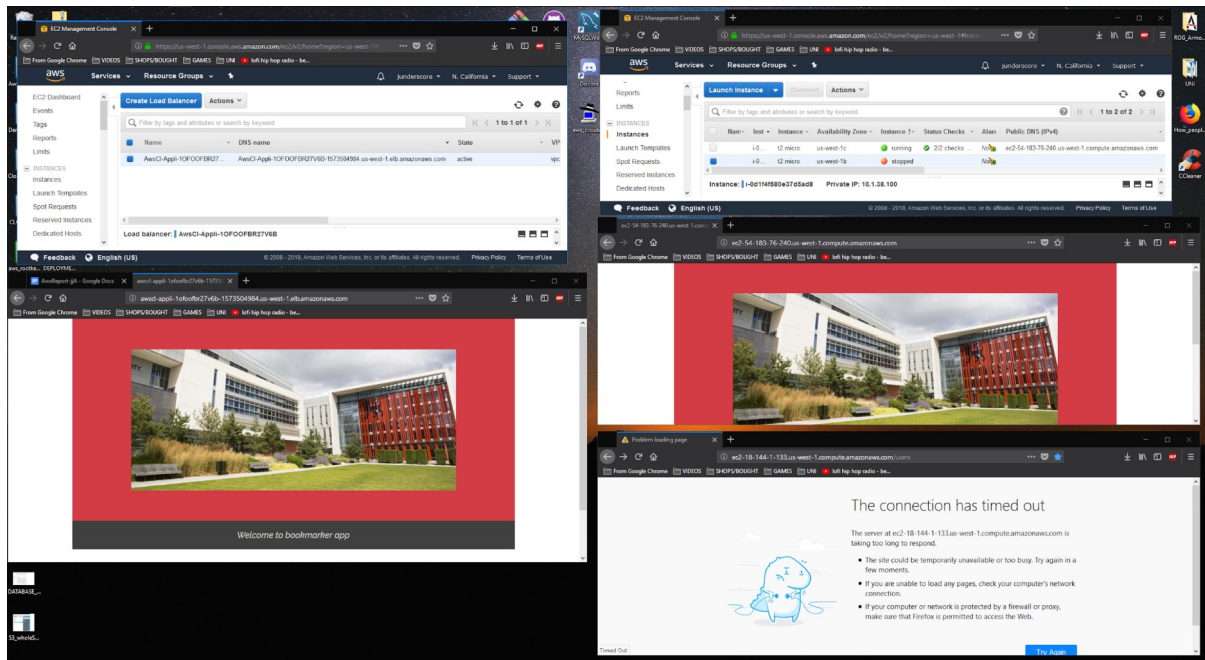
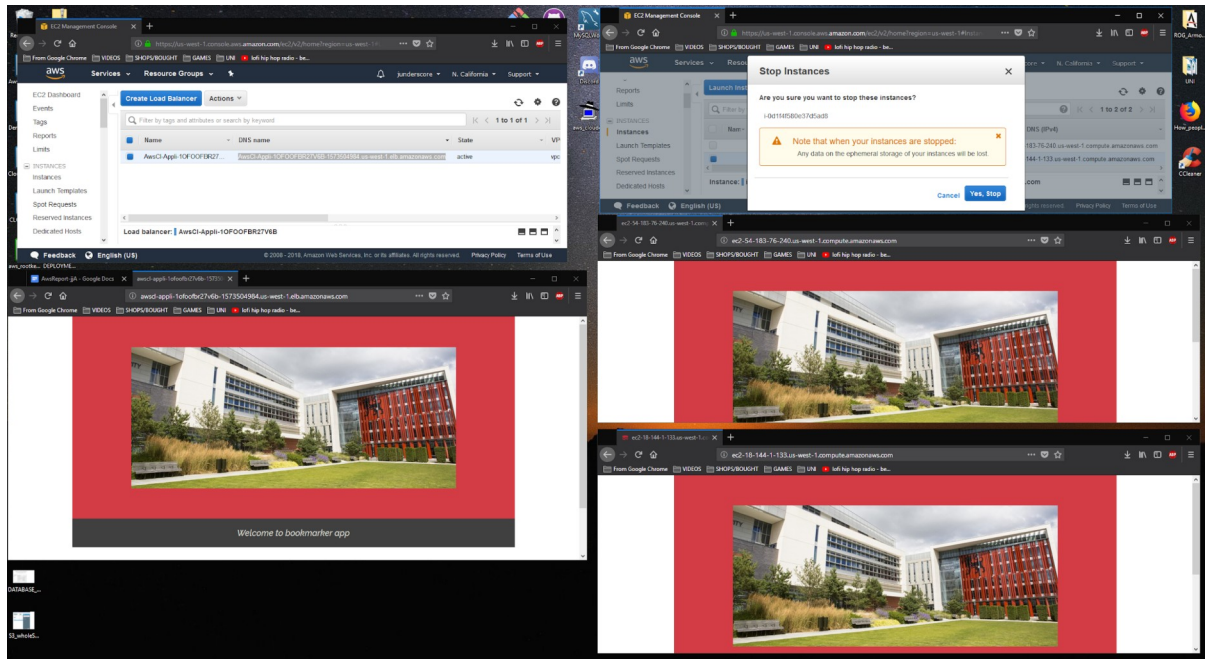
Illustration 5: Screenshot showing test data has been added

STAGE THREE: Load balancing tests (North California)

We have two Ubuntu EC2 instances on the right side of the screen running in N.Cali region, amazon zone 1b and 1c. Both of them can be accessed from the load balancer on the left side of the screen.

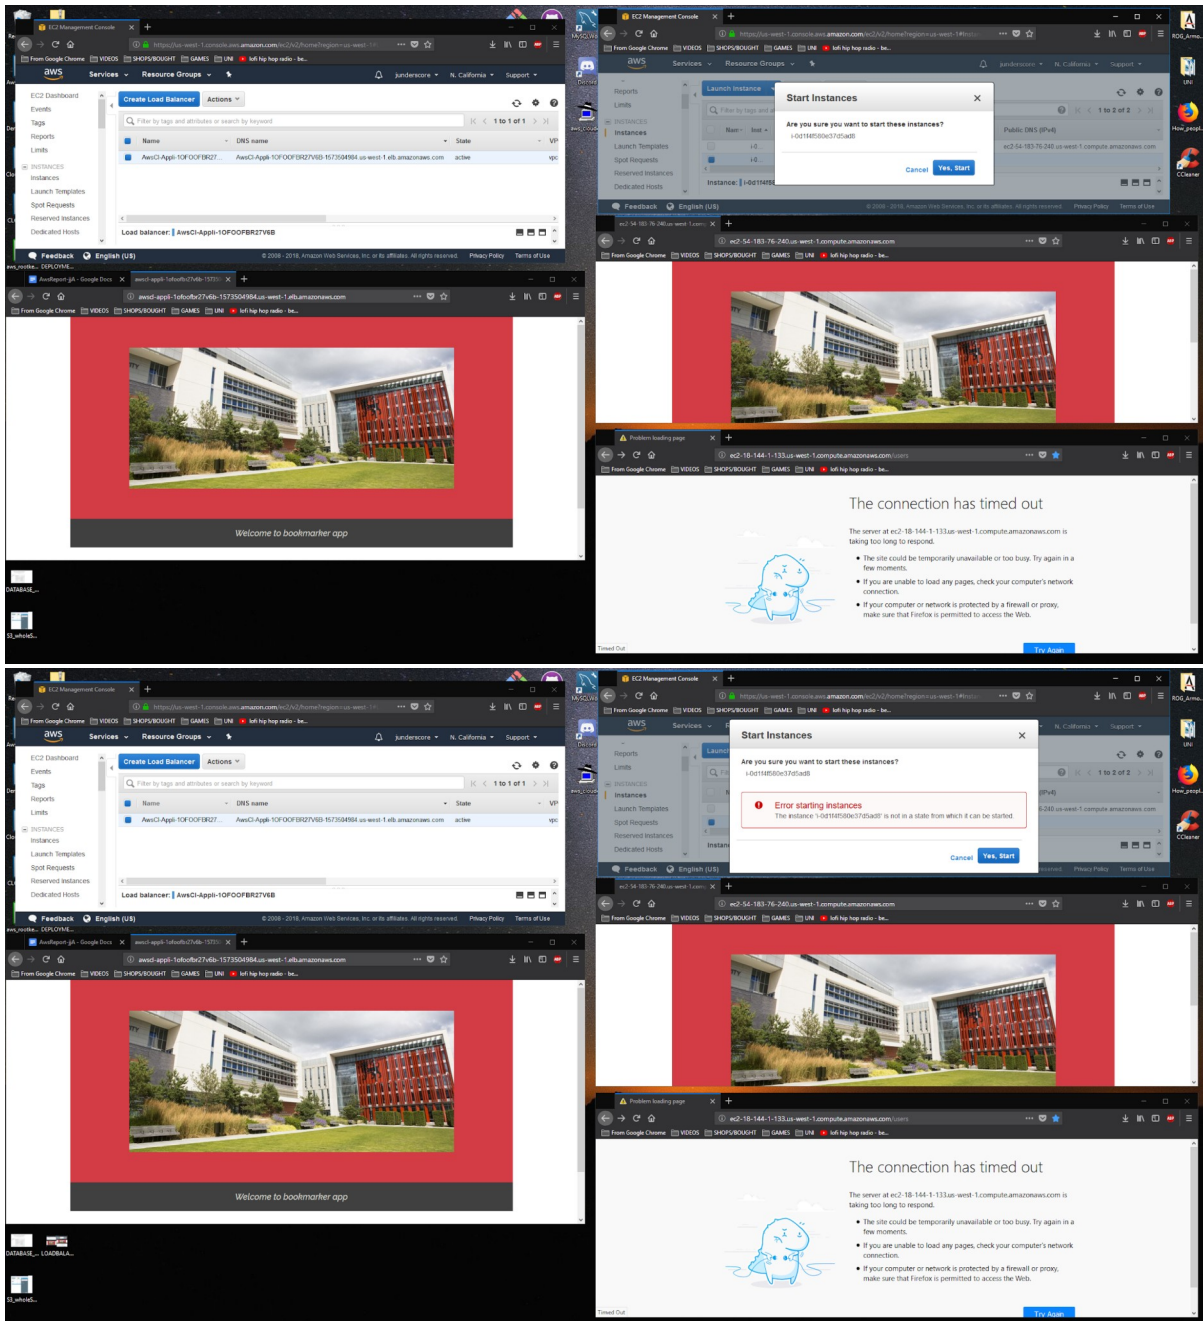


Stopping 1b gives us this:



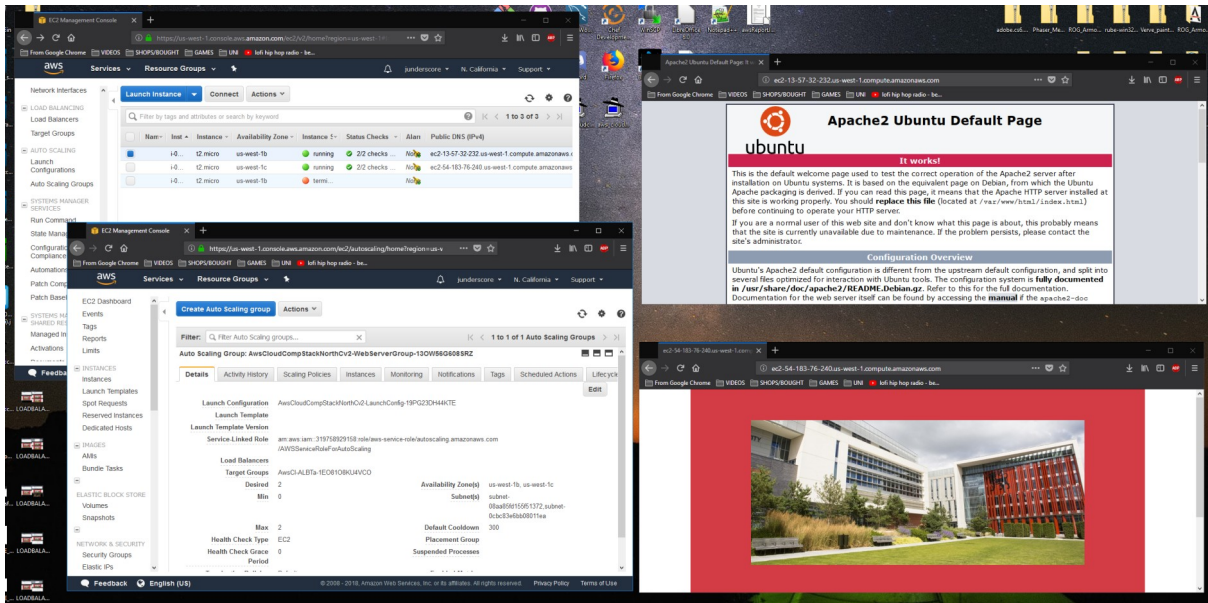
This shows us that the load balancer is still able to access the app through ubuntu EC2 instance in AZ 1c.

Restarting the instance gives this:



(Instances can't be recovered once stopped.)

The additional features we have included like the autoscaling group almost fixes this problem as it “defaults” to two running instances and boots another ubuntu EC2 instance ready for app deployment:



STAGE FOUR: Security practices (Oregon region)

N.B. All the current setups and tests done from the VPC in Oregon could be continued over into the VPC in North California so are still valid.

IAM service

One of the security features that were used in this project is the creation of an IAM service. The IAM service is one of the services within AWS. With IAM it's possible to manage users and their security credentials like their passwords, permissions and access keys, also it controls which AWS service the users account can use.

When a user creates an AWS account, the account has access to the AWS management console, it's possible to create access keys which contain an access ID and a secret access key. This is used to make programmatic calls to AWS using different features like the command line interface (CLI).

The IAM allows the creation of individual users with that AWS account, they each get their own username, access keys and password. Each account is tied to a URL link which allows the user to access their account. All the activities that are done within these account are billed to the main account. There are 4 access levels which a user can have which are list, read, write and permissions management. The list and read allows the user to just view and list different services like for example on S3 services the user with these access levels would only be able to view and list buckets. With the write access the user can delete or add objects in the buckets and with the permission management they have full access to all the services and its features such as manage bucket policies on S3.

As a best practice 3 different user accounts were created within 1 group for each of the team members, including the main accounts holder. In this case this was done on James's AWS account. What was done was James's user account had all the access levels since he's the administrator, whereas Ahmed's and cheng's account has the list, read and write access which limited.

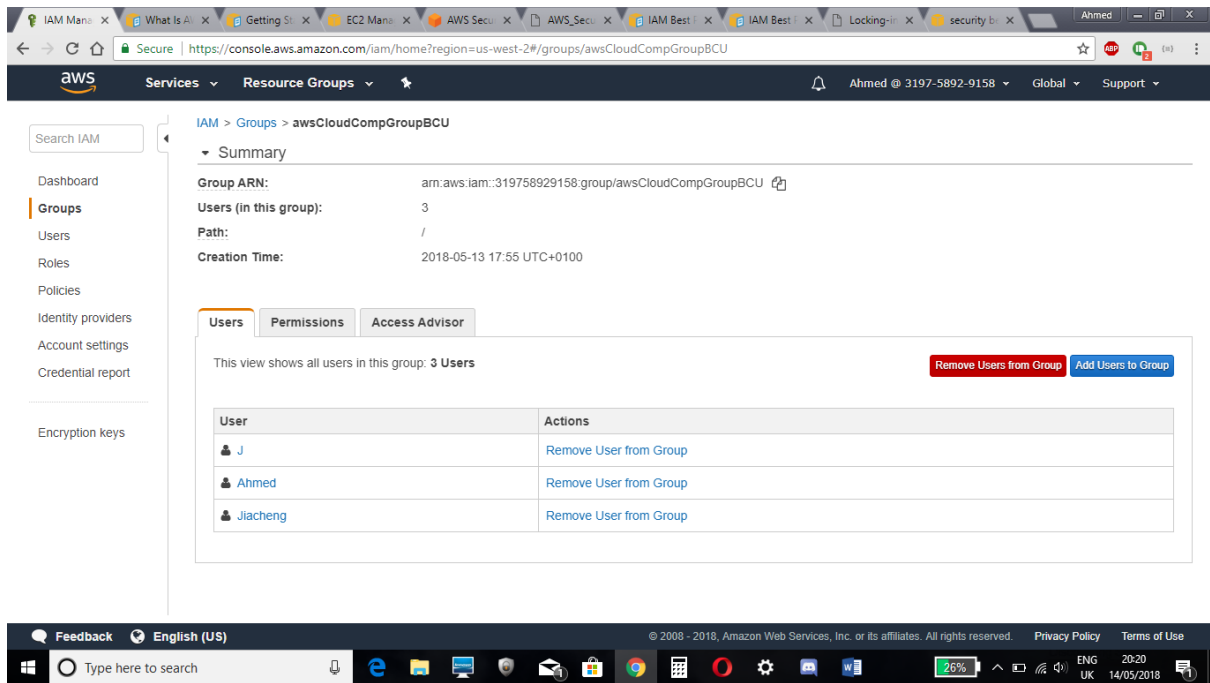


Figure 1: the IAM group with all the users contained.

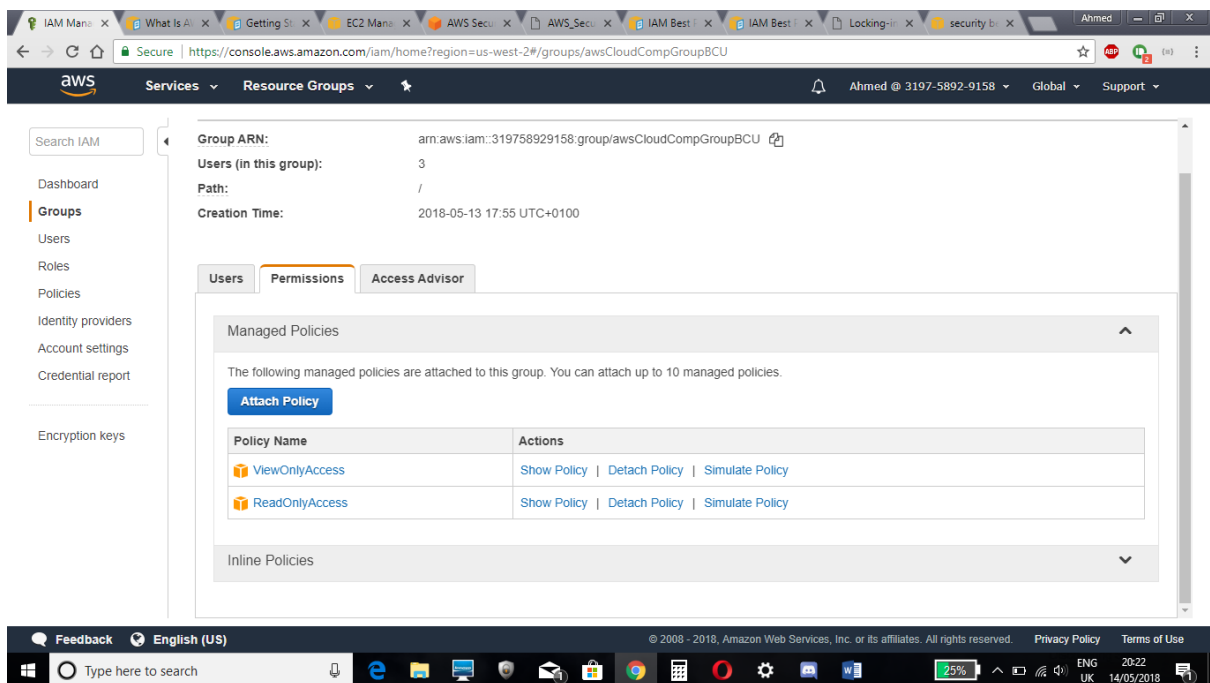


Figure 2: permissions included in the group

Another feature with the IAM service was an IAM password policy allowing the user to change their password. What was done was when one user wants to change their password they would need to create a strong password. The screenshot below shows what policies was added.

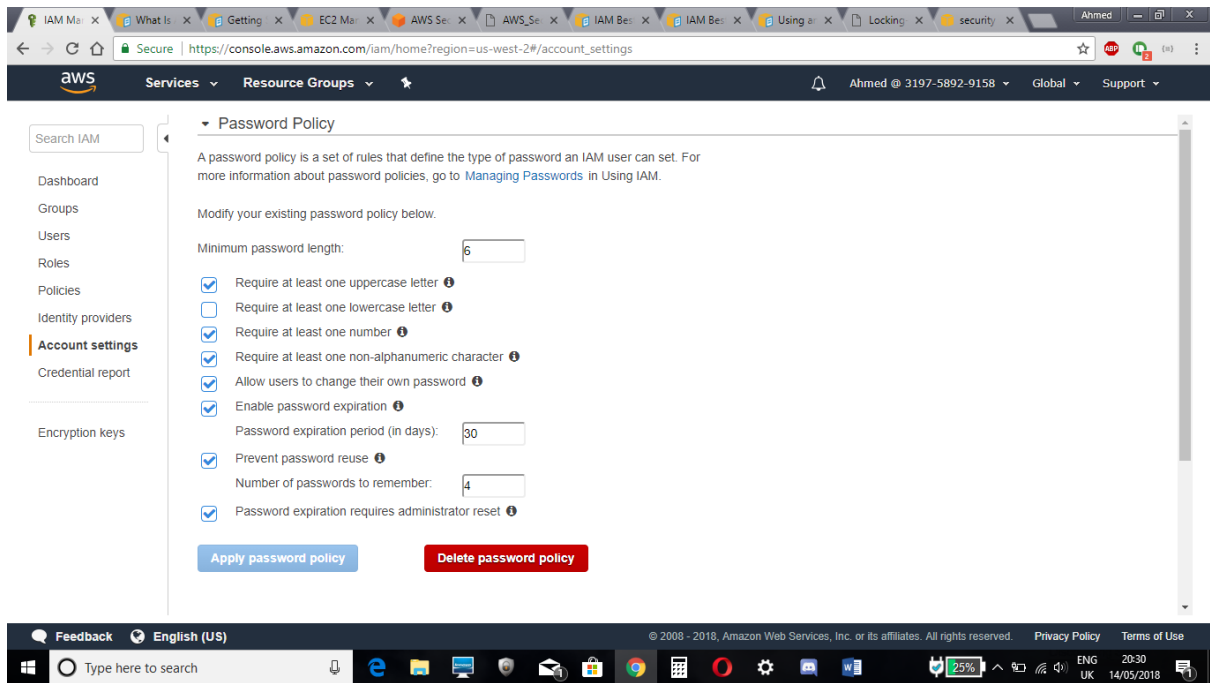


Figure 3: password policy.

IAM roles

When an EC2 instance is running applications that are run on it need certain IDs to make use of the AWS services they need. A role is basically a unit with its own set of permissions. Using IAM roles allows the EC2 instance to temporarily gain these credentials to allow them to use the services they need. An example can be application code running on an EC2 instance that needs to perform actions on AWS resources. IAM roles distribute keys that are valid for a short time, thus ensuring they are a more secure way to grant access. Below is a screen shot of the roles that are included in the EC2 instances within the account.

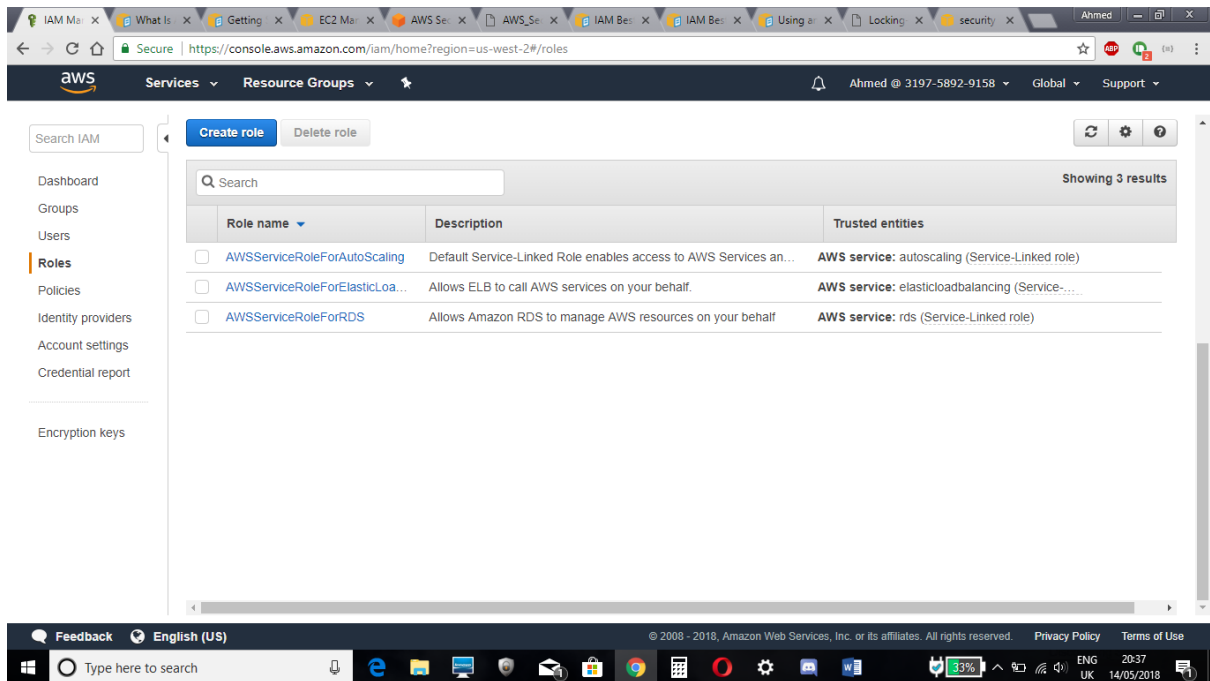


Figure 4: IAM Roles

Network ACLs

Network access control list (ACL) is an optional layer of security for the VPC which essentially acts as a firewall for controlling traffic in and out of one or more subnets. As a best practice the network ACLs that were created have the same outbound rules as the security groups that were created. The security groups have more specific inbound rules. The resemblance in the rules add extra layer of security to the VPC. Each subnet is related to a network ACL. A network ACL has a numbered list of rules are valued in order from lowest to highest, to control if traffic is allowed in or out of any subnet linked with the network ACL.

The highest number that can be used for a rule is 32766.

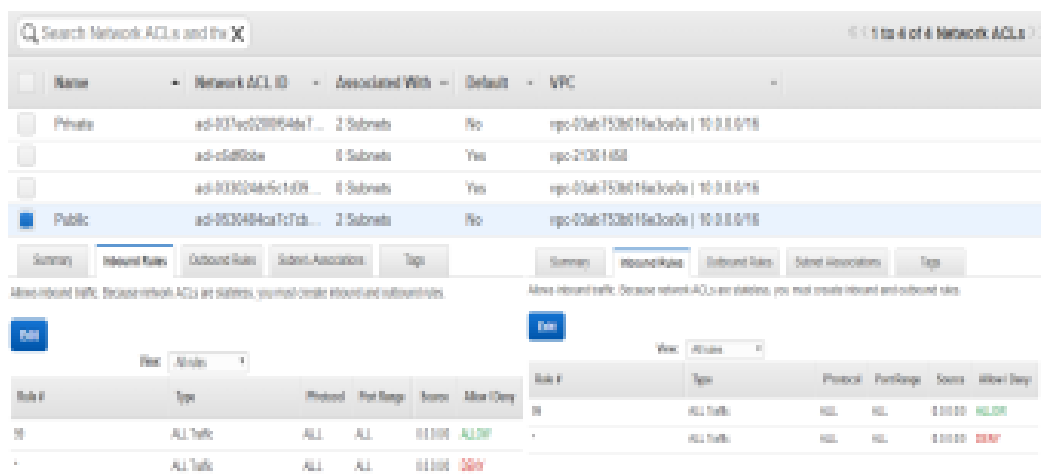
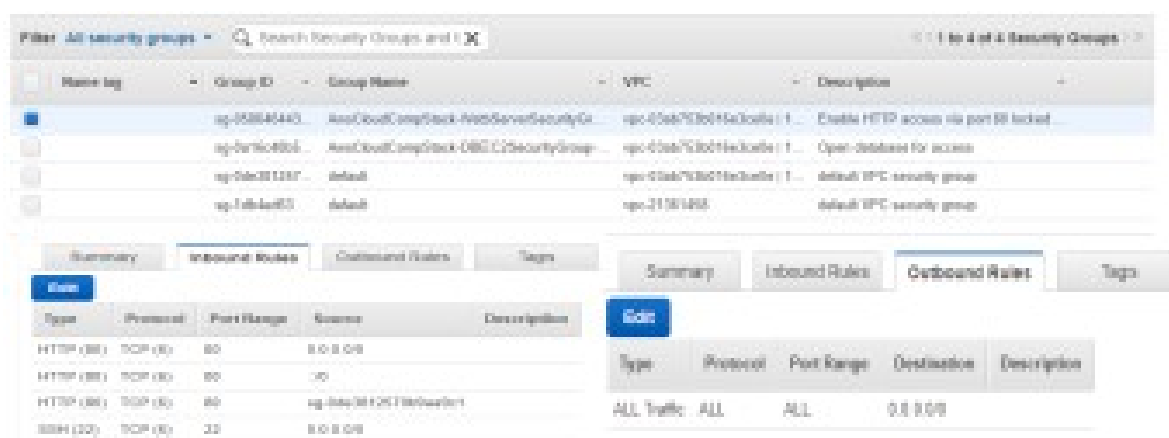


Figure 5: ACLs with inbound and outbound rules

Security groups

A security groups actions is being a simulated firewall that controls the traffic for one or more instances. When an instance is launched, it is good practice to add one or more security groups to that instance. adding rules to each security group allows traffic to or from its related instances. The rules of a security group control the inbound traffic that's allowed to get to the instances which link to the security group and the outbound traffic that's allowed to leave them. As a best security practice these security groups implement the least permissive rules.



The screenshot shows the AWS IAM console interface for security groups. At the top, there is a search bar and a filter for 'All security groups'. Below this is a table listing security groups with columns for Name tag, Group ID, Group Name, VPC, and Description. The table contains four entries: 'sg-05864440' (AmazonCloudWatch-AmazonServerSecurityGr...), 'sg-0a1c4906' (AmazonCloudWatch-DBEC2SecurityGroup...), 'sg-0a2c31287' (default), and 'sg-14b1a251' (default). Below the table, there are two tabs: 'Inbound Rules' and 'Outbound Rules'. The 'Inbound Rules' tab is active, showing a table with columns for Type, Protocol, Port Range, Source, and Description. It lists four rules: 'HTTP (80)' (TCP (8)) to port 80 from '0.0.0.0/0', 'HTTP (80)' (TCP (8)) to port 80 from '0/0', 'HTTP (80)' (TCP (8)) to port 80 from 'sg-0a2c31287/0.0.0.0/0', and 'SSH (22)' (TCP (8)) to port 22 from '0.0.0.0/0'. The 'Outbound Rules' tab is also visible, showing a single rule for 'ALL Traffic' (ALL) to 'ALL' destinations from '0.0.0.0/0'.

Figure 6: security groups with inbound and outbound rules

Additional options

There are other security options that can be added, one such security element is a multi-factor authentication. For extra security, multi-factor authentication can be enabled for the users who have all the access levels as they are allowed to access delicate resources. The user can obtain a device that would generate a onetime unique passcode, they would then need to sign in with both their user account password and this one time passcode.

Demonstrates high availability at different levels

Instances:

There were 2 instances created and with high availability in mind, these instances have different availability zones. By having them in different availability zones it improves the fault tolerance in the application. If on availability zone is not working because of an outage, all traffic is routed to the other availability zone.

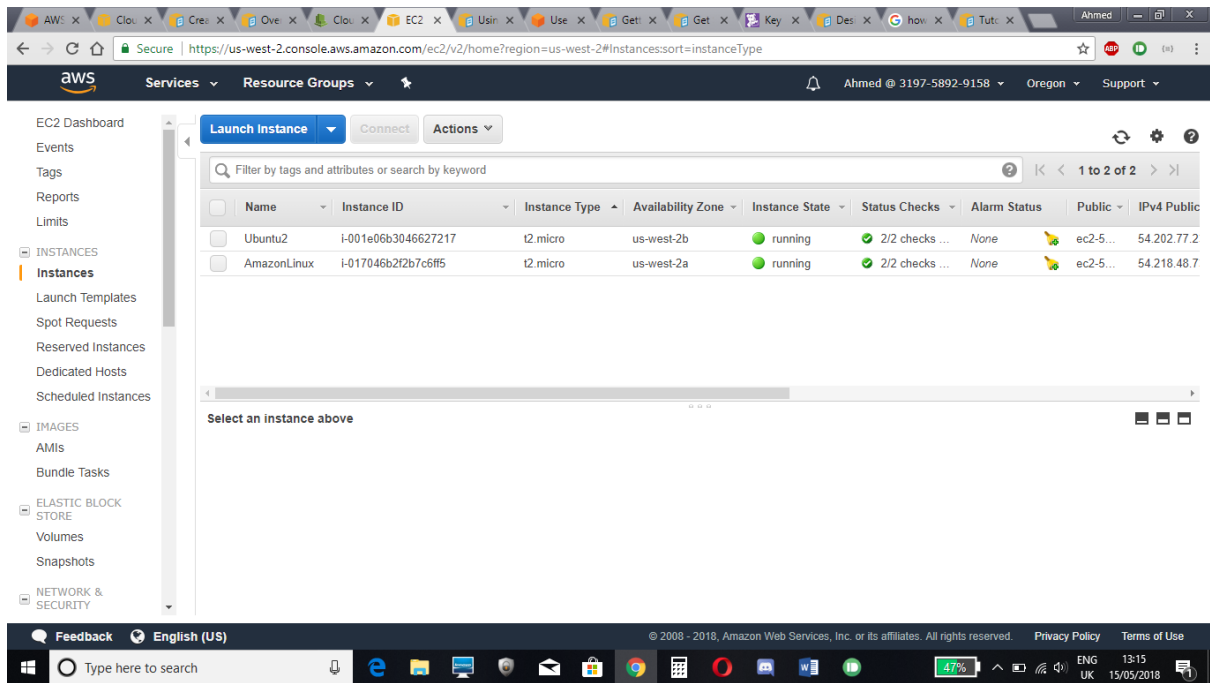


Figure 1: instances with different availability zones.

Auto scaling groups - transferred to North California.

This contains the EC2 instances that have a link between them, like in our example where the 2 instances that we created, one for the web application and another for the PHP application.

What the auto scaling group does it launches the necessary number of instances to meet the requirements we need which is for it to run a web application that can edit a RDS database. For our case one was enough. The auto scaling group keeps this one instance running and keeps tabs on that instance like doing health check, if that specific instance fails the auto scaling group terminates that instance in the group and launches the other instance that is available so that it switches it out.

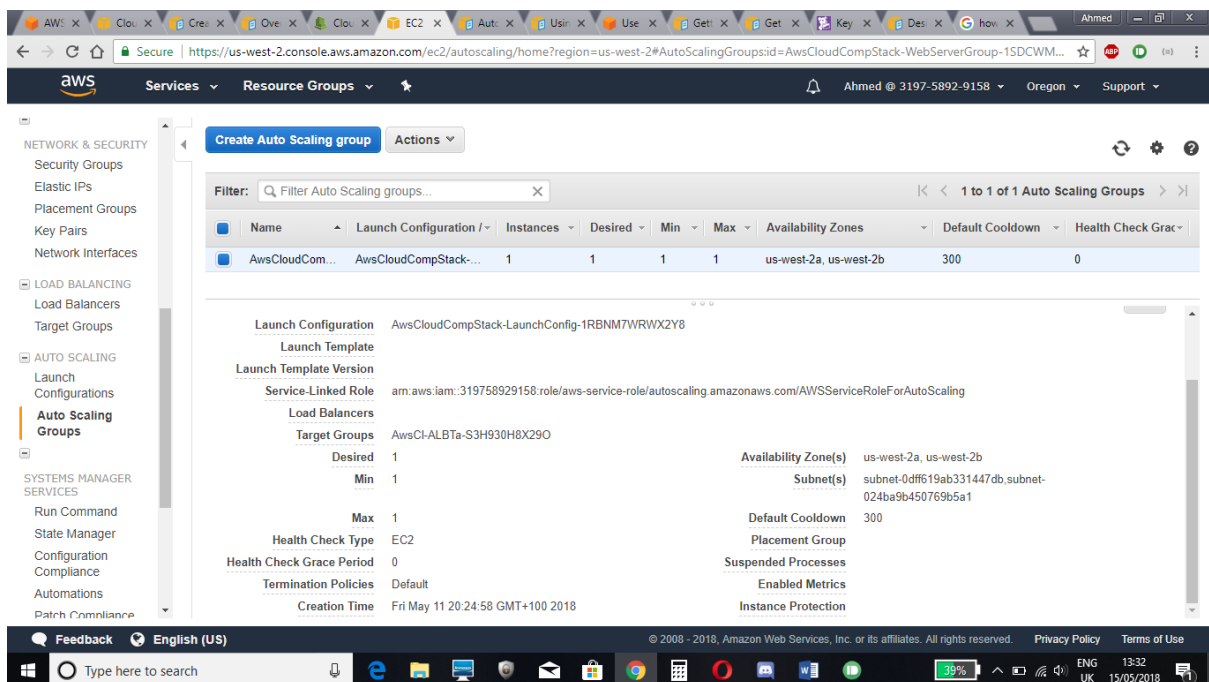


Figure 3: auto scaling for 2 different instances.

STAGE FIVE: Monitoring Instances Using CloudWatch

With CloudWatch it's possible to monitor the instances within the AWS account. The data that is collected with is in raw form is processed from EC2 instance and converted into readable, real-time metrics. These data help in having advanced perception on how the web application or the different services that are running within the instance is performing. By default, Amazon EC2 sends metric data to CloudWatch in 5-minute periods. Our group configured this option to send data in 1 minute periods. This allows more detailed monitoring on the instances created.

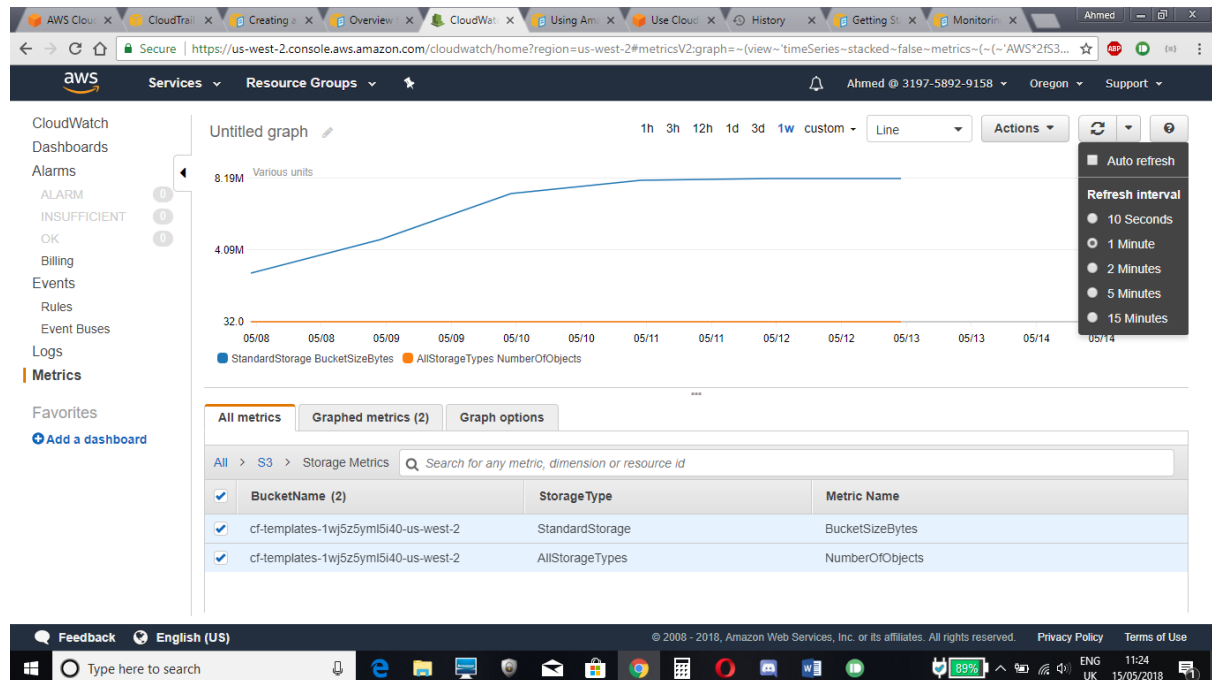


Figure 1: metric data refresh intervals

When monitoring the instances it's important to monitor the basic system-level metrics which are the core setup of the instances to make user they are healthy. Other things to consider tracking are the resources being used. For our project we covered 3 metrics types.

CPU

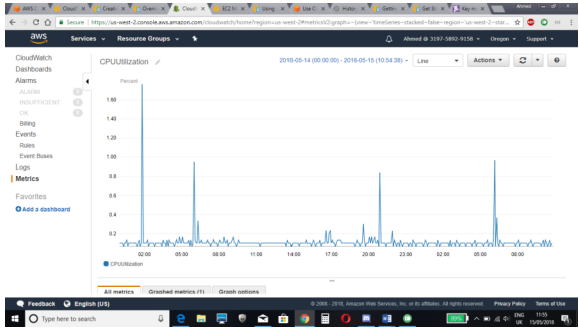

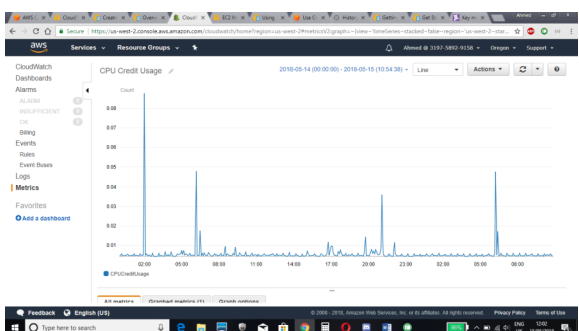
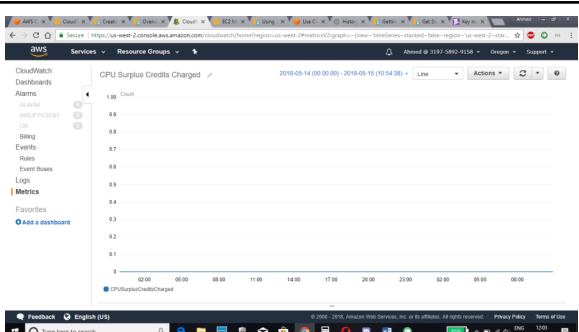
There are several CPU configurations within ECS instances, tracking the CPU usage allows us to know that the instance is appropriately sized for the job that it's configured for.

5 different metrics fall under this category which are:

1. CPU Utilization
2. CPU Credit Balance
3. CPU Credit Usage
4. CPU Surplus Credits Charged
5. CPU Surplus Credit Balance

The table below shows the graphs for these metrics. With description on what these metrics are and screenshots of the graphs corresponding with them. There will a table for each instance. The graphs show the usage between 2018/05/14 at 00:00:00 and 2018/05/15 at 10:54:38

Table 1: Ubuntu 2

Test	Screenshot of graph
<p>CPUUtilization: Percentage of allocated EC2 compute units that are currently in use on the instance.</p>	 <p>The screenshot shows the AWS CloudWatch console for a resource group. The main graph is titled 'CPU Utilization' and displays a line chart of CPU usage percentage over a 24-hour period from 02:00 to 06:00. The y-axis ranges from 0.0 to 1.8. The chart shows a baseline near 0% with several distinct spikes reaching approximately 1.0. The interface includes a left-hand navigation menu with options like Dashboards, Alarms, Billing, and Metrics.</p>
<p>CPUCreditBalance: Number of CPU credits that an instance has accumulated.</p>	 <p>The screenshot shows the AWS CloudWatch console for a resource group. The main graph is titled 'CPU Credit Balance' and displays a line chart of the credit balance over a 24-hour period from 02:00 to 06:00. The y-axis ranges from 164 to 168. The chart shows a steady, linear decrease in the credit balance from approximately 168 to 164. The interface includes a left-hand navigation menu with options like Dashboards, Alarms, Billing, and Metrics.</p>
<p>CPUCreditUsage Number of CPU credits consumed.</p>	 <p>The screenshot shows the AWS CloudWatch console for a resource group. The main graph is titled 'CPU Credit Usage' and displays a line chart of credit usage over a 24-hour period from 02:00 to 06:00. The y-axis ranges from 0.00 to 3.00. The chart shows a baseline near 0.00 with several distinct spikes reaching approximately 0.50. The interface includes a left-hand navigation menu with options like Dashboards, Alarms, Billing, and Metrics.</p>
<p>CPU Surplus Credits Charged Number of surplus credits not offset by earned CPU credits and that will incur charges</p>	 <p>The screenshot shows the AWS CloudWatch console for a resource group. The main graph is titled 'CPU Surplus Credits Charged' and displays a line chart of surplus credits charged over a 24-hour period from 02:00 to 06:00. The y-axis ranges from 0 to 1.00. The chart shows a flat line at 0.00 throughout the entire period. The interface includes a left-hand navigation menu with options like Dashboards, Alarms, Billing, and Metrics.</p>

CPU Surplus Credit Balance
 Number of credits consumed after CPU credit balance has reached 0.




Table 2: Amazon Linux

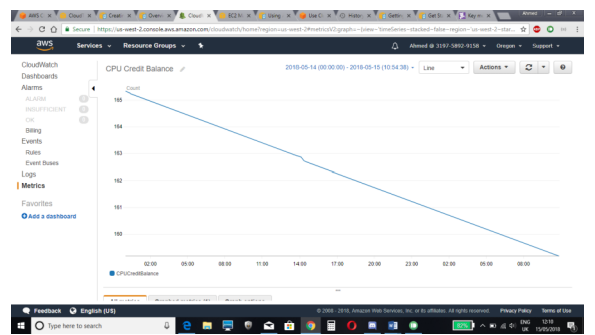
Test

Screenshot of graph

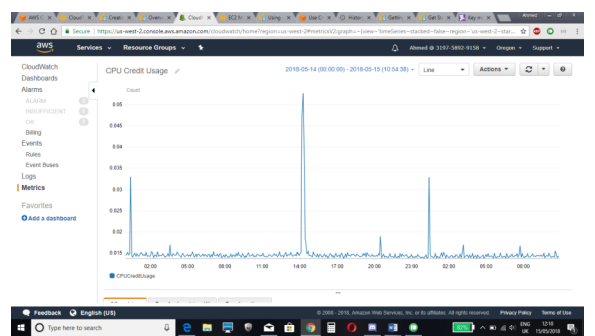
CPU Utilization:
 Percentage of allocated EC2 compute units that are currently in use on the instance.



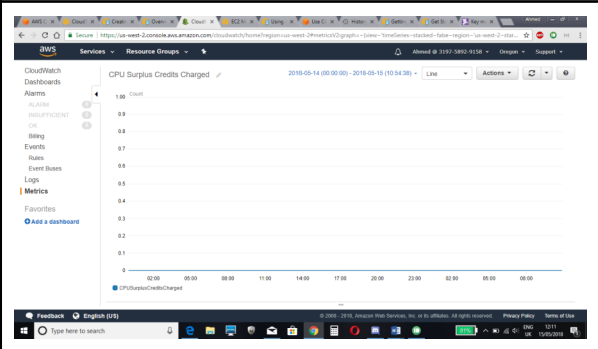
CPU Credit Balance:
 Number of CPU credits that an instance has accumulated.



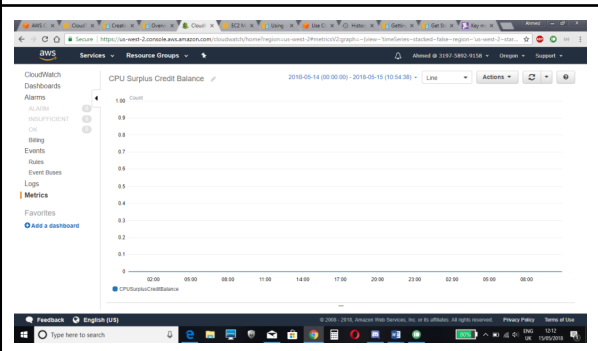
CPU Credit Usage
 Number of CPU credits consumed.



CPU Surplus Credits Charged
 Number of surplus credits not offset by earned CPU credits and that will incur charges



CPU Surplus Credit Balance
 Number of credits consumed after CPU credit balance has reached 0.



Network

The main reason behind the importance of network metrics are their importance for EC2 since it depends on solid network connections, which might be spread across various availability zones.

4 different metrics fall under this category which are:

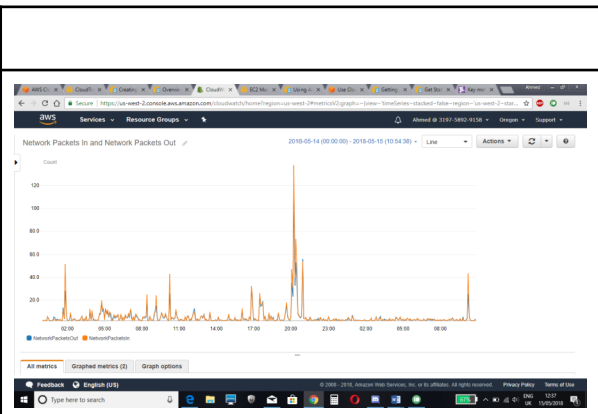
1. Network packs in
2. Network packs out
3. Network in
4. Network out

The table below shows the graphs for these metrics. With description on what these metrics are and screenshots of the graphs corresponding with them. There will be a table for each instance. The graphs show the usage between 2018/05/14 at 00:00:00 and 2018/05/15 at 10:54:38

Table 1: Ubuntu 2

Test

NetworkPacketsIn/NetworkPacketsOut:
 Number of packets received/sent out on all network interfaces by the instance. Only available at five-minute resolution.



<p>NetworkIn NetworkOut Number of bytes received/sent out on all network interfaces by the instance.</p>	
------------------------------------------------------------------------------------------------------------------	--

Table 2: Amazon Linux

<p>Test</p>	
<p>NetworkPacketsIn/NetworkPacketsOut: Number of packets received/sent out on all network interfaces by the instance. Only available at five-minute resolution.</p>	
<p>NetworkIn NetworkOut Number of bytes received/sent out on all network interfaces by the instance.</p>	

Status checks

These checks are simply what is used to check the status of the instance and the systems within it. 2 different metrics fall under this category which are:

1. Status Check Failed_System
2. Status Check Failed_Instance

The table below shows the graphs for these metrics. With description on what these metrics are and screenshots of the graphs corresponding with them. There will a table for each instance. The graphs show the usage between 2018/05/14 at 00:00:00 and 2018/05/15 at 10:54:38

Table 1: Ubuntu 2

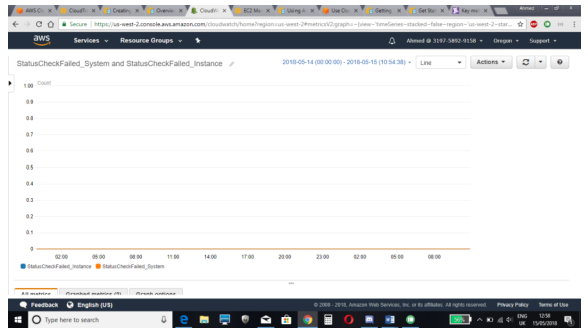
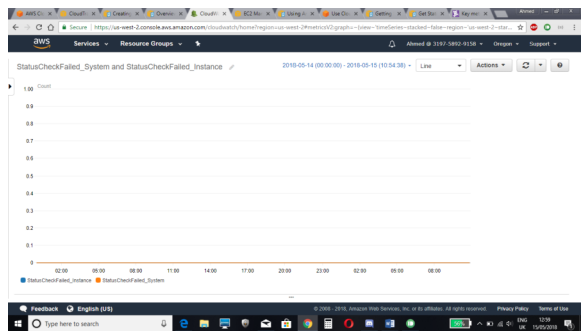
<p>StatusCheckFailed_System StatusCheckFailed_Instance Returns 1 if the instance has failed EC2's system/instance status check.</p>	

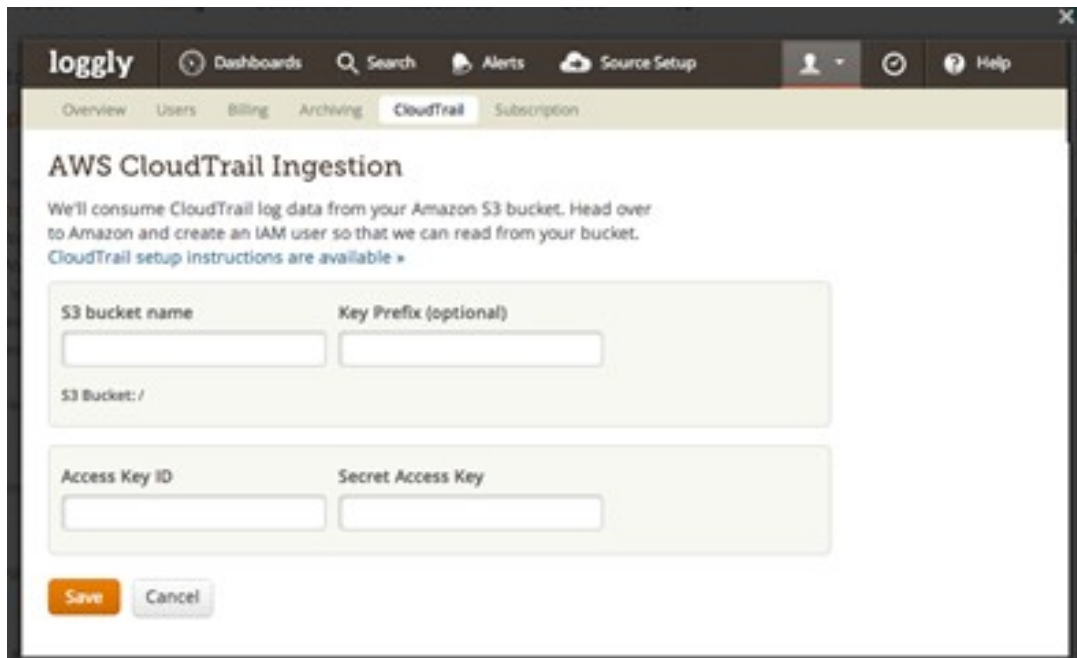
Table 2: Amazon Linux

<p>StatusCheckFailed_System StatusCheckFailed_Instance Returns 1 if the instance has failed EC2's system/instance status check.</p>	

Cloud trail


This is basically an API (application program interface) which creates log files and sends it to the S3 bucket for storage. They provide an audit trail of changes to the AWS instances and the interactions that happen within them.

We have used a third party website called Loggly to create graphs that represent different findings from these log files. What we did was we gave the site the S3 bucket name and an IAM user's access key ID and secret access key. With that user having the right permissions which are list buckets and get objects we managed to get all the S3 data into the site.



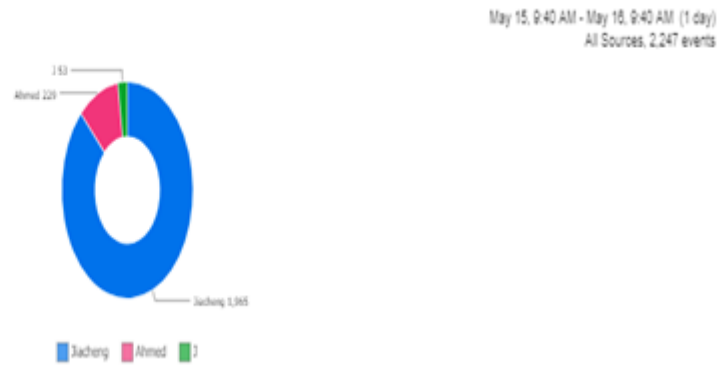
The graphs below show different charts and graphs that we were able to pull from the CloudTrail

Top AWS CloudTrail Event Sources in the Last Day.




This chart shows which of the AWS services were modified by the admin. This can be used to check for unexpected activity

AWS CloudTrail Top Users in the Last Day



This chart shows the activities the users within the AWS account. Unknown usernames can show a security breach

AWS CloudTrail Top Error Messages in the Last Day



This chart shows all the error messages that have come up from all sources.

Cost breakdown for our current setup:

If one person use data 50MB every month, 100 user will use 5GB every month.

Amazon Web Services - x
 https://calculator.s3.amazonaws.com/index.html?key=NONE

Language: English

Need Help? [Watch the Videos](#) or [Read How AWS Pricing Works](#) or [Contact Us](#)

Get Started with AWS: [Learn more about our Free Tier](#) or [Sign Up for an AWS Account >](#)

FREE USAGE TIER: New Customers get free usage tier for first 12 months

Reset All

Services **Estimate of your Monthly Bill (\$ 32.39)**

Estimate of Your Monthly Bill
 Show First Month's Bill (include all one-time fees, if any)

Below you will see an estimate of your monthly bill. Expand each line item to see cost breakout of each service. To save this bill and input values, click on 'Save and Share' button. To remove the service from the estimate, jump back to the service and clear the specific service's form.

Export to CSV Save and Share

Service	Item	Cost	Total
Amazon EC2	Amazon EC2 Service (US East (N. Virginia))		\$ 18.60
	Compute:	\$ 17.00	
	EBS Volumes:	\$ 1.60	
Amazon S3	Amazon S3 Service (US East (N. Virginia))		\$ 0.03
	Standard Storage:	\$ 0.03	
Amazon RDS	Amazon RDS Service (US East (N. Virginia))		\$ 12.95
	DB instances:	\$ 12.45	
	Storage:	\$ 0.50	
Amazon SNS	AWS Data Transfer In		\$ 0.00
	US East (N. Virginia) Region:	\$ 0.00	
Amazon Elastic Transcoder	AWS Data Transfer Out		\$ 0.81
	US East (N. Virginia) Region:	\$ 0.81	
AWS Support (Basic)	AWS Support (Basic)		\$ 0.00
	Support for all AWS services:	\$ 0.00	
Total Monthly Payment:			\$ 32.39

Common Customer Samples

- Free Website on AWS
- AWS Elastic Beanstalk Default
- Marketing Web Site
- Large Web Application (All On-Demand)
- Media Application
- European Web Application
- Disaster Recovery and Backup

5GB data translate will need \$32.39 every month.
 10 thousand user will be 500GB/month

Amazon Web Services - x
 Google 翻译 x
 https://calculator.s3.amazonaws.com/index.html?key=NONE

Language: English

Need Help? [Watch the Videos](#) or [Read How AWS Pricing Works](#) or [Contact Us](#)

Get Started with AWS: [Learn more about our Free Tier](#) or [Sign Up for an AWS Account >](#)

FREE USAGE TIER: New Customers get free usage tier for first 12 months

Reset All

Services **Estimate of your Monthly Bill (\$ 133.09)**

Estimate of Your Monthly Bill
 Show First Month's Bill (include all one-time fees, if any)

Below you will see an estimate of your monthly bill. Expand each line item to see cost breakout of each service. To save this bill and input values, click on 'Save and Share' button. To remove the service from the estimate, jump back to the service and clear the specific service's form.

Export to CSV Save and Share

Service	Item	Cost	Total
Amazon EC2	Amazon EC2 Service (US East (N. Virginia))		\$ 18.60
	Compute:	\$ 17.00	
	EBS Volumes:	\$ 1.60	
Amazon S3	Amazon S3 Service (US East (N. Virginia))		\$ 11.63
	Standard Storage:	\$ 11.50	
	Standard - IA Storage:	\$ 0.13	
Amazon RDS	Amazon RDS Service (US East (N. Virginia))		\$ 12.95
	DB instances:	\$ 12.45	
	Storage:	\$ 0.50	
Amazon SNS	AWS Data Transfer In		\$ 0.00
	US East (N. Virginia) Region:	\$ 0.00	
Amazon Elastic Transcoder	AWS Data Transfer Out		\$ 89.91
	US East (N. Virginia) Region:	\$ 89.91	
AWS Support (Basic)	AWS Support (Basic)		\$ 0.00
	Support for all AWS services:	\$ 0.00	
Total Monthly Payment:			\$ 133.09

Common Customer Samples

- Free Website on AWS
- AWS Elastic Beanstalk Default
- Marketing Web Site
- Large Web Application (All On-Demand)
- Media Application
- European Web Application
- Disaster Recovery and Backup

About 500GB/month, it will be use \$133.09

At last, if one million user, the data transfer is 50 TB per month

The screenshot shows the AWS Simple Monthly Calculator interface. The main heading is "Estimate of your Monthly Bill (\$ 10256.69)". Below this, there is a table of services and their associated costs. The table is as follows:

Service	Cost
Amazon EC2 Service (US East (N. Virginia))	\$ 18.60
Compute:	\$ 17.00
EBS Volumes:	\$ 1.60
EBS IOPS:	\$ 0.00
Amazon S3 Service (US East (N. Virginia))	\$ 1305.60
Standard Storage:	\$ 1177.60
Standard - IA Storage:	\$ 128.00
Amazon RDS Service (US East (N. Virginia))	\$ 12.95
DB instances:	\$ 12.45
Storage:	\$ 0.50
AWS Data Transfer In	\$ 0.00
US East (N. Virginia) Region:	\$ 0.00
AWS Data Transfer Out	\$ 7987.11
US East (N. Virginia) Region:	\$ 7987.11
AWS Support (Business)	\$ 932.43
Support for all AWS services:	\$ 932.43
Total Monthly Payment:	\$ 10256.69

That Estimate of monthly bill is \$10256.69

STAGE SIX: Remarks about project improvements

AWS systems Manager

Systems manager is Powerful tools. Systems Manager provides a unified user interface that lets you view the operational data of multiple AWS services and automate operational tasks on AWS resources. With Systems Manager, you can group resources (for example, Amazon EC2 instances, Amazon S3 buckets, or Amazon RDS instances) by application, view operational data for monitoring and troubleshooting, and act on resource groups. Systems Manager simplifies resource and application management, reduces the time required to detect and resolve operational problems, and enables you to run and manage your infrastructure on a large scale with ease and security. (Mathew, 2014a)

AWS Config

AWS Config is a service that evaluates, reviews, and evaluates the configuration of AWS resources. Config continuously monitors and records AWS resource configurations and supports automatic configuration of records based on configuration requirements. And you can view the relationship between configuration changes and AWS resources, drill down into detailed resource configuration history, and determine if the configuration meets the configuration requirements specified in the internal guide. Absconding basically has the following features, configurable and customizable rules, configuration history of AWS

resources, software configuration history, configuration snapshots, resource relationship tracking, cloud management control panel, partner solution ecosystem. As a result, compliance audits, security analysis, change management, and operational troubleshooting will simplify these features.(Mathew, 2014b)

References

Ref - AWS CloudFormation [WWW Document], n.d. URL

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/intrinsic-function-reference-ref.html> (accessed 5.16.18).

“Ref - AWS CloudFormation.” Accessed May 16, 2018._

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/intrinsic-function-reference-ref.html>.

Mathew, S. (2014). Overview of Amazon Web Services. [online] Ajsnetworking.com.

Available at: <http://www.ajsnetworking.com/wp-content/uploads/2015/01/aws-overview.pdf> [Accessed 16 May 2018].

Docs.aws.amazon.com. (2018). Designing for High Availability - Exchange Server on AWS.

[online] Available at: <https://docs.aws.amazon.com/quickstart/latest/exchange/design-ha.html> [Accessed 14 May 2018].

Docs.aws.amazon.com. (2018). Creating a Trail - AWS CloudTrail. [online] Available at:

https://docs.aws.amazon.com/awscloudtrail/latest/userguide/cloudtrail-create-a-trail-using-the-console-first-time.html?icmpid=docs_cloudtrail_console [Accessed 14 May 2018].

Appendix I: Cloud Formation Design Diagram

