

Implantando um Application Load Balancer (ALB) na AWS com EC2

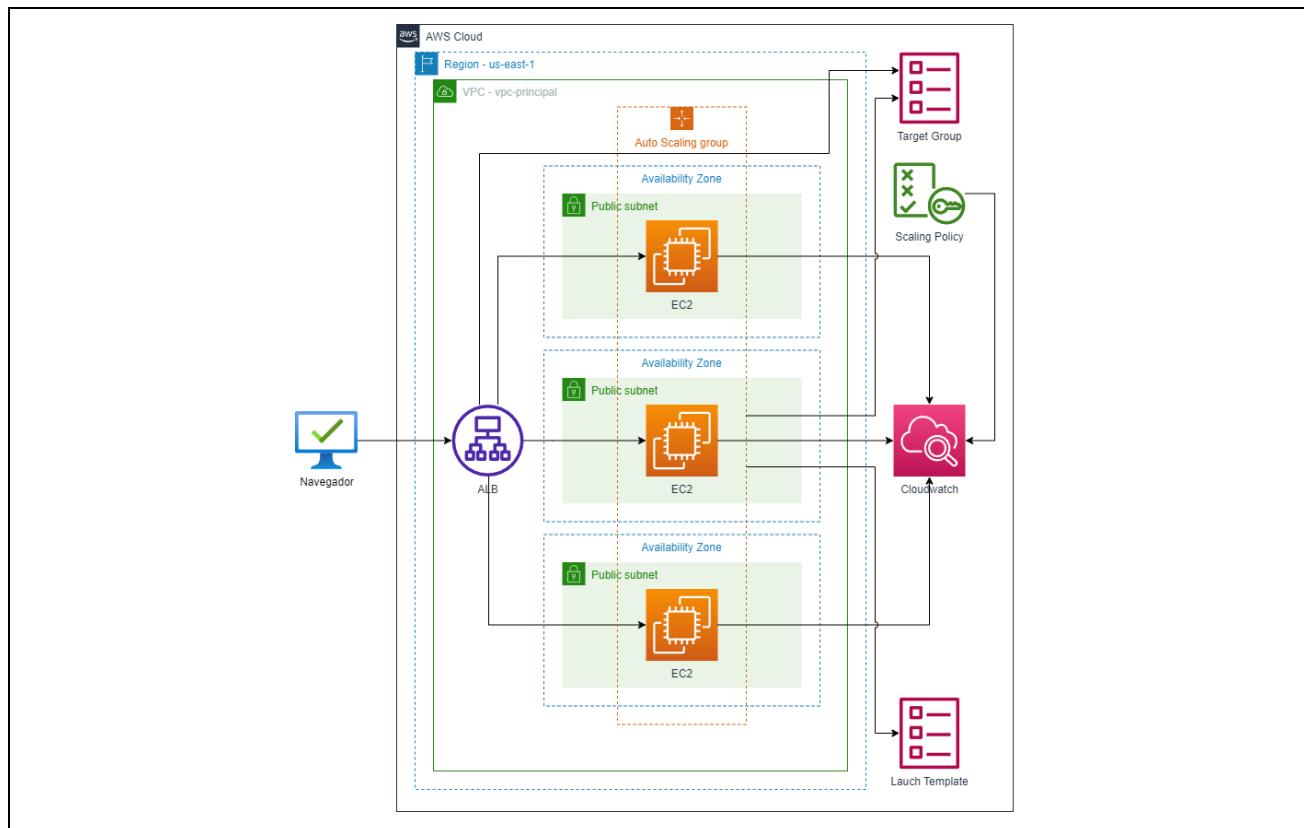


Prof. Thomás da Costa - <https://thomasdacosta.com.br>

Objetivo:

Implantar um Application Load Balancer na AWS direcionando o tráfego para as instâncias EC2 que possuem o servidor Apache HTTP instalado. Neste tutorial, iremos demonstrar como criar o ALB, Auto Scaling Group, Target Group e Launch Template.

Desenho da Solução:



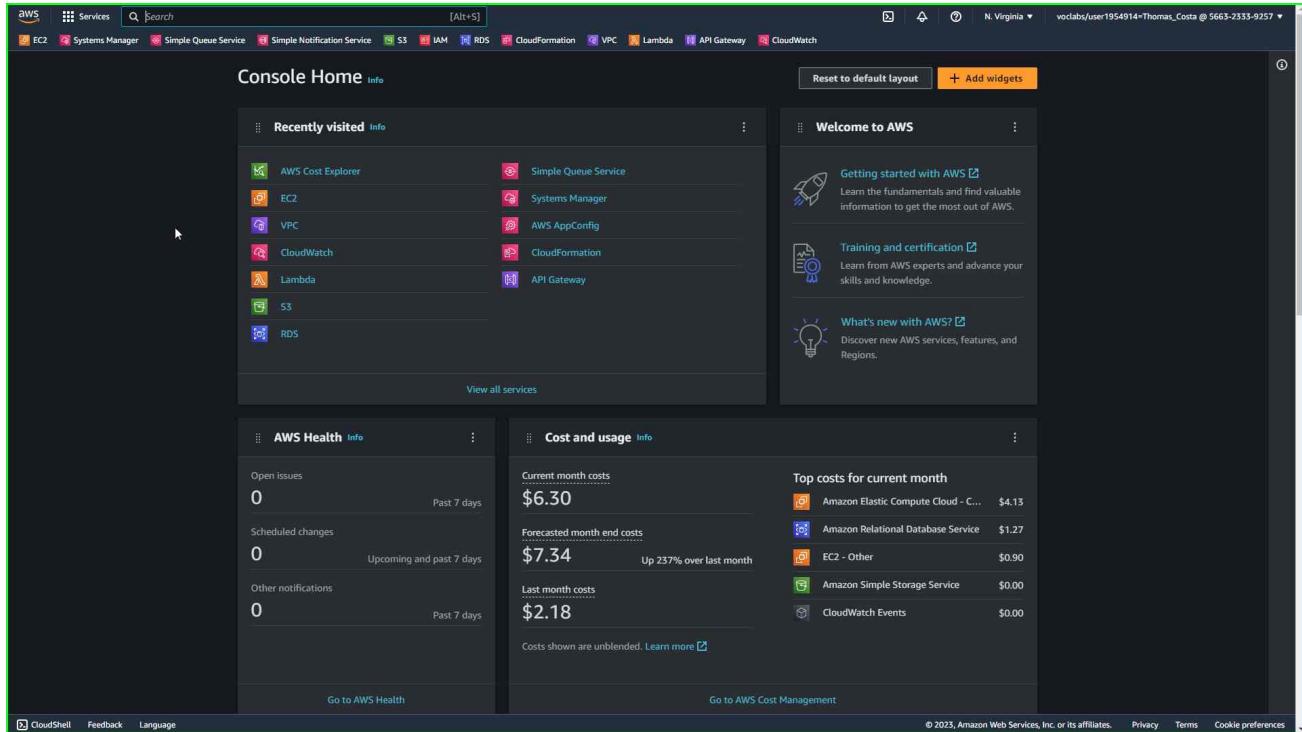
Premissas:

- Possuir um Security Group com acesso de entrada para a porta HTTP 80 para todos os endereços chamado SG-webserver;
- Possuir uma IAM Role chamada LabRole com as permissões e Policy necessárias para a execução dos recursos;
- Conhecer previamente o CloudWatch Logs e seu funcionamento.

Implantando um Application Load Balancer (ALB) na AWS com EC2

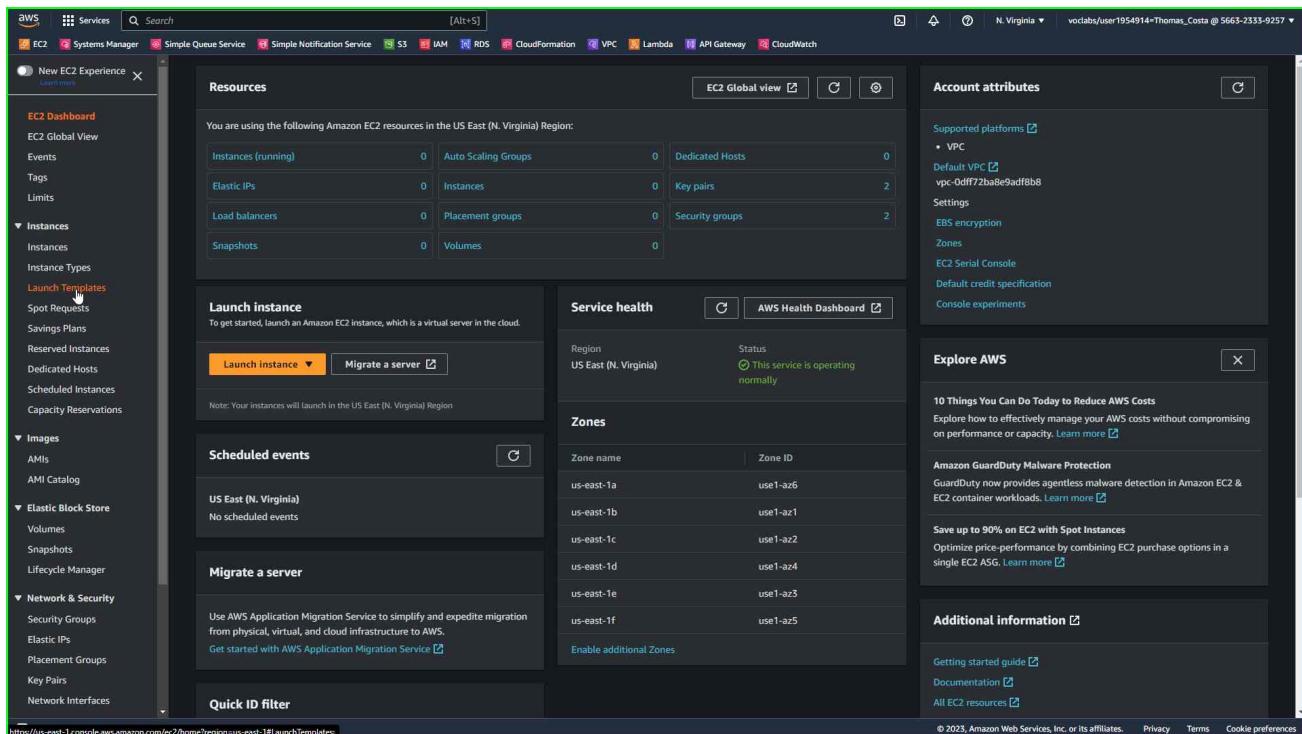
Parte 1 – Criando o Launch Template

Na tela principal, pesquisar a opção EC2:



The screenshot shows the AWS Console Home page. The top navigation bar has 'EC2' selected. The main content area includes sections for 'Recently visited' services (AWS Cost Explorer, EC2, VPC, CloudWatch, Lambda, S3, RDS) and 'Welcome to AWS' (Getting started with AWS, Training and certification, What's new with AWS). Below these are sections for 'AWS Health' (Open issues: 0, Scheduled changes: 0, Other notifications: 0) and 'Cost and usage' (Current month costs: \$6.30, Forecasted month end costs: \$7.34, Last month costs: \$2.18). A sidebar on the left shows the 'EC2 Dashboard' with various metrics: Instances (running) 0, Auto Scaling Groups 0, Dedicated Hosts 0, Elastic IPs 0, Instances 0, Key pairs 2, Load balancers 0, Placement groups 0, Security groups 2, Snapshots 0, Volumes 0. The URL in the address bar is <https://us-east-1.console.aws.amazon.com/ec2/home?region=us-east-1#LaunchTemplates>.

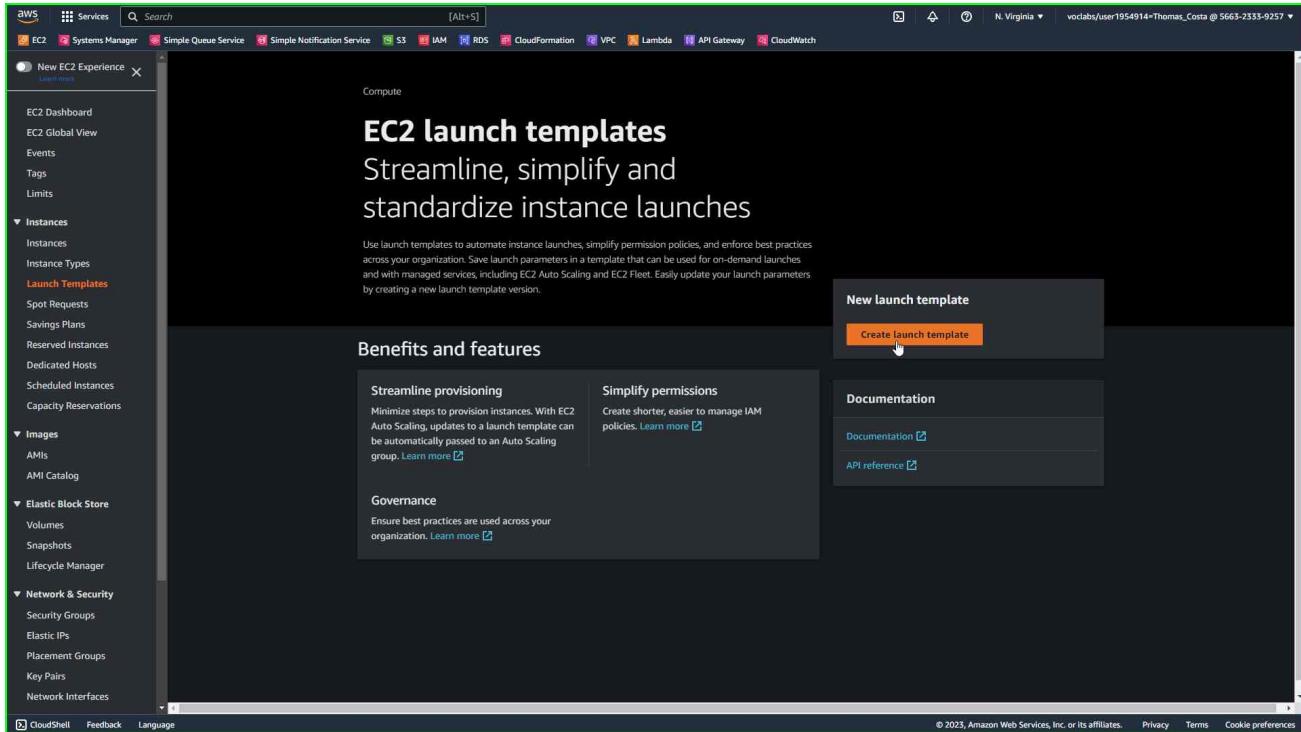
Selecione a opção Launch Templates:



The screenshot shows the AWS EC2 Dashboard. The left sidebar has 'Launch Templates' selected. The main content area includes sections for 'Resources' (listing Instances (running) 0, Auto Scaling Groups 0, Dedicated Hosts 0, Elastic IPs 0, Instances 0, Key pairs 2, Load balancers 0, Placement groups 0, Security groups 2, Snapshots 0, Volumes 0), 'Launch instance' (with 'Launch instance' and 'Migrate a server' buttons), 'Scheduled events' (No scheduled events), 'Migrate a server' (with 'Get started with AWS Application Migration Service' link), and 'Quick ID filter'. The right sidebar shows 'Account attributes' (Supported platforms: VPC, Default VPC: vpc-0dff72ba8e9adff8b8, Settings, EBS encryption, Zones, EC2 Serial Console, Default credit specification, Console experiments) and 'Explore AWS' (10 Things You Can Do Today to Reduce AWS Costs, Amazon GuardDuty Malware Protection, Save up to 90% on EC2 with Spot Instances, Additional information: Getting started guide, Documentation, All EC2 resources). The URL in the address bar is <https://us-east-1.console.aws.amazon.com/ec2/home?region=us-east-1#LaunchTemplates>.

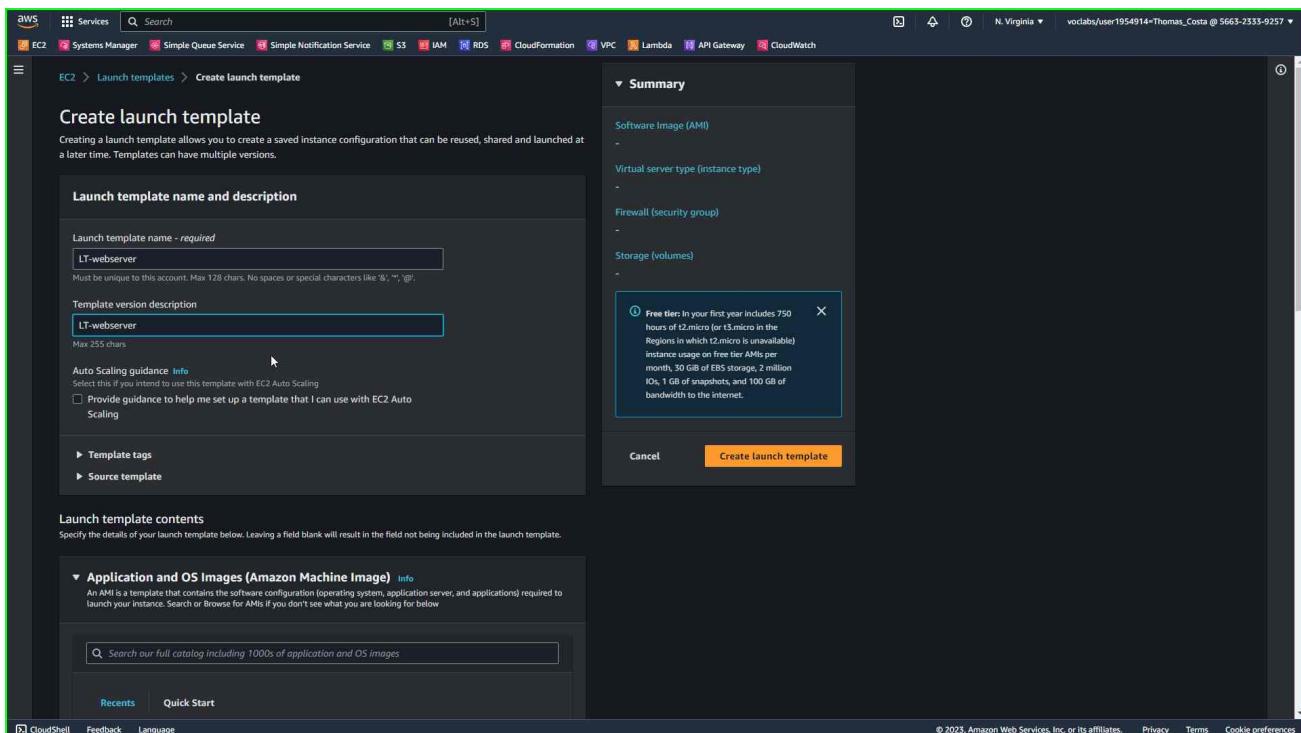
Implantando um Application Load Balancer (ALB) na AWS com EC2

Selecione o botão Create launch template:



The screenshot shows the AWS EC2 Launch Templates page. On the left, a sidebar lists various EC2 features: EC2 Dashboard, EC2 Global View, Events, Tags, Limits, Instances (Instances, Instance Types, Launch Templates), Images (AMIs, AMI Catalog), Elastic Block Store (Volumes, Snapshots, Lifecycle Manager), Network & Security (Security Groups, Elastic IPs, Placement Groups, Key Pairs, Network Interfaces). The 'Launch Templates' section is currently selected. The main content area is titled 'EC2 launch templates' and 'Streamline, simplify and standardize instance launches'. It includes a section on 'Benefits and features' with three items: 'Streamline provisioning', 'Simplify permissions', and 'Governance'. A large orange 'Create launch template' button is prominently displayed. The top right corner shows the region 'N. Virginia' and a user profile.

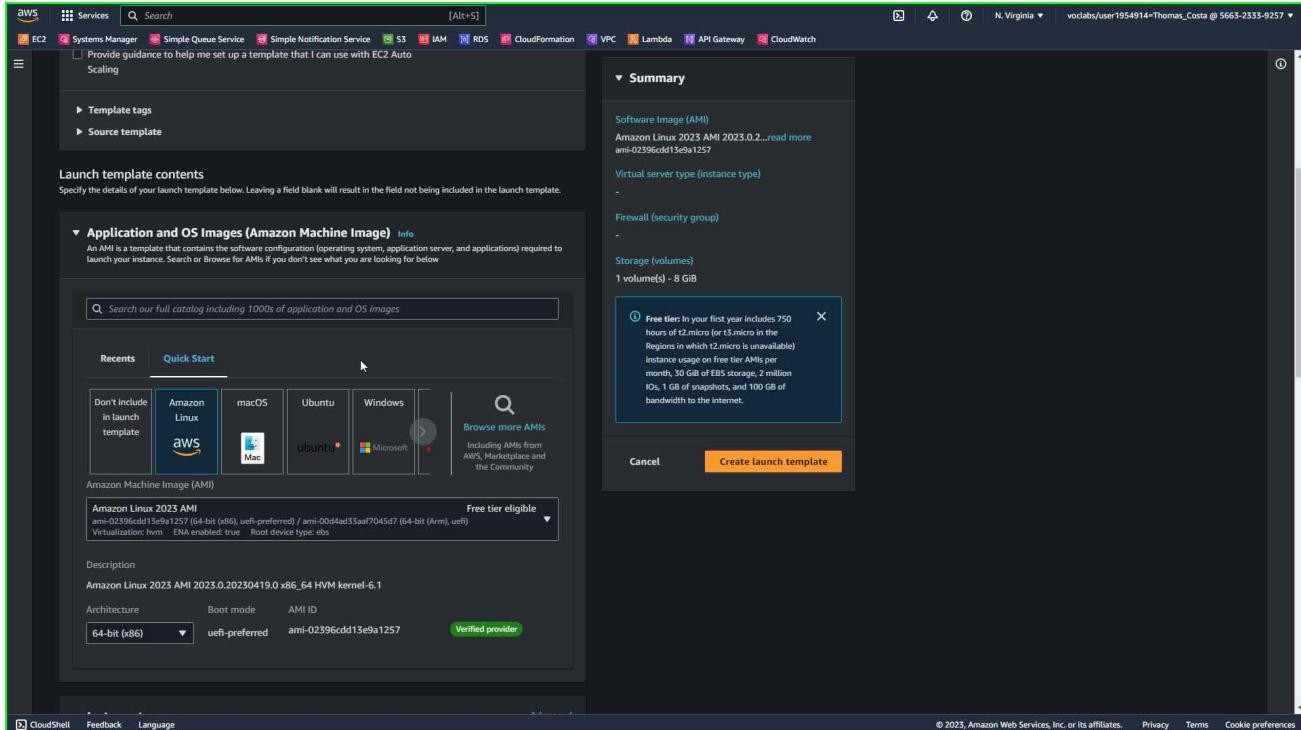
Incluir o nome do template, No nosso caso LT-webservice:



The screenshot shows the 'Create launch template' wizard. The first step, 'Launch template name and description', is displayed. The 'Launch template name - required' field contains 'LT-webservice'. The 'Template version description' field also contains 'LT-webservice'. Below these fields is an 'Auto Scaling guidance' section with a checkbox for 'Provide guidance to help me set up a template that I can use with EC2 Auto Scaling'. The right side of the screen shows a 'Summary' section with collapsed sections for 'Software Image (AMI)', 'Virtual server type (instance type)', 'Firewall (security group)', and 'Storage (volumes)'. A tooltip for the 'Free tier' is visible, stating: 'Free tier: In your first year includes 750 hours of t2.micro (or t3.micro in the Regions in which t2.micro is unavailable) instance usage on free tier AMIs per month, 30 GB of EBS storage, 2 million I/Os, 1 GB of snapshots, and 100 GB of bandwidth to the internet.' At the bottom are 'Cancel' and 'Create launch template' buttons. The bottom of the screen shows the AWS navigation bar and a footer with copyright information.

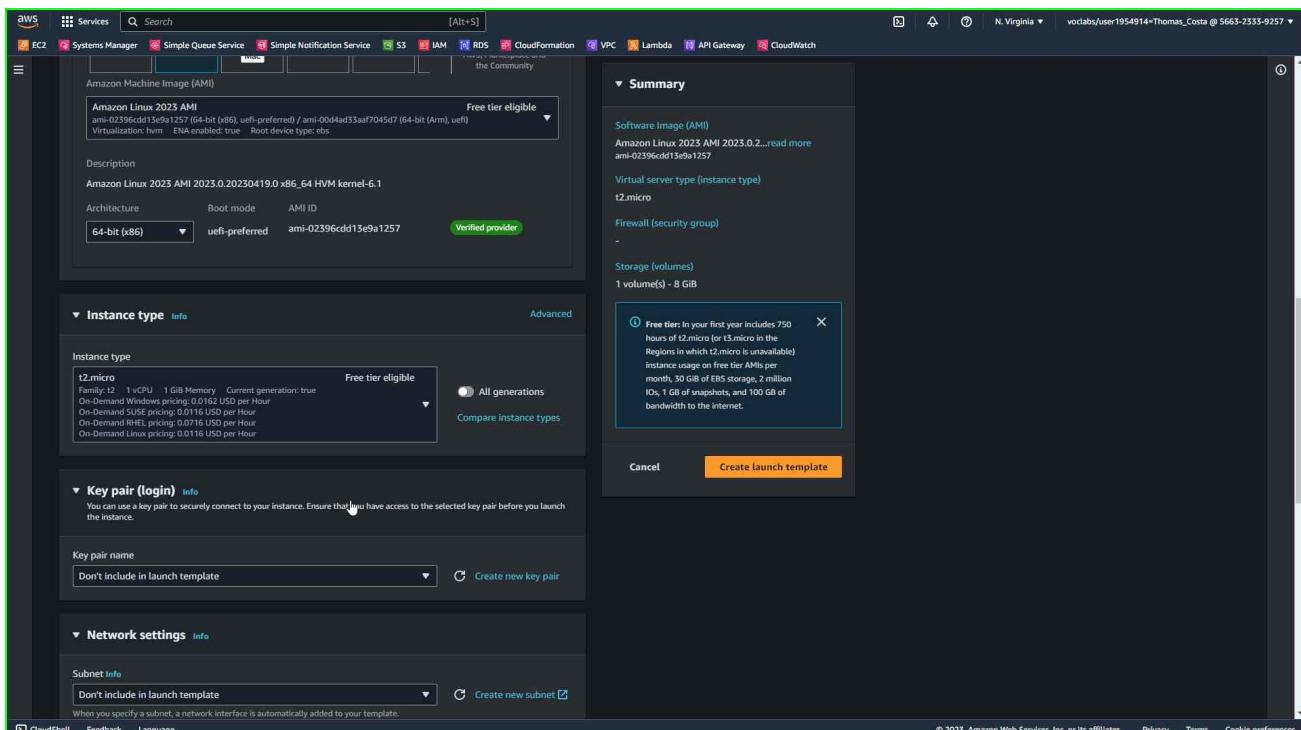
Implantando um Application Load Balancer (ALB) na AWS com EC2

Selecionar a imagem do Amazon Linux:



The screenshot shows the AWS EC2 console with the 'Launch template' creation wizard. In the 'Application and OS Images (Amazon Machine Image) Info' section, the 'Amazon Linux 2023 AMI' is selected. The instance is described as 'Free tier eligible'. Below this, the 'Architecture' is set to '64-bit (x86)', 'Boot mode' to 'uefi-preferred', and 'AMI ID' to 'ami-02396cd15e9a1257'. The 'Verified provider' button is visible. The 'Summary' section on the right provides a summary of the selected AMI, including its software image (Amazon Linux 2023 AMI 2023.0.2), virtual server type (t2.micro), and storage (1 volume(s) - 8 GiB). A tooltip for the 'Free tier' is displayed, stating: 'Free tier: In your first year includes 750 hours of t2.micro (or t3.micro in the Regions in which t2.micro is unavailable) instance usage on free tier AMIs per month, 30 GiB of EBS storage, 2 million I/Os, 1 GiB of snapshots, and 100 GiB of bandwidth to the internet.'

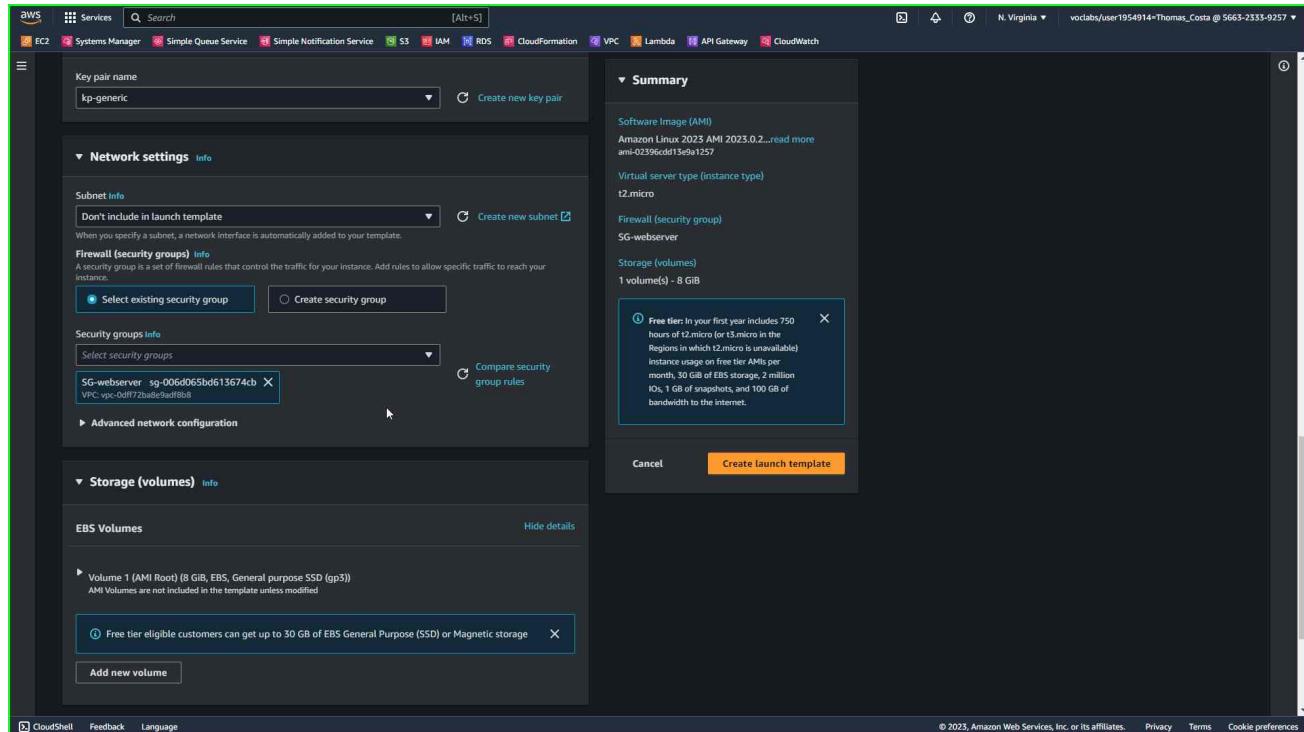
Escolher a instância do tipo t2.micro:



The screenshot shows the continuation of the AWS EC2 'Launch template' creation wizard. In the 'Instance type' section, the 't2.micro' instance type is selected. It is described as 'Free tier eligible' and part of the 't2.micro' family. The 'Current generation' is noted as 'true'. Pricing information is provided for On-Demand Windows, SUSE, RHEL, and Linux. The 'All generations' and 'Compare instance types' buttons are available. The 'Summary' section on the right reiterates the instance type (t2.micro), software image (Amazon Linux 2023 AMI 2023.0.2), and storage (1 volume(s) - 8 GiB). A tooltip for the 'Free tier' is also present.

Implantando um Application Load Balancer (ALB) na AWS com EC2

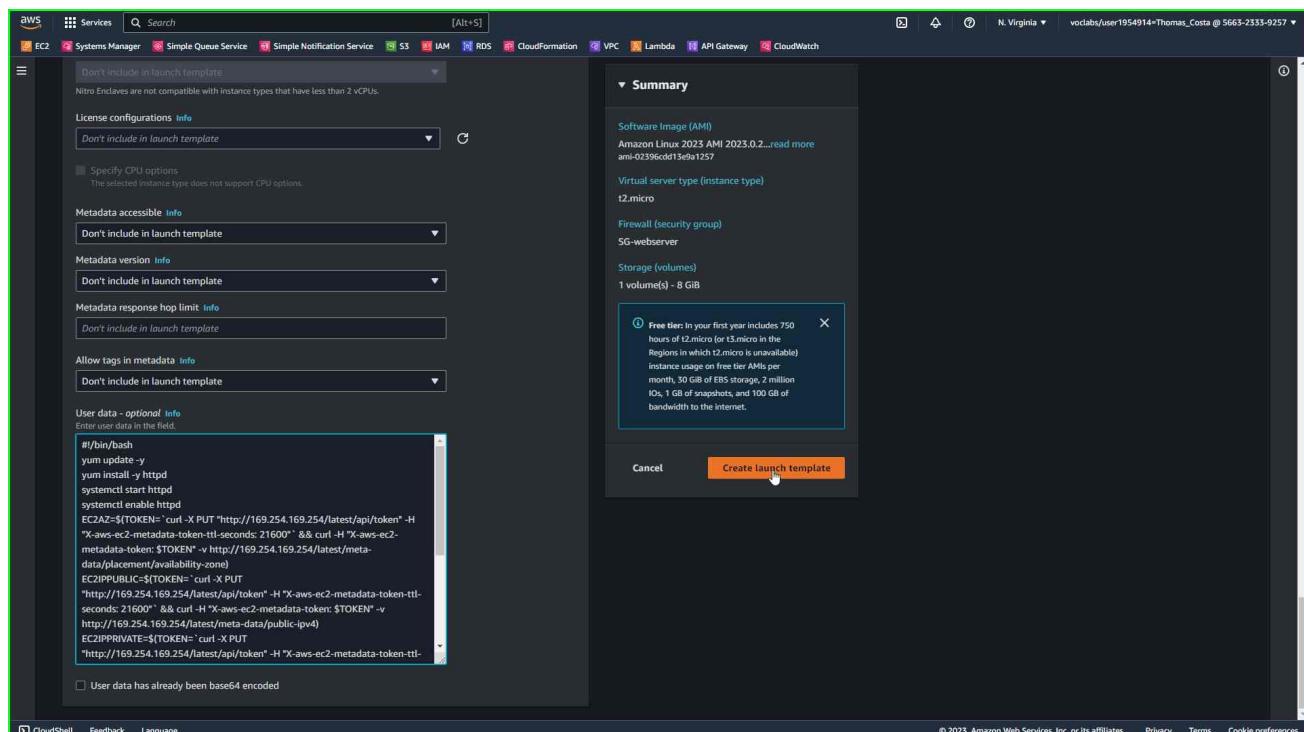
Selecionar o Security Group chamado **SG-webserver**:



The screenshot shows the AWS EC2 Launch Wizard. In the 'Firewall (security group)' section, the 'Select existing security group' radio button is selected, and the dropdown menu shows 'SG-webserver'. A tooltip for the 'Free tier' is visible, stating: 'Free tier: In your first year includes 750 hours of t2.micro (or t3.micro in the Regions in which t2.micro is unavailable) instance usage on free tier AMIs per month, 30 GiB of EBS storage, 2 million I/Os, 1 GiB of snapshots, and 100 GiB of bandwidth to the internet.' At the bottom right of the wizard, the 'Create launch template' button is highlighted.

Incluir o script de inicialização no User data e clicar no botão Create launch template.

O script se encontra no final desse documento:

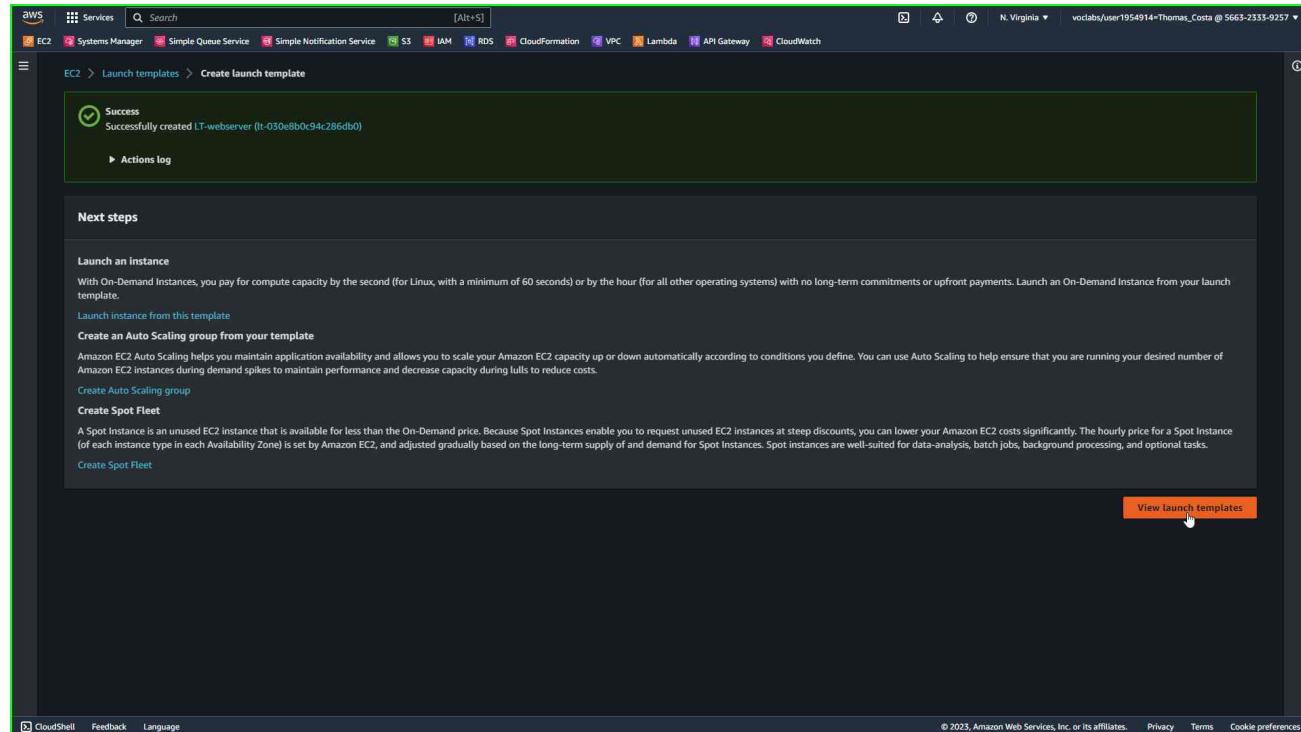


The screenshot shows the AWS EC2 Launch Wizard with the 'User data - optional' section expanded. The 'User data' field contains a shell script for initializing an EC2 instance. The script includes commands for updating yum, installing httpd, starting the httpd service, enabling systemctl, and performing AWS metadata token retrieval. A tooltip for the 'Free tier' is visible. At the bottom right of the wizard, the 'Create launch template' button is highlighted.

```
#!/bin/bash
yum update -y
yum install -y httpd
systemctl start httpd
systemctl enable httpd
EC2AZ=$TOKEN= curl -X PUT "http://169.254.169.254/latest/api/token" -H
"X-aws-ec2-metadata-token-ttl-seconds: 21600" && curl -H "X-aws-ec2-
metadata-token: $TOKEN" -v http://169.254.169.254/latest/meta-
data/placement/availability-zone
EC2IPPUBLIC=$TOKEN= curl -X PUT
"http://169.254.169.254/latest/api/token" -H "X-aws-ec2-metadata-token-ttl-
seconds: 21600" && curl -H "X-aws-ec2-metadata-token: $TOKEN" -v
http://169.254.169.254/latest/meta-data/public-ipv4
EC2IPPRIVATE=$TOKEN= curl -X PUT
"http://169.254.169.254/latest/api/token" -H "X-aws-ec2-metadata-token-ttl-
```

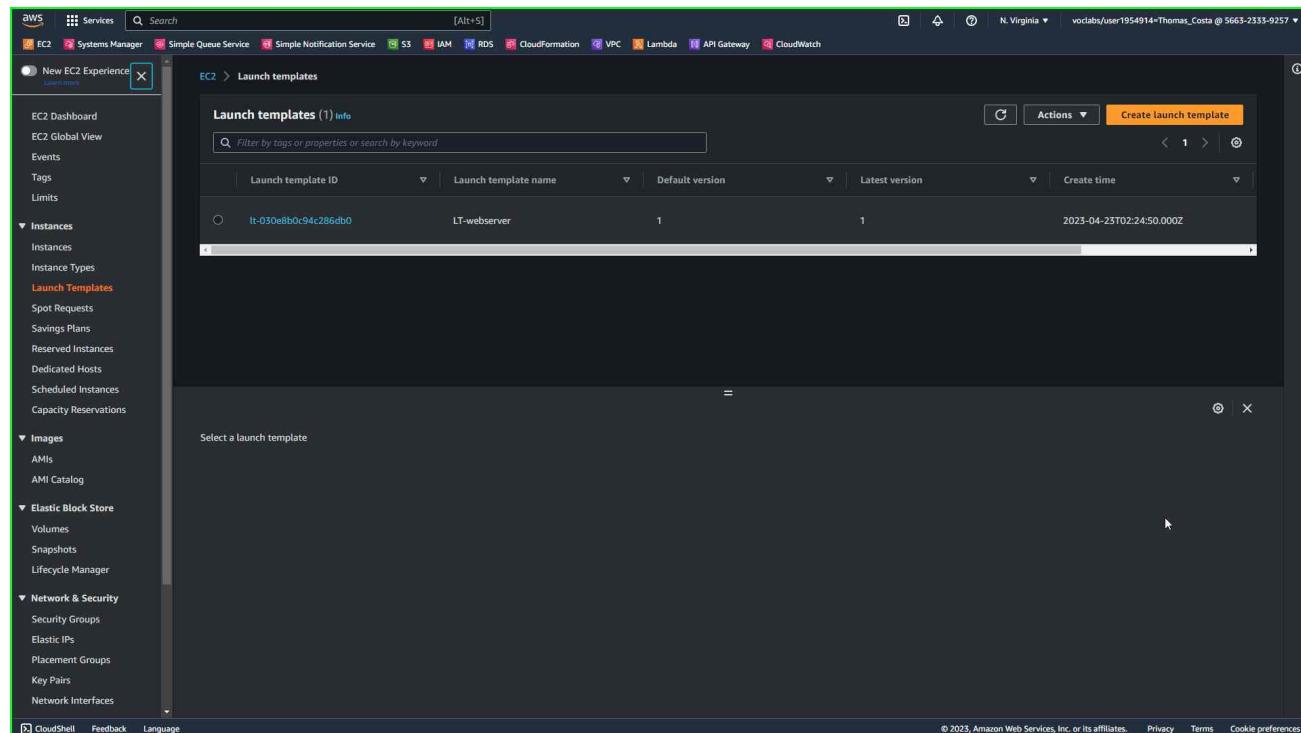
Implantando um Application Load Balancer (ALB) na AWS com EC2

Launch Template criado com sucesso:



The screenshot shows the AWS EC2 Launch Templates creation page. At the top, a green success message box displays: "Successfully created LT-webserver (lt-030e8b0c94c286db0)". Below this, an "Actions log" button is visible. The main content area is titled "Next steps" and includes several options: "Launch an instance", "Launch instance from this template", "Create an Auto Scaling group from your template", "Create Auto Scaling group", and "Create Spot Fleet". Each option has a brief description. At the bottom right of the content area is an orange "View launch templates" button, which is highlighted with a mouse cursor. The bottom of the page includes standard AWS navigation links: CloudShell, Feedback, Language, and copyright information: © 2023, Amazon Web Services, Inc. or its affiliates. Privacy Terms Cookie preferences.

Entre na tela inicial do Launch templates e temos o template criado:

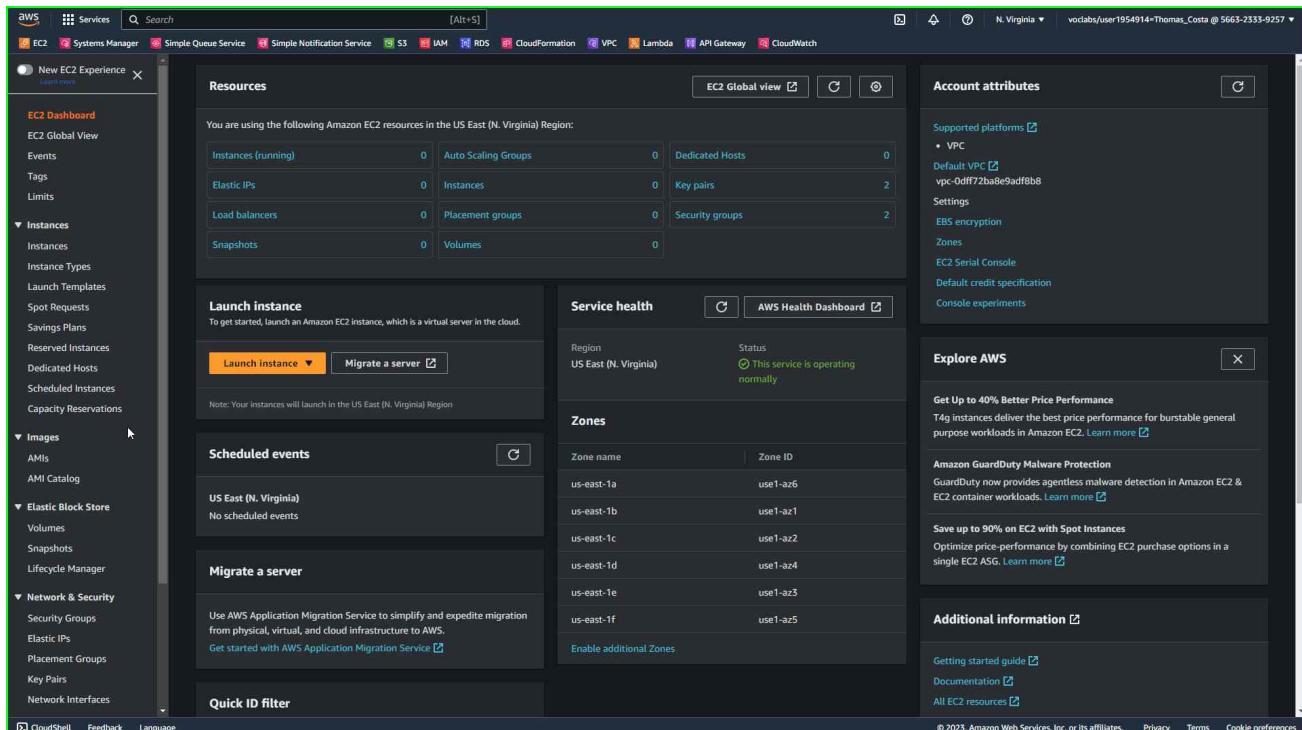


The screenshot shows the AWS EC2 Launch Templates list page. The left sidebar is expanded, showing the "Launch Templates" section under "Instances". The main content area is titled "Launch templates (1) Info" and displays a table with one row. The table columns are: Launch template ID, Launch template name, Default version, Latest version, and Create time. The data in the table is: lt-030e8b0c94c286db0, LT-webserver, 1, 1, 2023-04-23T02:24:50.000Z. Below the table, a message says "Select a launch template". The bottom of the page includes standard AWS navigation links: CloudShell, Feedback, Language, and copyright information: © 2023, Amazon Web Services, Inc. or its affiliates. Privacy Terms Cookie preferences.

Implantando um Application Load Balancer (ALB) na AWS com EC2

Parte 2 – Criando o EC2 Auto Scaling Group (ASG)

Selecione a opção Auto Scaling Groups:

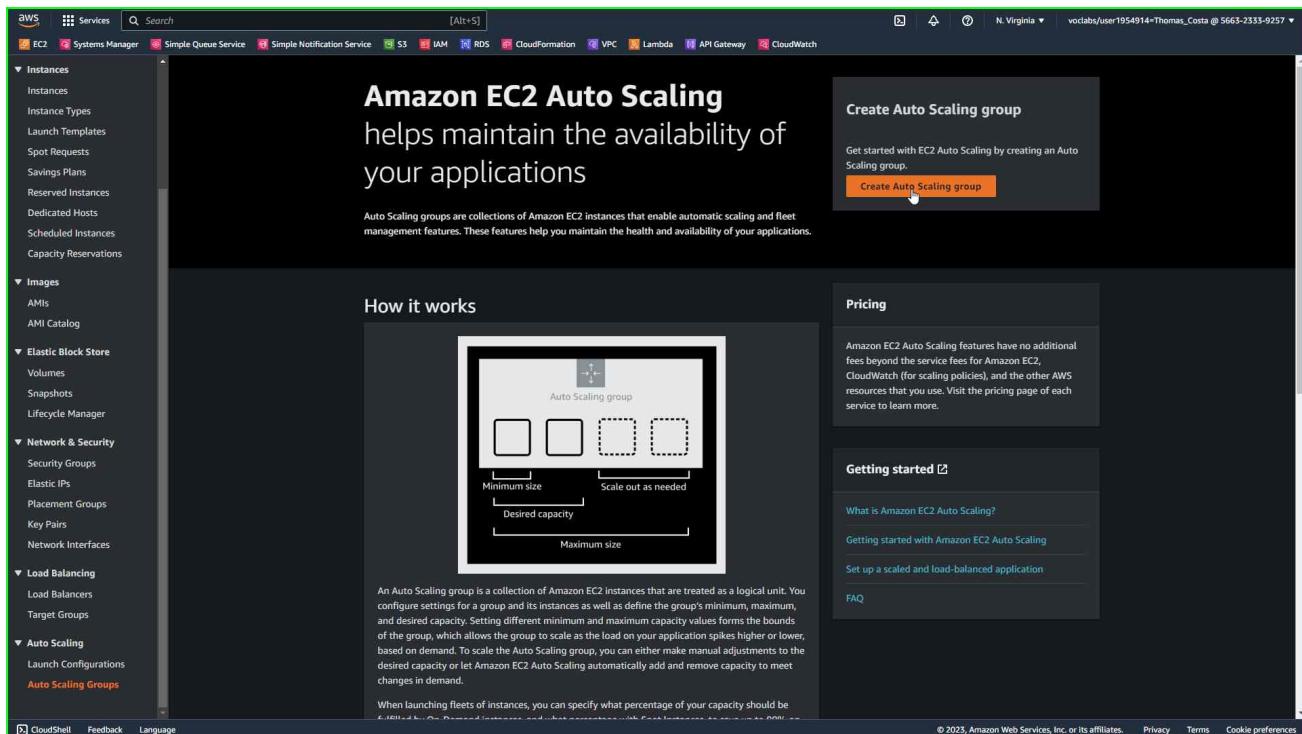


The screenshot shows the AWS EC2 Dashboard. The left sidebar is expanded to show the 'Instances' section, with 'Auto Scaling Groups' highlighted. The main content area displays a table of resources:

Instances (running)	0	Auto Scaling Groups	0	Dedicated Hosts	0
Elastic IPs	0	Instances	0	Key pairs	2
Load balancers	0	Placement groups	0	Security groups	2
Snapshots	0	Volumes	0		

Below the table are links for 'Launch instance', 'Scheduled events', 'Migrate a server', and 'Quick ID filter'.

Clique no botão Create Auto Scaling group:



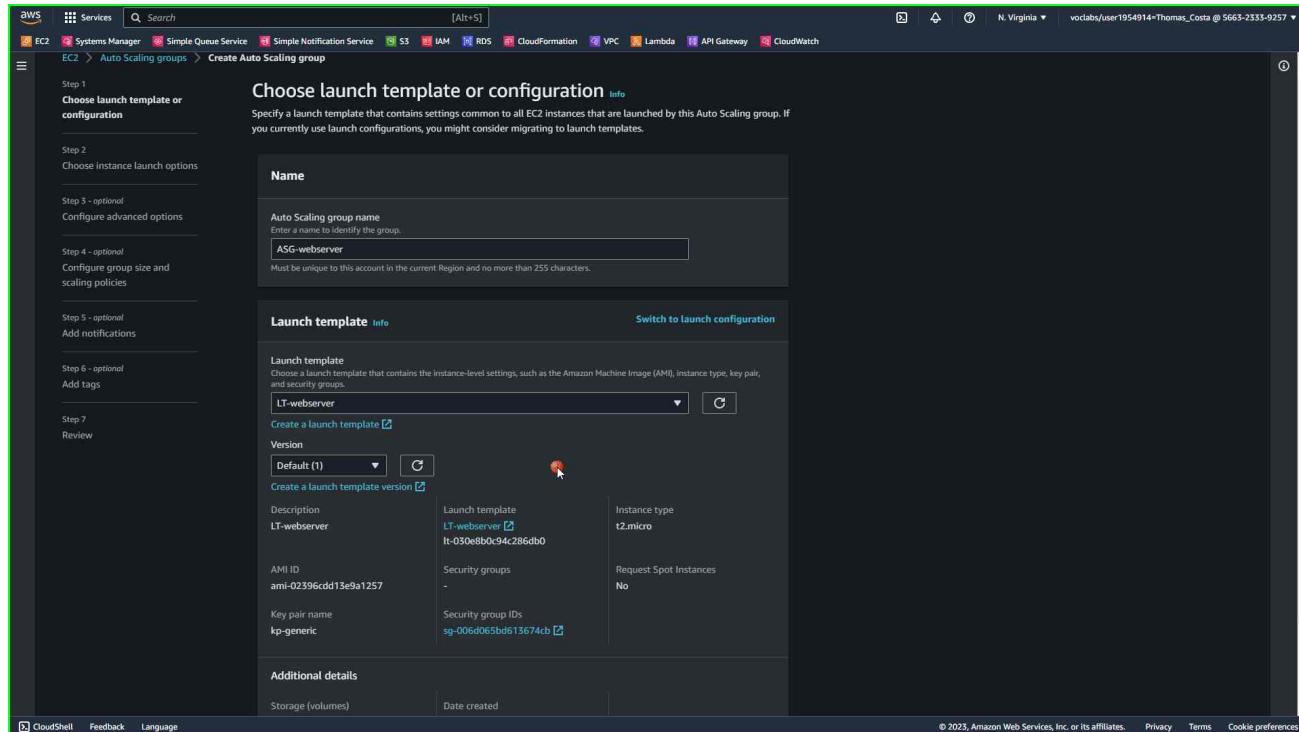
The screenshot shows the 'Create Auto Scaling group' wizard. The first step, 'How it works', is displayed. It includes a diagram of an Auto Scaling group with four instances and labels for 'Minimum size', 'Desired capacity', and 'Maximum size'. Text explains the concept of an Auto Scaling group:

An Auto Scaling group is a collection of Amazon EC2 instances that are treated as a logical unit. You configure settings for a group and its instances as well as define the group's minimum, maximum, and desired capacity. Setting different minimum and maximum capacity values forms the bounds of the group, which allows the group to scale as the load on your application spikes higher or lower, based on demand. To scale the Auto Scaling group, you can either make manual adjustments to the desired capacity or let Amazon EC2 Auto Scaling automatically add and remove capacity to meet changes in demand.

When launching fleets of instances, you can specify what percentage of your capacity should be

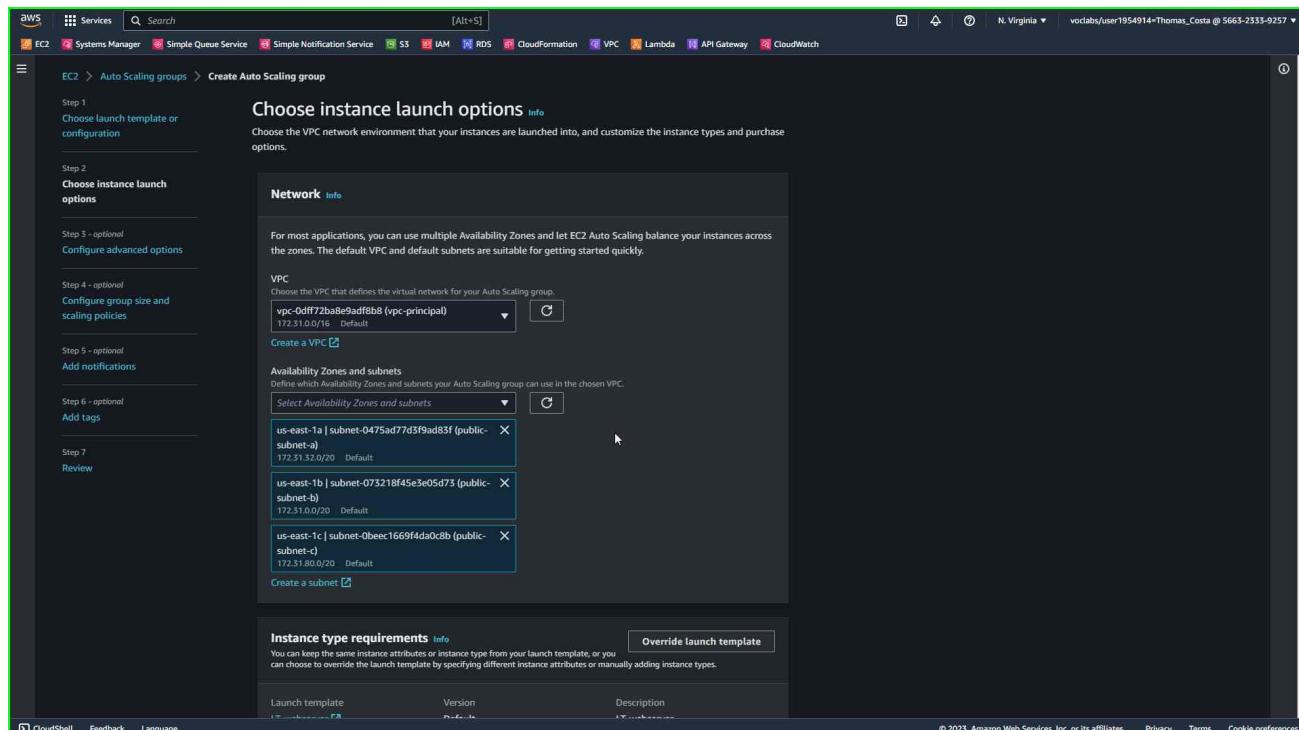
Implantando um Application Load Balancer (ALB) na AWS com EC2

Inclua o nome do ASG. No nosso exemplo chama-se ASG-webserver:



The screenshot shows the 'Create Auto Scaling group' wizard in the AWS EC2 console. The current step is 'Step 1: Choose launch template or configuration'. The 'Name' field is populated with 'ASG-webserver'. Below it, the 'Launch template' section shows a table with instance details: Description 'LT-webserver', Launch template 'LT-webserver', Instance type 't2.micro', AMI ID 'ami-02396cdd13e9a1257', Security groups ' - ', and Key pair name 'kp-generic'. The 'Additional details' section shows 'Storage (volumes)' and 'Date created'.

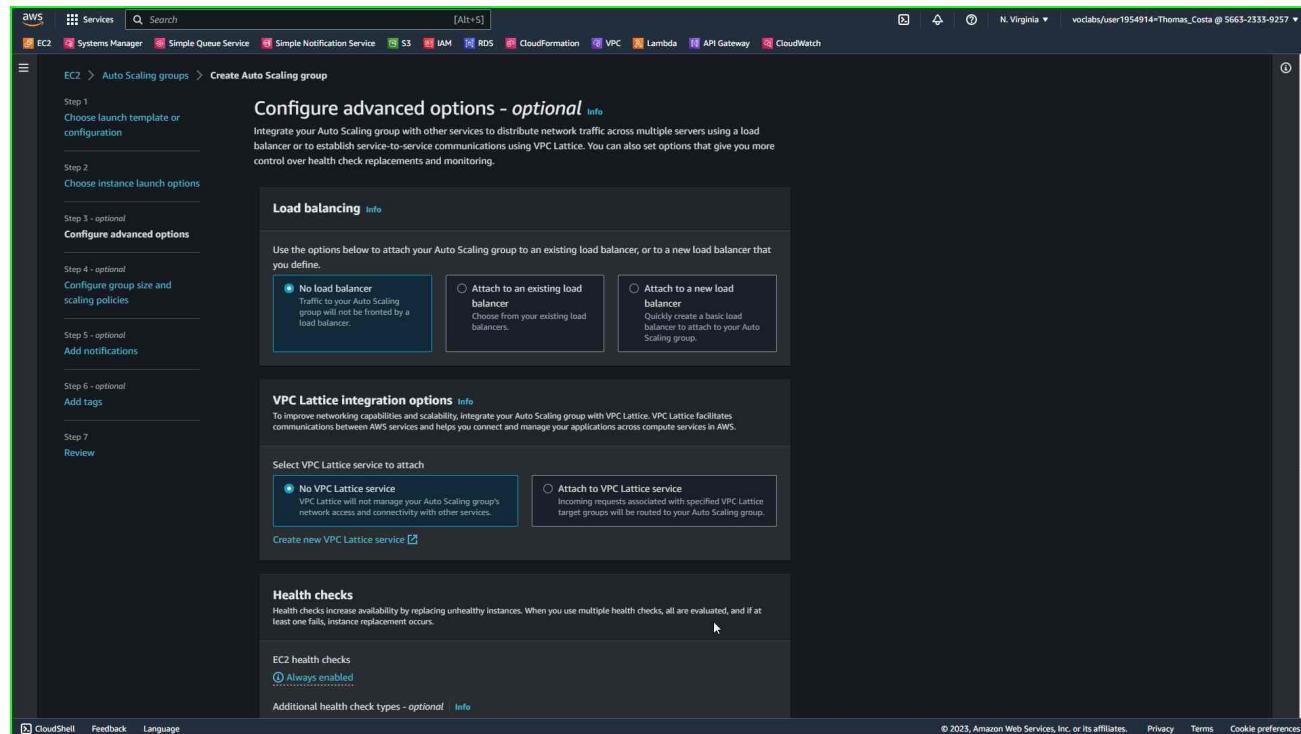
Selecione no mínimo 3 AZ para o ASG:



The screenshot shows the 'Create Auto Scaling group' wizard in the AWS EC2 console, progressing to 'Step 2: Choose instance launch options'. The 'Network' section shows the VPC 'vpc-0df77bab9e9adffbb8 (vpc-principal)'. The 'Availability Zones and subnets' section lists three subnets: 'us-east-1a | subnet-0475ad77d3f9ad83f (public-subnet-a)', 'us-east-1b | subnet-073218f45e3e05d73 (public-subnet-b)', and 'us-east-1c | subnet-0beecc1669f4da0c8b (public-subnet-c)'. The 'Instance type requirements' section is at the bottom.

Implantando um Application Load Balancer (ALB) na AWS com EC2

Selecione a opção No load balancer:



Configure advanced options - *optional* [Info](#)

Integrate your Auto Scaling group with other services to distribute network traffic across multiple servers using a load balancer or to establish service-to-service communications using VPC Lattice. You can also set options that give you more control over health check replacements and monitoring.

Load balancing [Info](#)

Use the options below to attach your Auto Scaling group to an existing load balancer, or to a new load balancer that you define.

No load balancer
Traffic to your Auto Scaling group will not be fronted by a load balancer.

Attach to an existing load balancer
Choose from your existing load balancers.

Attach to a new load balancer
Quickly create a basic load balancer to attach to your Auto Scaling group.

VPC Lattice integration options [Info](#)

To improve networking capabilities and scalability, integrate your Auto Scaling group with VPC Lattice. VPC Lattice facilitates communications between AWS services and helps you connect and manage your applications across compute services in AWS.

Select VPC Lattice service to attach

No VPC Lattice service
VPC Lattice will not manage your Auto Scaling group's network access and connectivity with other services.

Attach to VPC Lattice service
Incoming requests associated with specified VPC Lattice target groups will be routed to your Auto Scaling group.

[Create new VPC Lattice service](#)

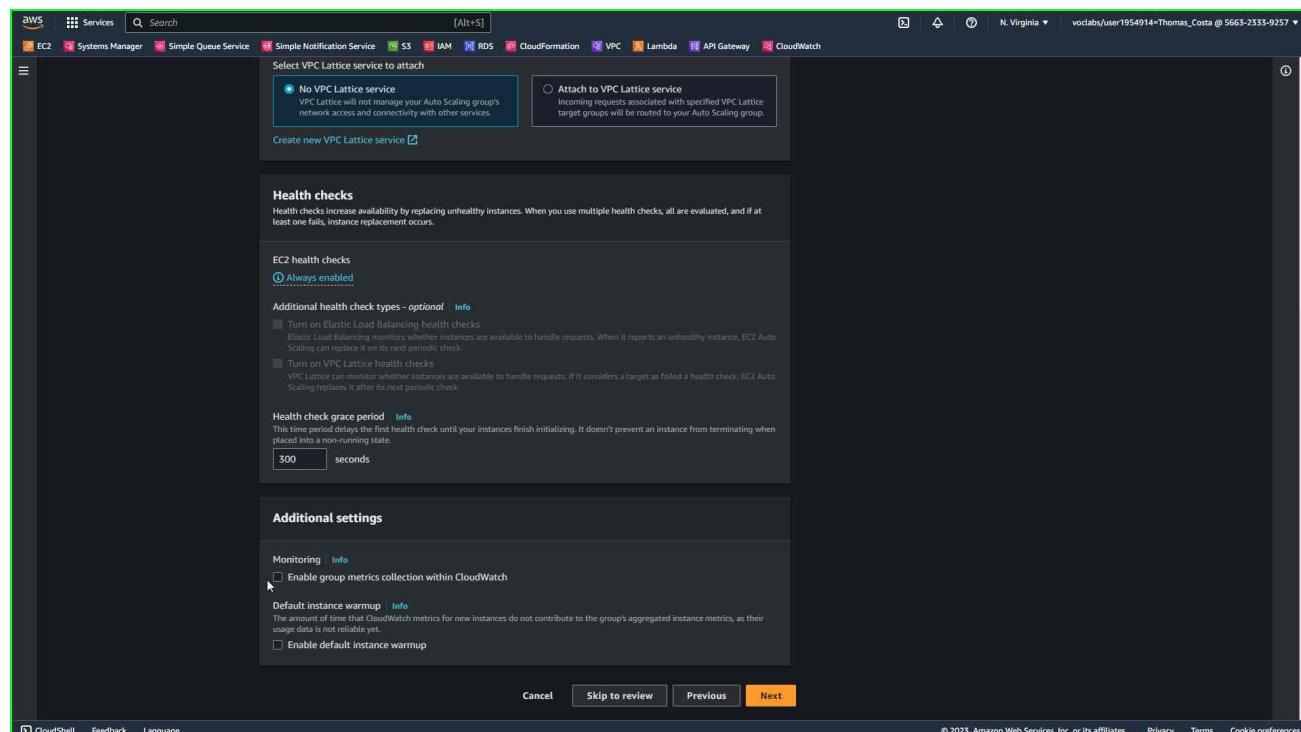
Health checks

Health checks increase availability by replacing unhealthy instances. When you use multiple health checks, all are evaluated, and if at least one fails, instance replacement occurs.

EC2 health checks
Always enabled

Additional health check types - optional [Info](#)

Selecione a opção Enable group metrics collection within CloudWatch:



Select VPC Lattice service to attach

No VPC Lattice service
VPC Lattice will not manage your Auto Scaling group's network access and connectivity with other services.

Attach to VPC Lattice service
Incoming requests associated with specified VPC Lattice target groups will be routed to your Auto Scaling group.

[Create new VPC Lattice service](#)

Health checks

Health checks increase availability by replacing unhealthy instances. When you use multiple health checks, all are evaluated, and if at least one fails, instance replacement occurs.

EC2 health checks
Always enabled

Additional health check types - optional [Info](#)

Turn on Elastic Load Balancing health checks
Elastic Load Balancing periodically checks instances are available to handle requests. When it reports an unhealthy instance, EC2 Auto Scaling can replace it on its next periodic check.

Turn on VPC Lattice health checks
VPC Lattice can monitor whether instances are available to handle requests. If it considers a target as failed a health check, EC2 Auto Scaling replaces it after its next periodic check.

Health check grace period [Info](#)

This time period delays the first health check until your instances finish initializing. It doesn't prevent an instance from terminating when placed into a non-running state.

300 seconds

Additional settings

Monitoring [Info](#)

Enable group metrics collection within CloudWatch

Default instance warmup [Info](#)

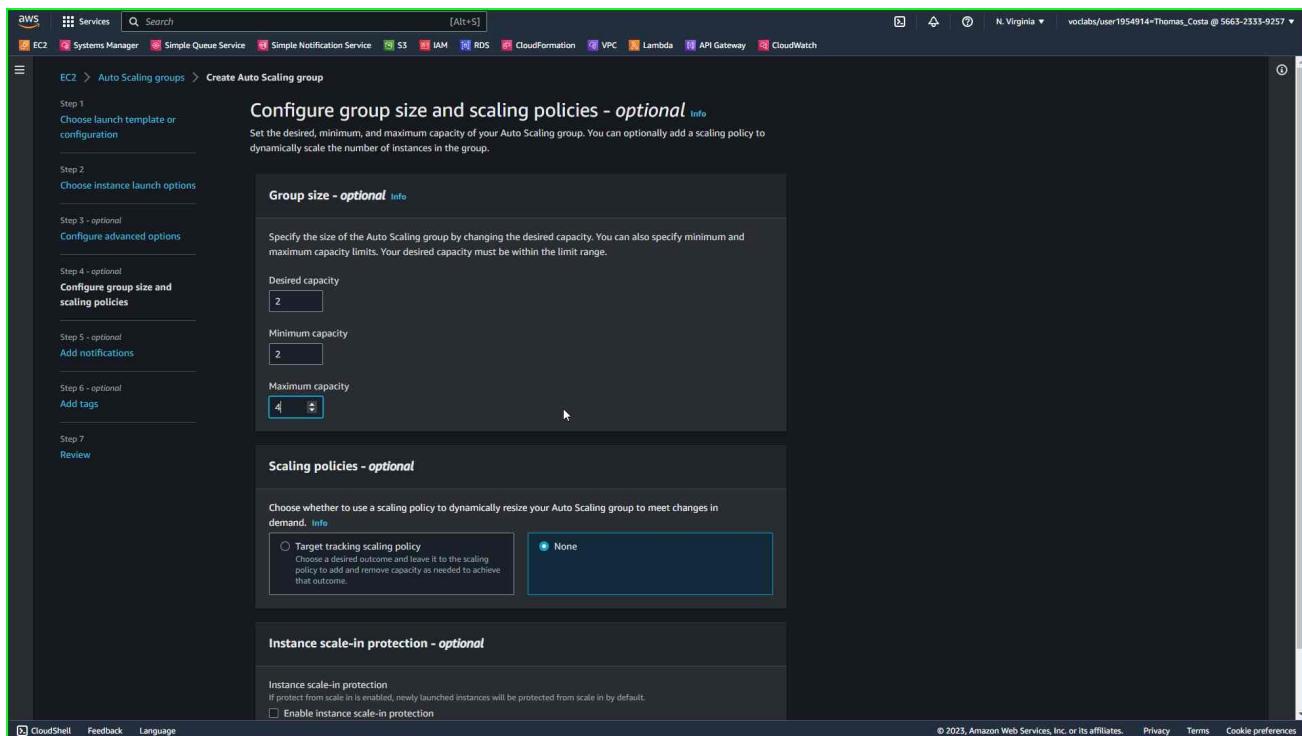
The amount of time that CloudWatch metrics for new instances do not contribute to the group's aggregated instance metrics, as their usage data is not reliable yet.

Enable default instance warmup

[Cancel](#) [Skip to review](#) [Previous](#) [Next](#)

Implantando um Application Load Balancer (ALB) na AWS com EC2

Inclua os seguintes valores para os servidores: Desired Capacity e Minimum Capacity é igual a 2
Maximum Capacity é igual a 4:



Configure group size and scaling policies - *optional* [Info](#)
Set the desired, minimum, and maximum capacity of your Auto Scaling group. You can optionally add a scaling policy to dynamically scale the number of instances in the group.

Group size - *optional* [Info](#)

Specify the size of the Auto Scaling group by changing the desired capacity. You can also specify minimum and maximum capacity limits. Your desired capacity must be within the limit range.

Desired capacity: 2

Minimum capacity: 2

Maximum capacity: 4

Scaling policies - *optional*

Choose whether to use a scaling policy to dynamically resize your Auto Scaling group to meet changes in demand. [Info](#)

Target tracking scaling policy
Choose a desired outcome and leave it to the scaling policy to add and remove capacity as needed to achieve that outcome.

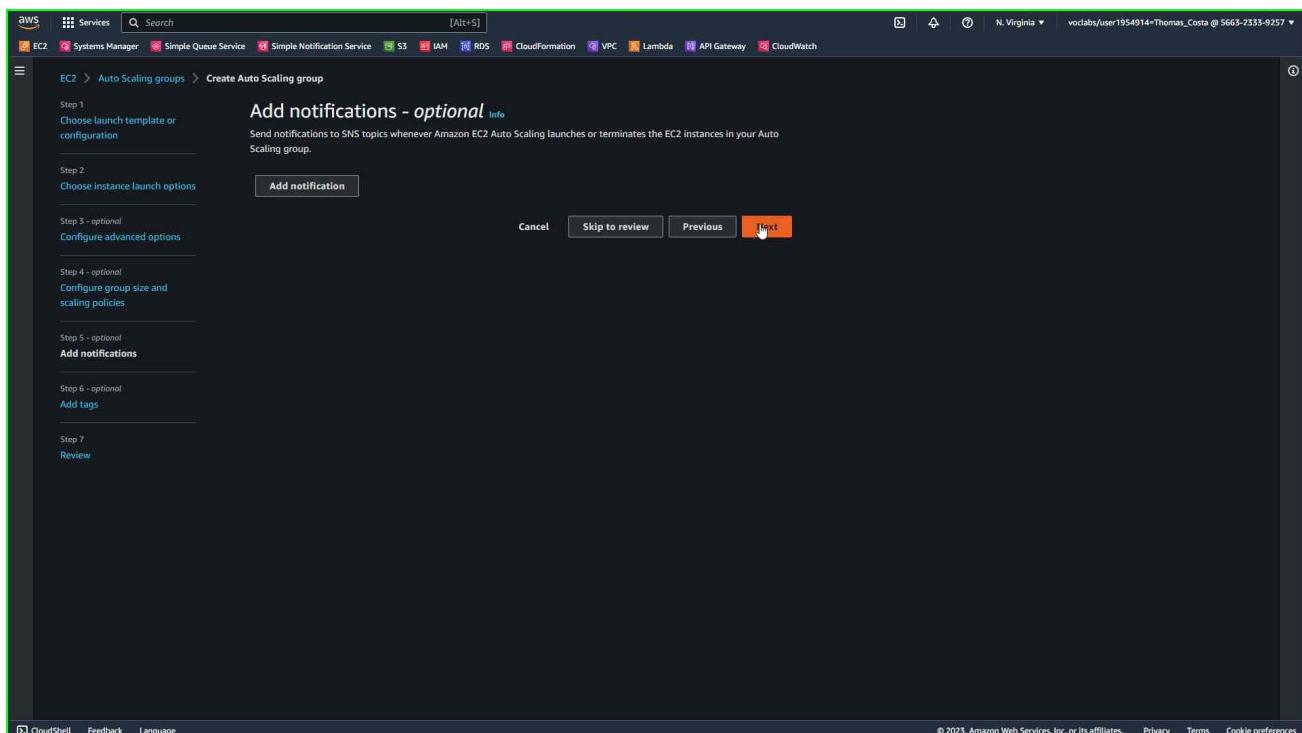
None

Instance scale-in protection - *optional*

Instance scale-in protection
If protect from scale in is enabled, newly launched instances will be protected from scale in by default.

Enable instance scale-in protection

Na próxima tela não tem ação somente clique em **Next**:



Add notifications - *optional* [Info](#)

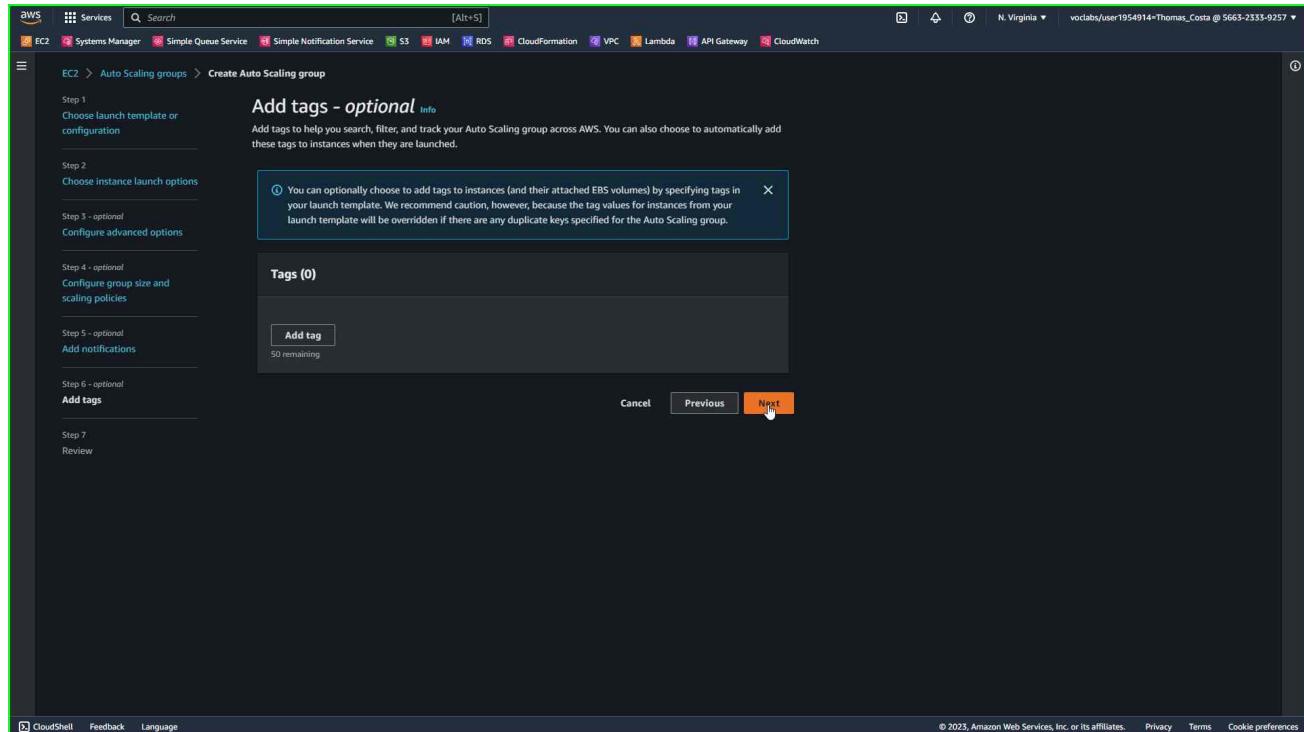
Send notifications to SNS topics whenever Amazon EC2 Auto Scaling launches or terminates the EC2 instances in your Auto Scaling group.

Add notification

Next

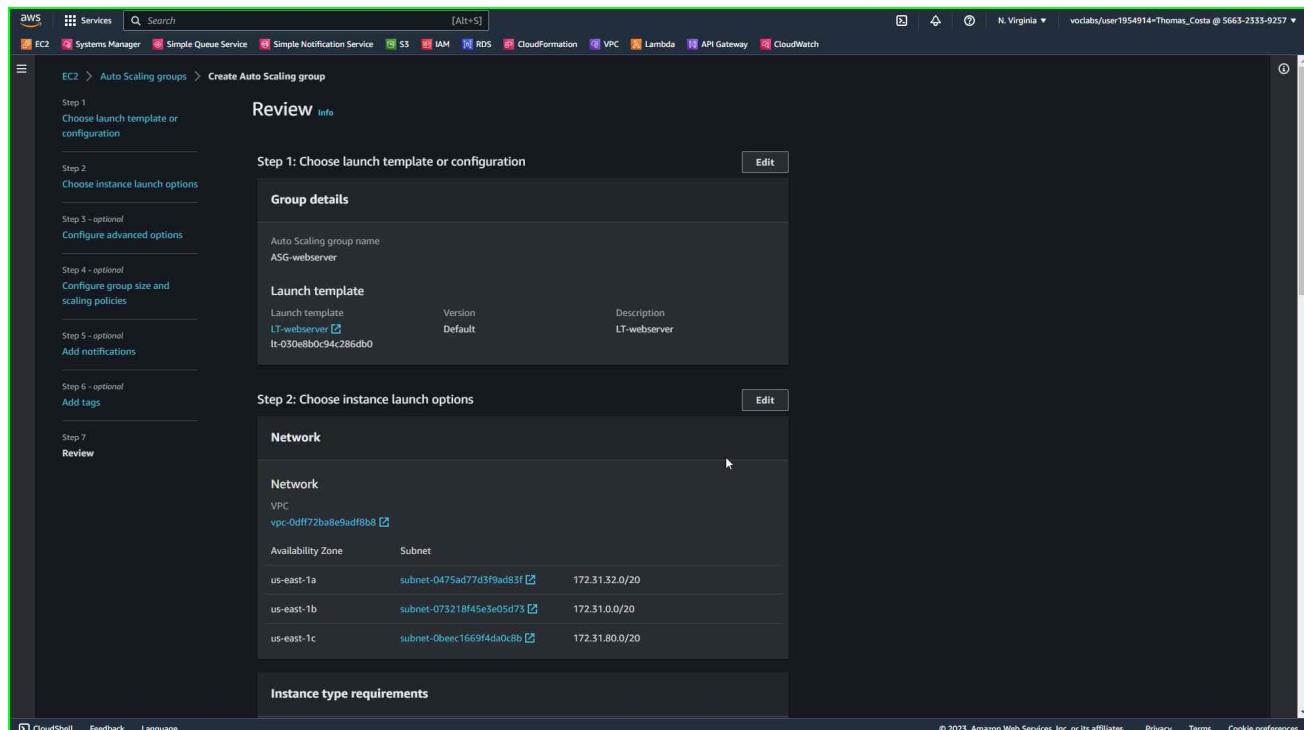
Implantando um Application Load Balancer (ALB) na AWS com EC2

Na próxima tela não tem ação somente clique em Next:



The screenshot shows the 'Add tags - optional' step of the 'Create Auto Scaling group' wizard. The 'Tags (0)' section contains an 'Add tag' button. The 'Next' button at the bottom right is highlighted with a red box.

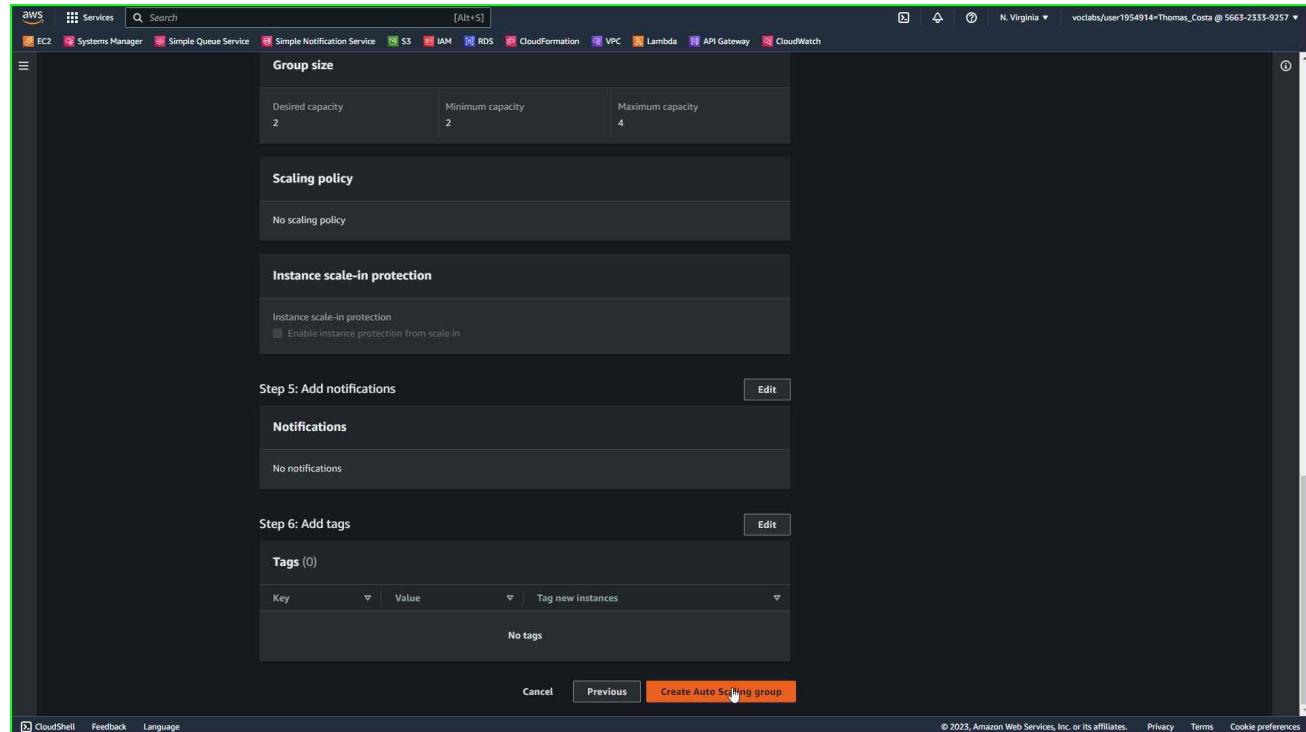
Faça o Review da criação do ASG:



The screenshot shows the 'Review' step of the 'Create Auto Scaling group' wizard. It displays the configuration from the previous steps: Step 1: Choose launch template or configuration (Group details: Auto Scaling group name 'ASG-webserver', Launch template 'LT-webserver' (version Default, description 'LT-webserver')) and Step 2: Choose instance launch options (Network: VPC 'vpc-0dff72b8e9adff8b8', Availability Zone: us-east-1a, us-east-1b, us-east-1c). The 'Next Step' button at the bottom right is highlighted with a red box.

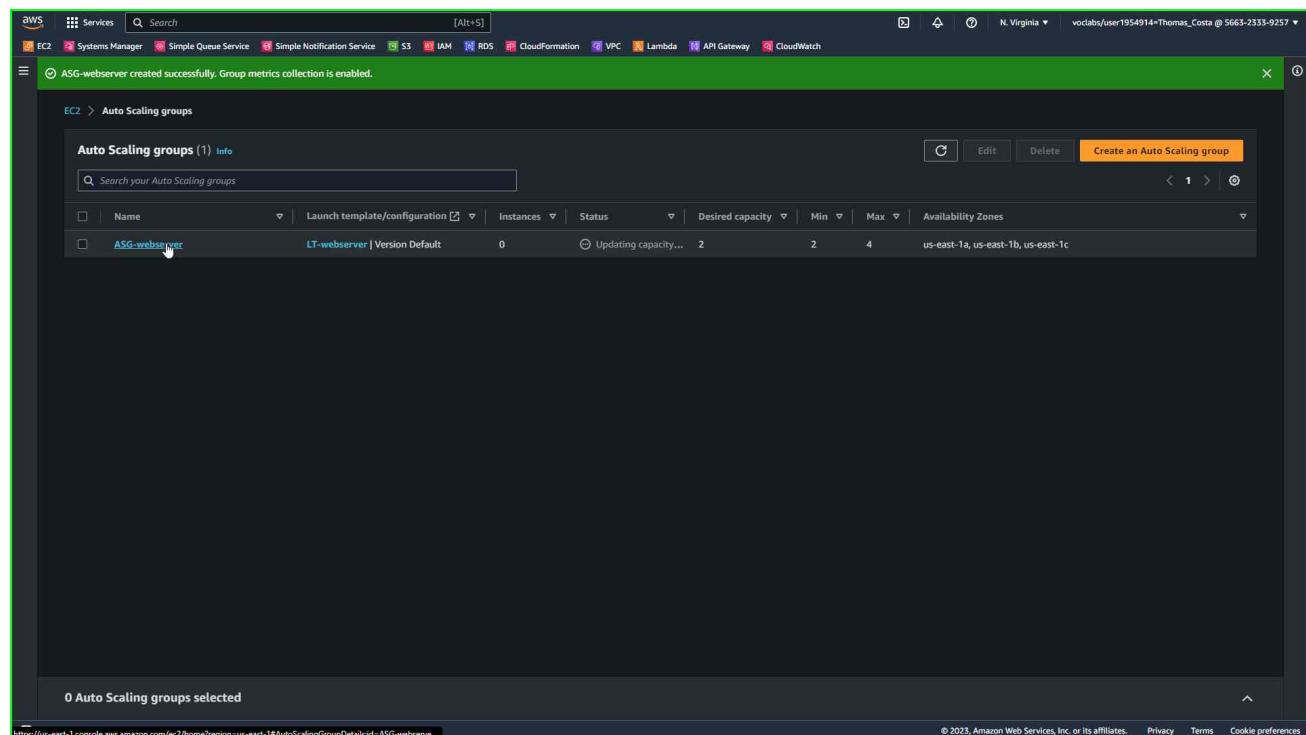
Implantando um Application Load Balancer (ALB) na AWS com EC2

Confirme a criação:



The screenshot shows the final step of creating an Auto Scaling group. The 'Create Auto Scaling group' button is highlighted in orange at the bottom right of the page. The page includes sections for Group size, Scaling policy, Instance scale-in protection, Step 5: Add notifications, and Step 6: Add tags.

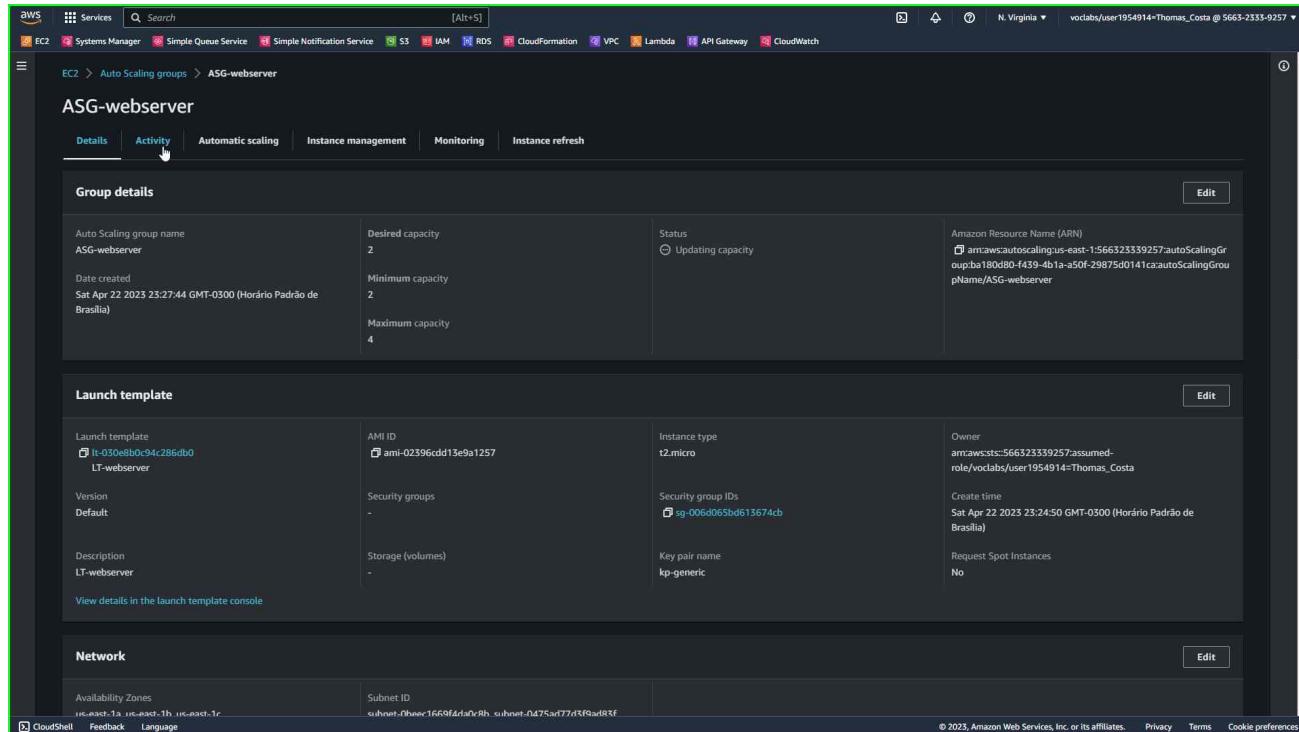
Com ASG criado com sucesso, selecione para ver suas propriedades:



The screenshot shows the 'Auto Scaling groups' page with one group listed: 'ASG-webservice'. The group is in the 'Updating capacity...' state with a desired capacity of 2. It is associated with the 'LT-webserver' launch template and is located in the 'us-east-1a, us-east-1b, us-east-1c' availability zones. The 'Create an Auto Scaling group' button is visible at the top right.

Implantando um Application Load Balancer (ALB) na AWS com EC2

Selecione a aba Activity:



ASG-webserver

Group details

Auto Scaling group name	ASG-webserver	Desired capacity	2
Date created	Sat Apr 22 2023 23:27:44 GMT-0300 (Horário Padrão de Brasília)	Minimum capacity	2
		Maximum capacity	4

Status: Updating capacity

Amazon Resource Name (ARN): arn:aws:autoscaling:us-east-1:566323339257:autoScalingGroup:ba180d80-f439-4b1a-a50f-29875d0141a:autoScalingGroupName/ASG-webserver

Launch template

Launch template	lt-030e8b0c94c286db0	AMI ID	ami-02396cdd13e9a1257	Instance type	t2.micro
Version	LT-webserver	Security groups	-	Security group IDs	sg-006d065bd613674cb
Description	LT-webserver	Storage (volumes)	-	Key pair name	kp-generic

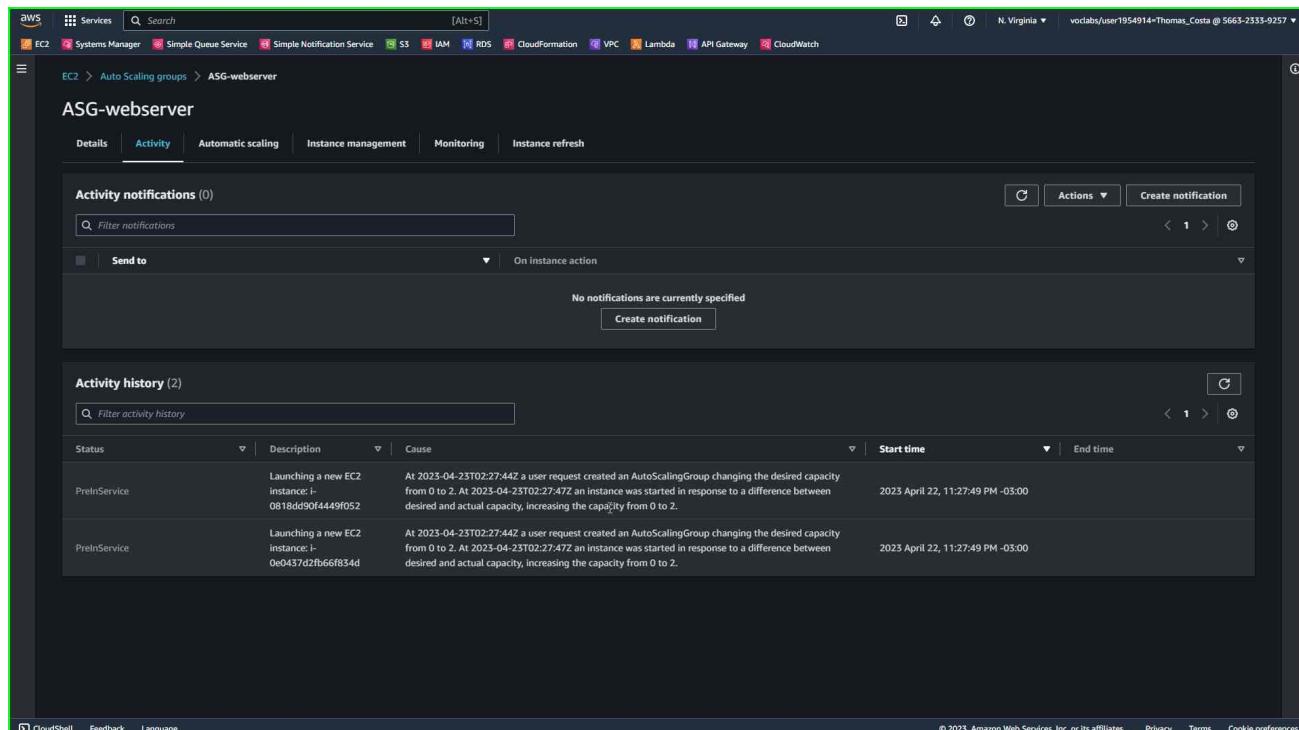
Network

Availability Zones	us-east-1a, us-east-1b, us-east-1c	Subnet ID	subnet-07eef16f0d4da0c8b, subnet-0475ad77ef3f6ad8f6
--------------------	------------------------------------	-----------	---

CloudShell Feedback Language

© 2023, Amazon Web Services, Inc. or its affiliates. Privacy Terms Cookie preferences

Em Activity History podemos ver que a criação dos EC2 foi disparada:



ASG-webserver

Activity notifications (0)

Activity history (2)

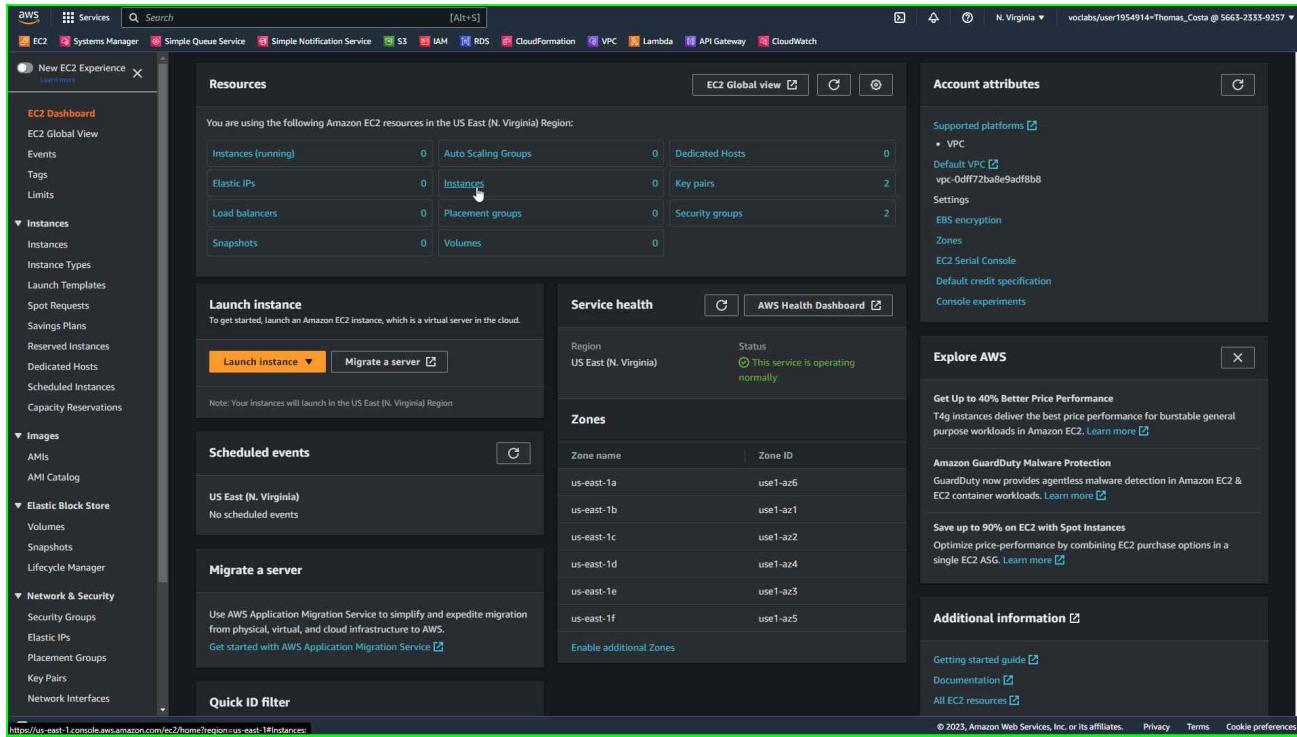
Status	Description	Cause	Start time	End time
PreInService	Launching a new EC2 instance: i-0818dd90f4449f052	At 2023-04-23T02:27:44Z a user request created an AutoScalingGroup changing the desired capacity from 0 to 2. At 2023-04-23T02:27:47Z an instance was started in response to a difference between desired and actual capacity, increasing the capacity from 0 to 2.	2023 April 22, 11:27:49 PM -03:00	
PreInService	Launching a new EC2 instance: i-0e0437d2fb66f834d	At 2023-04-23T02:27:44Z a user request created an AutoScalingGroup changing the desired capacity from 0 to 2. At 2023-04-23T02:27:47Z an instance was started in response to a difference between desired and actual capacity, increasing the capacity from 0 to 2.	2023 April 22, 11:27:49 PM -03:00	

CloudShell Feedback Language

© 2023, Amazon Web Services, Inc. or its affiliates. Privacy Terms Cookie preferences

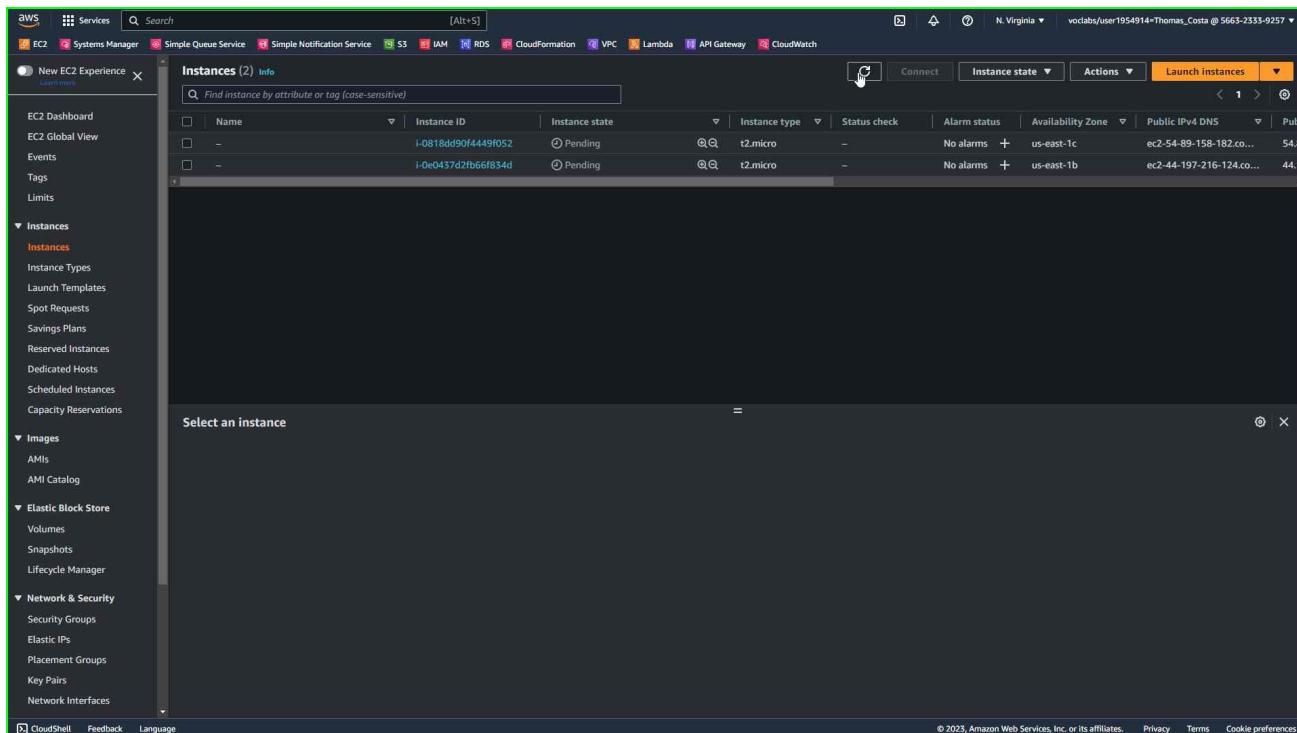
Implantando um Application Load Balancer (ALB) na AWS com EC2

Voltando para a tela principal do EC2 selecionando a opção Instances:



The screenshot shows the AWS EC2 Dashboard. The left sidebar is collapsed, and the main content area is the 'Instances' section. The 'Instances' link in the sidebar is highlighted with a blue box. The main area displays a table of EC2 resources, with the 'Instances' row being the focus, indicated by a mouse cursor. The table includes columns for Instances (running), Auto Scaling Groups, Dedicated Hosts, Elastic IPs, Instances, Key pairs, Load balancers, Placement groups, Security groups, Snapshots, and Volumes. The 'Instances' row shows 0 for all categories except 'Instances' which is 2. The 'Instances' link in the sidebar is also highlighted with a blue box.

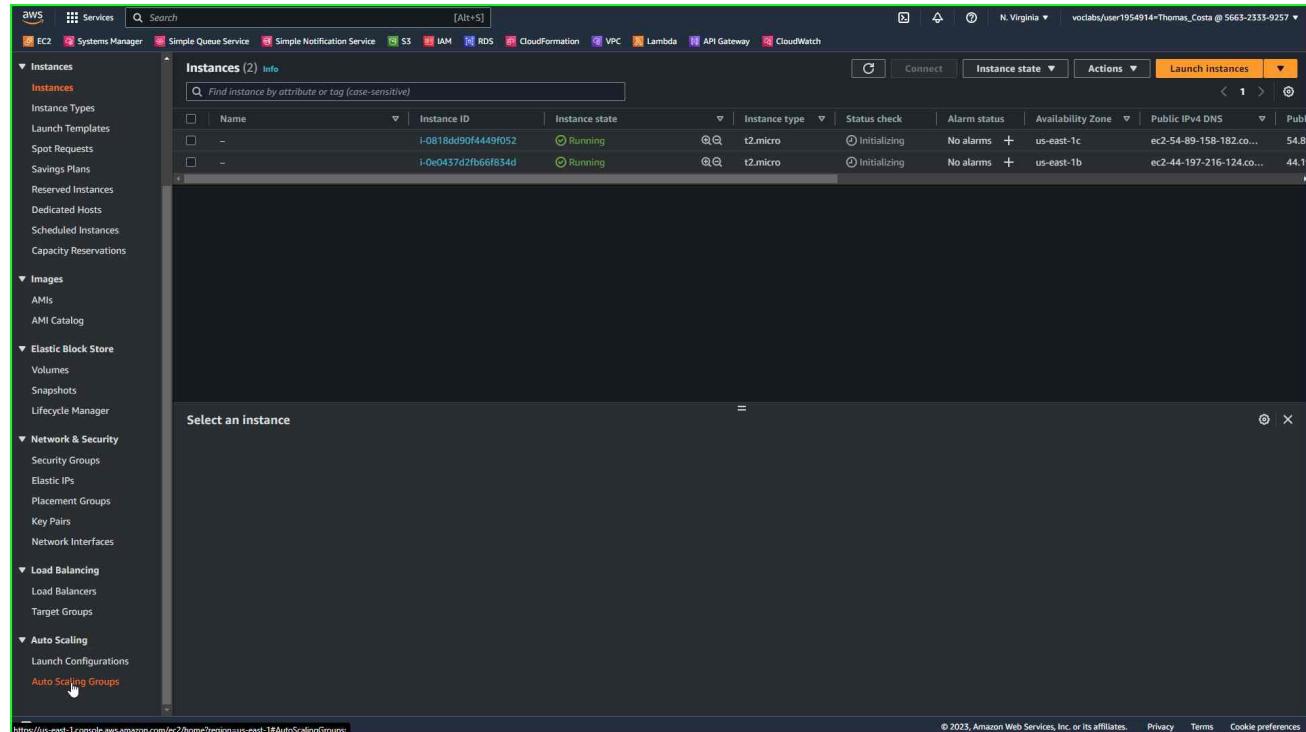
Duas instâncias estão sendo criadas de acordo com os parâmetros do ASG:



The screenshot shows the AWS EC2 Instances list. The left sidebar is collapsed, and the main content area is the 'Instances' list. The 'Instances' link in the sidebar is highlighted with a blue box. The main area displays a table of instances, with the first two rows being the focus, indicated by a mouse cursor. The table includes columns for Name, Instance ID, Instance state, Instance type, Status check, Alarm status, Availability Zone, Public IPv4 DNS, and Public IPv6 DNS. The first instance is 'i-0818ad50f449f052' and the second is 'i-0e457d2fb66f834d', both in the 'Pending' state. The 'Instances' link in the sidebar is also highlighted with a blue box.

Implantando um Application Load Balancer (ALB) na AWS com EC2

Instâncias criadas com sucesso:

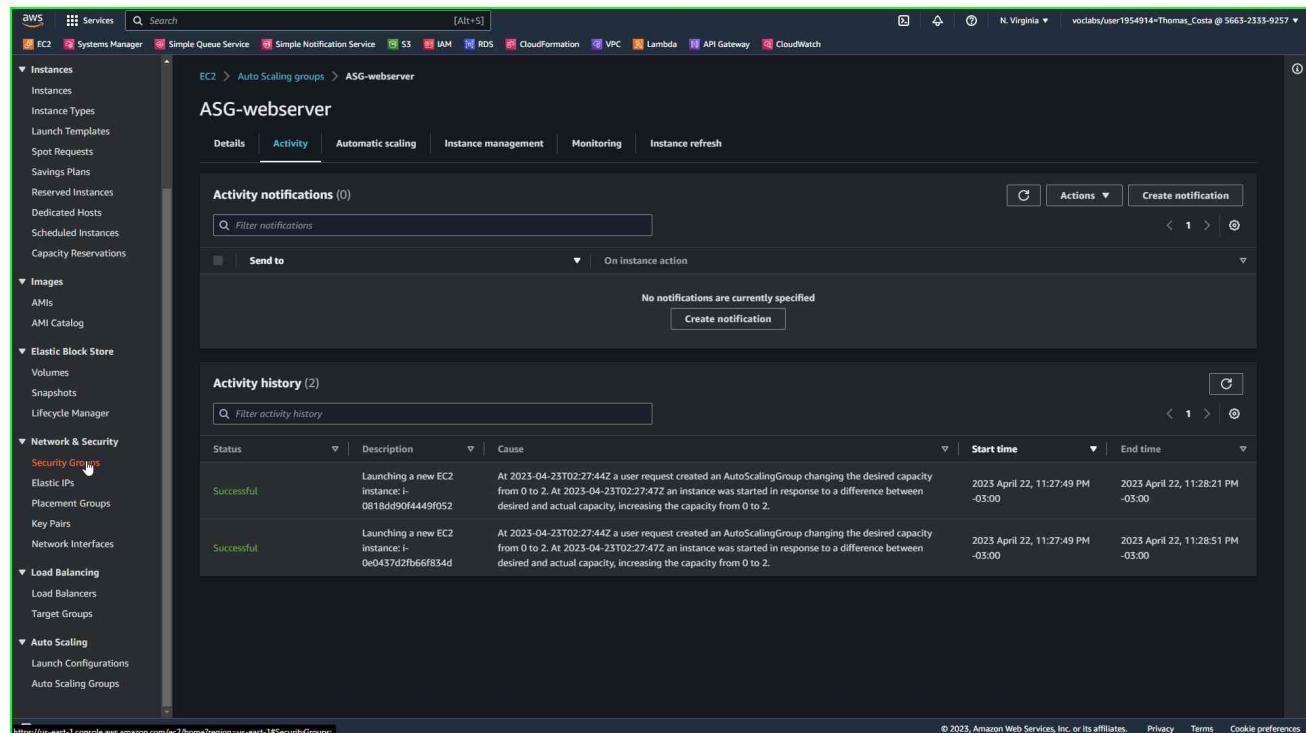


The screenshot shows the AWS EC2 Instances page. The left sidebar is collapsed. The main table displays two instances:

Name	Instance ID	Instance state	Instance type	Status check	Alarm status	Availability Zone	Public IPv4 DNS	Publ...
-	i-0818dd90f4449f052	Running	t2.micro	OK	Initializing	No alarms	+ us-east-1c	ec2-54-89-158-182.co...
-	i-0e0437d2fb66f834d	Running	t2.micro	OK	Initializing	No alarms	+ us-east-1b	ec2-44-197-216-124.co...

A modal window titled "Select an instance" is open in the foreground, listing the two instances.

Em Activity History do ASG foi notificado que as duas instâncias foram criadas:

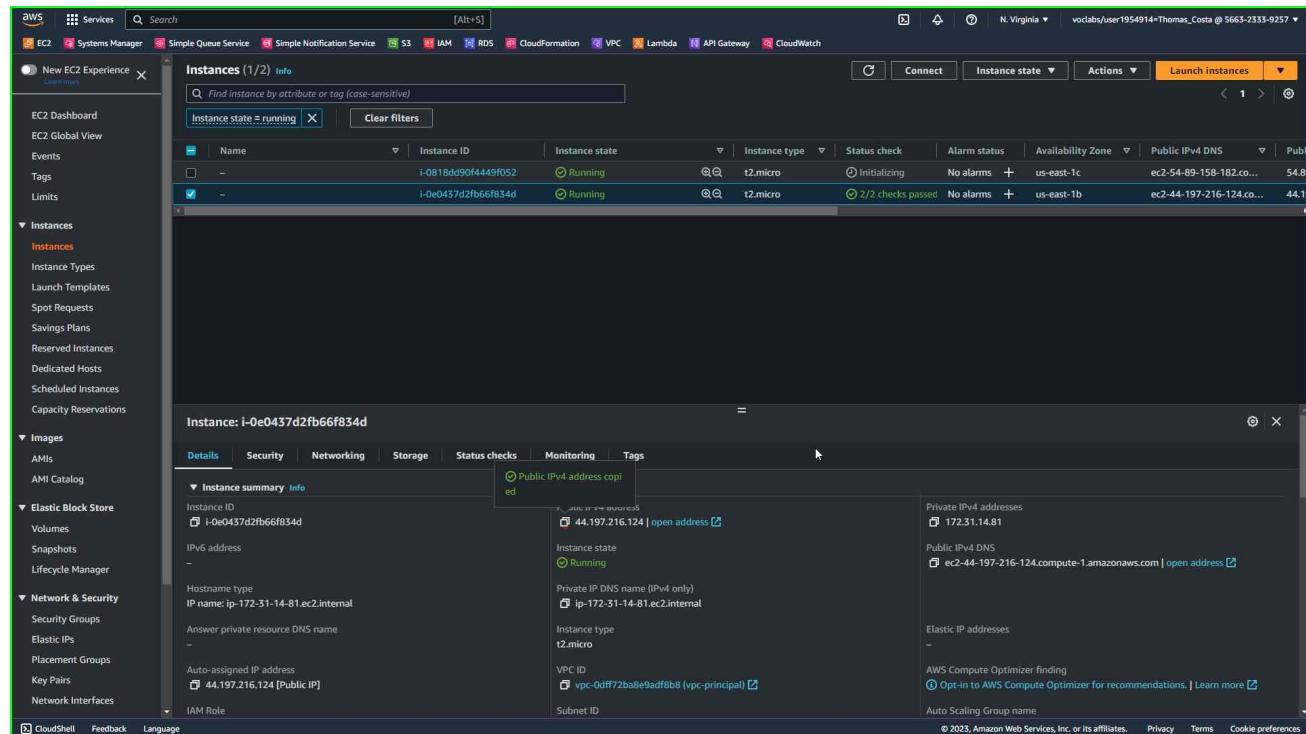


The screenshot shows the AWS Auto Scaling Groups Activity History page for the "ASG-webserver" group. The left sidebar is collapsed. The "Activity" tab is selected. The "Activity history" table shows two entries:

Status	Description	Cause	Start time	End time
Successful	Launching a new EC2 instance: i-0818dd90f4449f052	At 2023-04-23T02:27:44Z a user request created an AutoScalingGroup changing the desired capacity from 0 to 2. At 2023-04-23T02:27:47Z an instance was started in response to a difference between desired and actual capacity, increasing the capacity from 0 to 2.	2023 April 22, 11:27:49 PM -03:00	2023 April 22, 11:28:21 PM -03:00
Successful	Launching a new EC2 instance: i-0e0437d2fb66f834d	At 2023-04-23T02:27:44Z a user request created an AutoScalingGroup changing the desired capacity from 0 to 2. At 2023-04-23T02:27:47Z an instance was started in response to a difference between desired and actual capacity, increasing the capacity from 0 to 2.	2023 April 22, 11:27:49 PM -03:00	2023 April 22, 11:28:51 PM -03:00

Implantando um Application Load Balancer (ALB) na AWS com EC2

Selecione o endereço público de uma das instâncias para acessar o servidor Apache HTTP:



Instances (1/2) Info

Find instance by attribute or tag (case-sensitive)

Instance state = running

Clear filters

Name	Instance ID	Instance state	Instance type	Status check	Alarm status	Availability Zone	Public IPv4 DNS	Public IPv6 DNS
-	i-0818dd90f4449f052	Running	t2.micro	Initializing	No alarms	us-east-1c	ec2-54-89-158-182.co...	54.89.158.182
<input checked="" type="checkbox"/>	i-0e0437d2fb66f834d	Running	t2.micro	2/2 checks passed	No alarms	us-east-1b	ec2-44-197-216-124.co...	44.197.216.124

Instance: i-0e0437d2fb66f834d

Details Security Networking Storage Status checks Monitoring Tags

Instance summary Info

Public IPv4 address copied

Instance ID: i-0e0437d2fb66f834d

IPv6 address:

Hostname type: IP name: ip-172-31-14-81.ec2.internal

Answers private resource DNS name:

Auto-assigned IP address: 44.197.216.124 [Public IP]

IAM Role:

Private IP DNS name (IPv4 only): ip-172-31-14-81.ec2.internal

Instance type: t2.micro

VPC ID: vpc-0dff72ba8e9adff8b (vpc-principal)

Subnet ID:

Private IPv4 addresses: 172.31.14.81

Public IPv4 DNS: ec2-44-197-216-124.compute-1.amazonaws.com [open address]

Elastic IP addresses:

AWS Compute Optimizer finding: Opt-in to AWS Compute Optimizer for recommendations. | Learn more

Auto Scaling Group name:

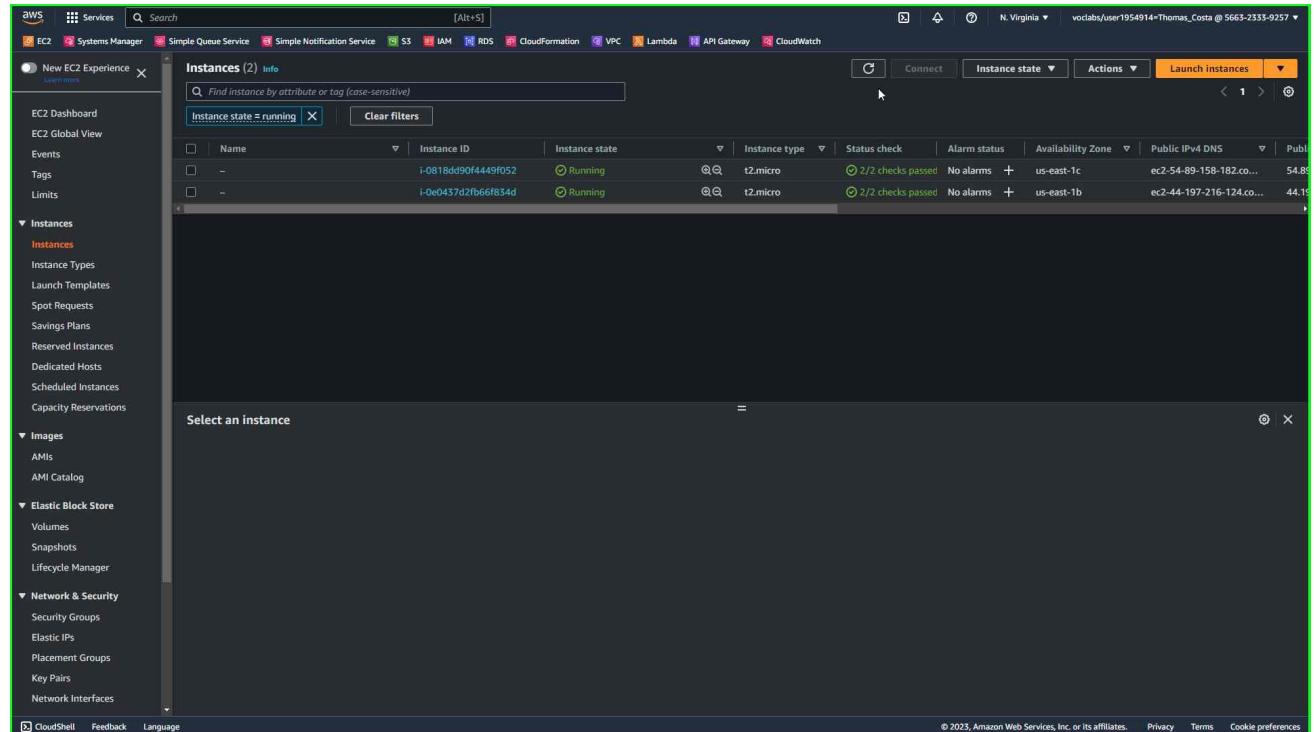
Resposta do servidor Apache HTTP no navegador:

Zona de Disponibilidade e IPs

Zona de Disponibilidade: us-east-1b
IP Público: 44.197.216.124
IP Privado: 172.31.14.81

Implantando um Application Load Balancer (ALB) na AWS com EC2

Instâncias criadas com sucesso e Status check em ordem:



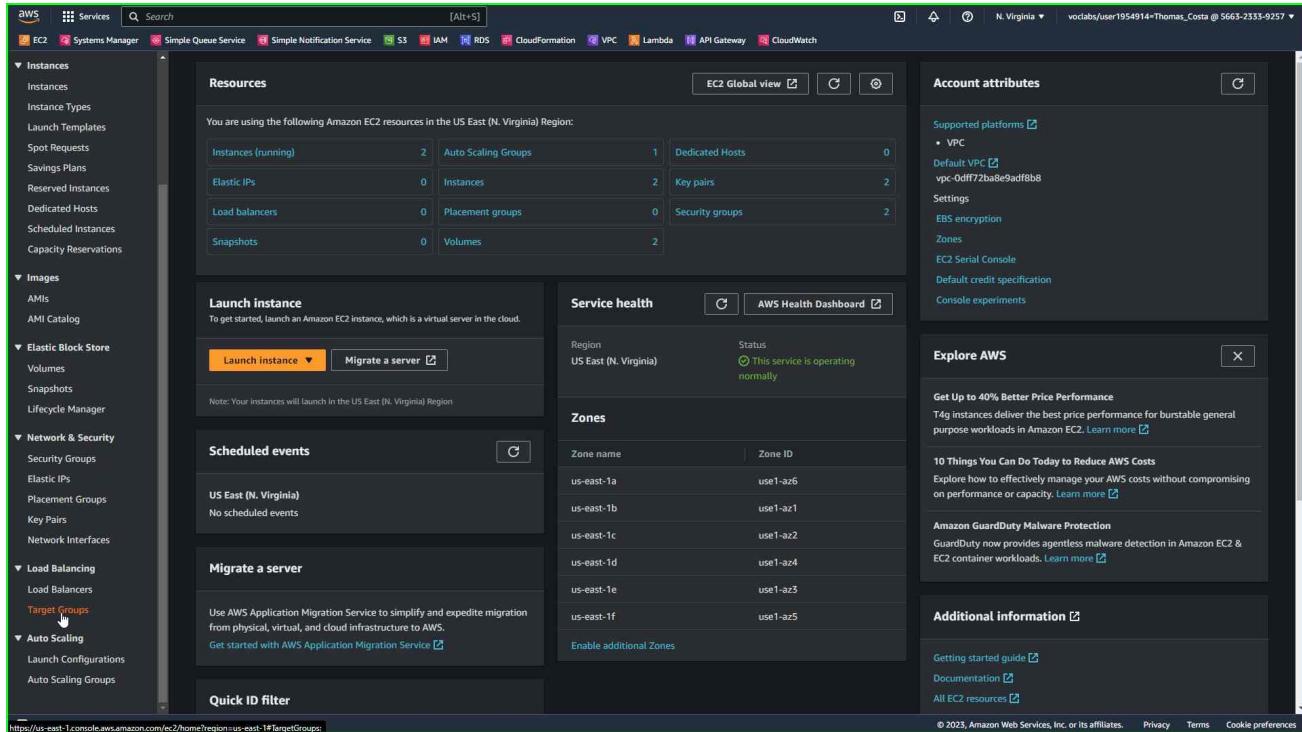
The screenshot shows the AWS EC2 Instances page with two running instances listed. The table includes columns for Name, Instance ID, Instance state, Instance type, Status check, Alarm status, Availability Zone, Public IPv4 DNS, and Public IPv6 DNS. Both instances are in the 'Running' state, t2.micro type, and have 2/2 checks passed. They are located in us-east-1c and us-east-1b respectively, with their respective public IP addresses and port numbers.

Name	Instance ID	Instance state	Instance type	Status check	Alarm status	Availability Zone	Public IPv4 DNS	Public IPv6 DNS
-	i-0818dd90f4449f052	Running	t2.micro	2/2 checks passed	No alarms	us-east-1c	ec2-54-89-158-182.co...	54.89.158.182:80
-	i-0e0437d2fb66f834d	Running	t2.micro	2/2 checks passed	No alarms	us-east-1b	ec2-44-197-216-124.co...	44.197.216.124:80

Implantando um Application Load Balancer (ALB) na AWS com EC2

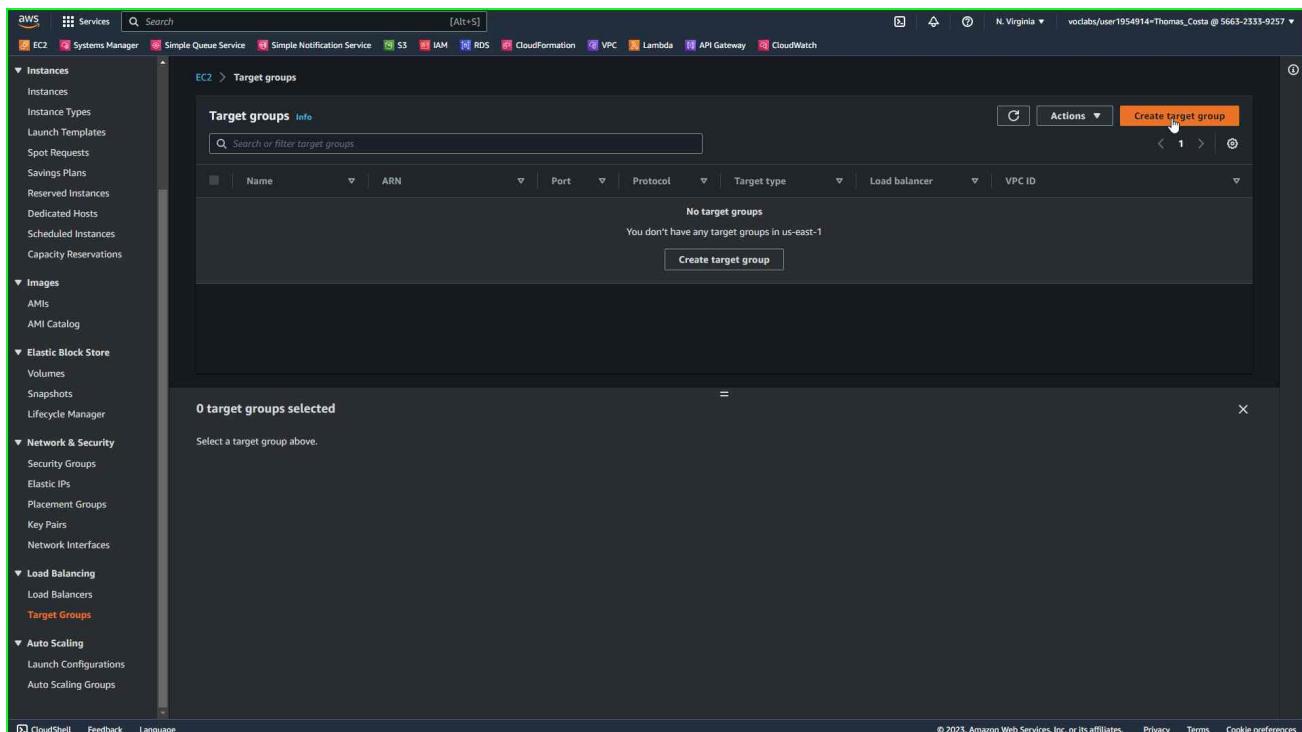
Parte 3 – Criando o Target Group

Selecionar a opção Target Groups:



The screenshot shows the AWS EC2 home page in the N. Virginia region. The left sidebar is expanded to show the 'Load Balancing' section, with 'Target Groups' highlighted. The main content area displays various EC2 resources: 2 instances (running), 1 auto scaling group, 0 dedicated hosts, 0 elastic IPs, 2 instances, 2 key pairs, 0 load balancers, 0 placement groups, 0 security groups, 0 snapshots, and 2 volumes. Below this, there are sections for 'Launch instance', 'Scheduled events', 'Migrate a server', and a 'Quick ID filter'. On the right, there is an 'Account attributes' sidebar with options like 'Supported platforms', 'Default VPC', and 'Settings', and a 'Explore AWS' sidebar with links to price performance, cost reduction, and GuardDuty. The URL in the address bar is <https://us-east-1.console.aws.amazon.com/ec2/home?region=us-east-1#TargetGroups>.

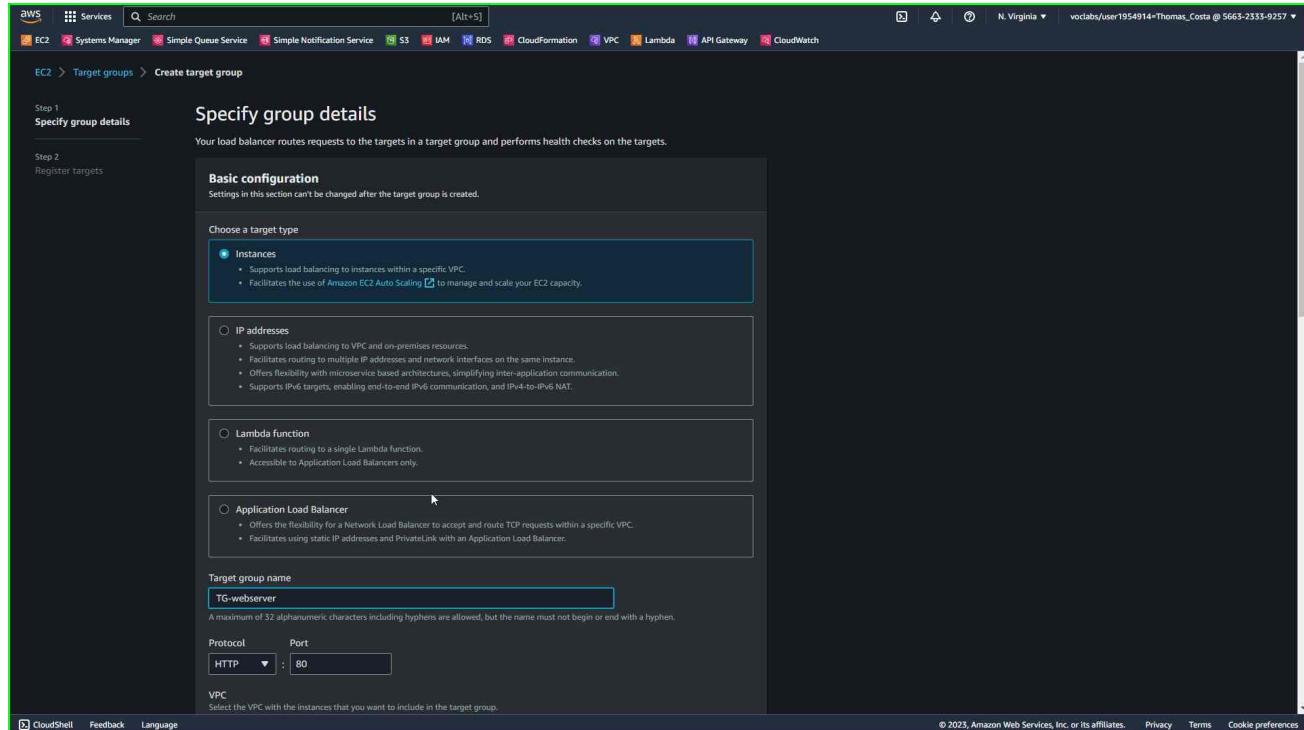
Na tela principal, clicar no botão Create target group:



The screenshot shows the 'Target groups' page in the N. Virginia region. The left sidebar is expanded to show the 'Load Balancing' section, with 'Target Groups' highlighted. The main content area shows a table with 0 target groups selected. A modal window is open at the bottom, prompting the user to 'Select a target group above.' The 'Create target group' button is visible in the top right corner of the main table area. The URL in the address bar is <https://us-east-1.console.aws.amazon.com/ec2/target-groups?region=us-east-1>.

Implantando um Application Load Balancer (ALB) na AWS com EC2

Colocar o nome do Target group de **TG-webserver**:



Step 1
Specify group details

Your load balancer routes requests to the targets in a target group and performs health checks on the targets.

Step 2
Register targets

Basic configuration
Settings in this section can't be changed after the target group is created.

Choose a target type

Instances

- Supports load balancing to instances within a specific VPC.
- Facilitates the use of Amazon EC2 Auto Scaling  to manage and scale your EC2 capacity.

IP addresses

- Supports load balancing to VPC and on-premises resources.
- Facilitates routing to multiple IP addresses and network interfaces on the same instance.
- Offers flexibility with microservice based architectures, simplifying inter-application communication.
- Supports IPv6 targets, enabling end-to-end IPv6 communication, and IPv4-to-IPv6 NAT.

Lambda function

- Facilitates routing to a single Lambda function.
- Accessible to Application Load Balancers only.

Application Load Balancer

- Offers the flexibility for a Network Load Balancer to accept and route TCP requests within a specific VPC.
- Facilitates using static IP addresses and PrivateLink with an Application Load Balancer.

Target group name

TG-webserver

A maximum of 32 alphanumeric characters including hyphens are allowed, but the name must not begin or end with a hyphen.

Protocol

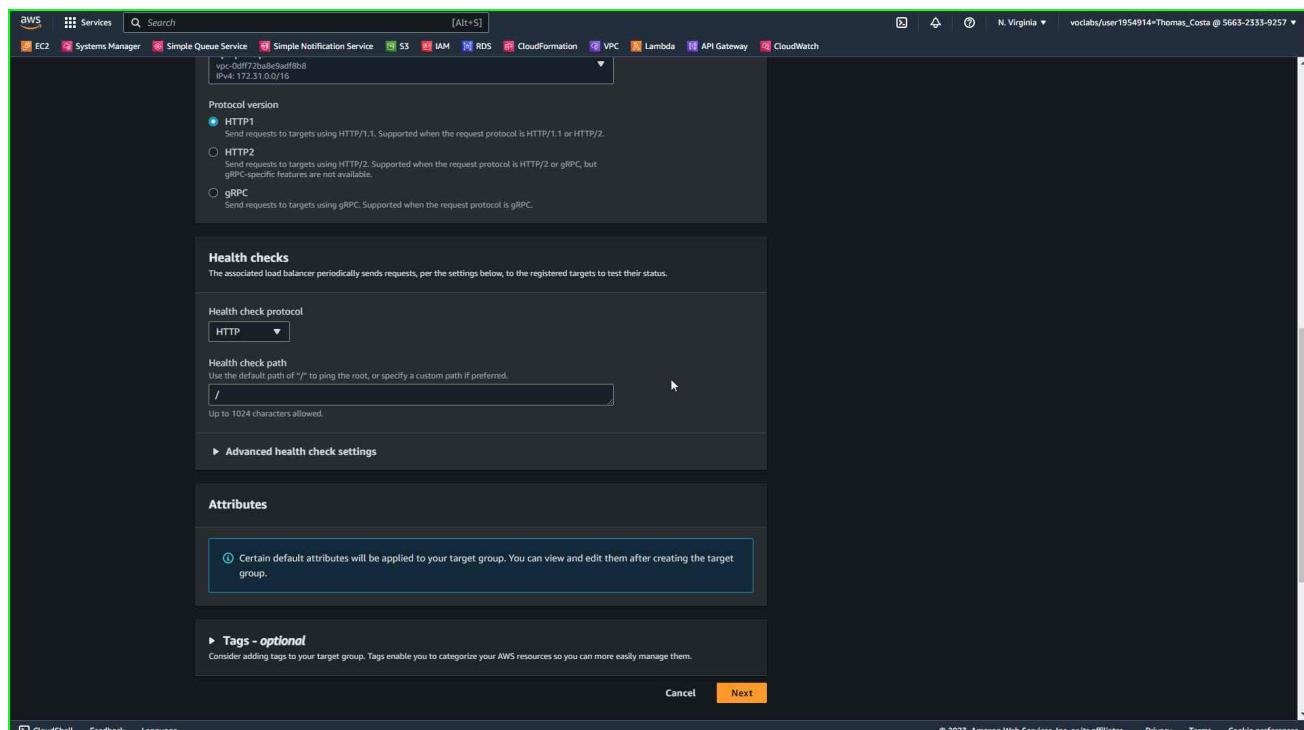
Port

HTTP : 80

VPC

Select the VPC with the instances that you want to include in the target group.

Selecionar as opções conforme imagem abaixo:



Protocol version

HTTP1

Send requests to targets using HTTP/1.1. Supported when the request protocol is HTTP/1.1 or HTTP/2.

HTTP2

Send requests to targets using HTTP/2. Supported when the request protocol is HTTP/2 or gRPC, but gRPC-specific features are not available.

gRPC

Send requests to targets using gRPC. Supported when the request protocol is gRPC.

Health checks

The associated load balancer periodically sends requests, per the settings below, to the registered targets to test their status.

Health check protocol

HTTP

Health check path

Use the default path of "/" to ping the root, or specify a custom path if preferred.

/

Up to 1024 characters allowed.

Advanced health check settings

Attributes

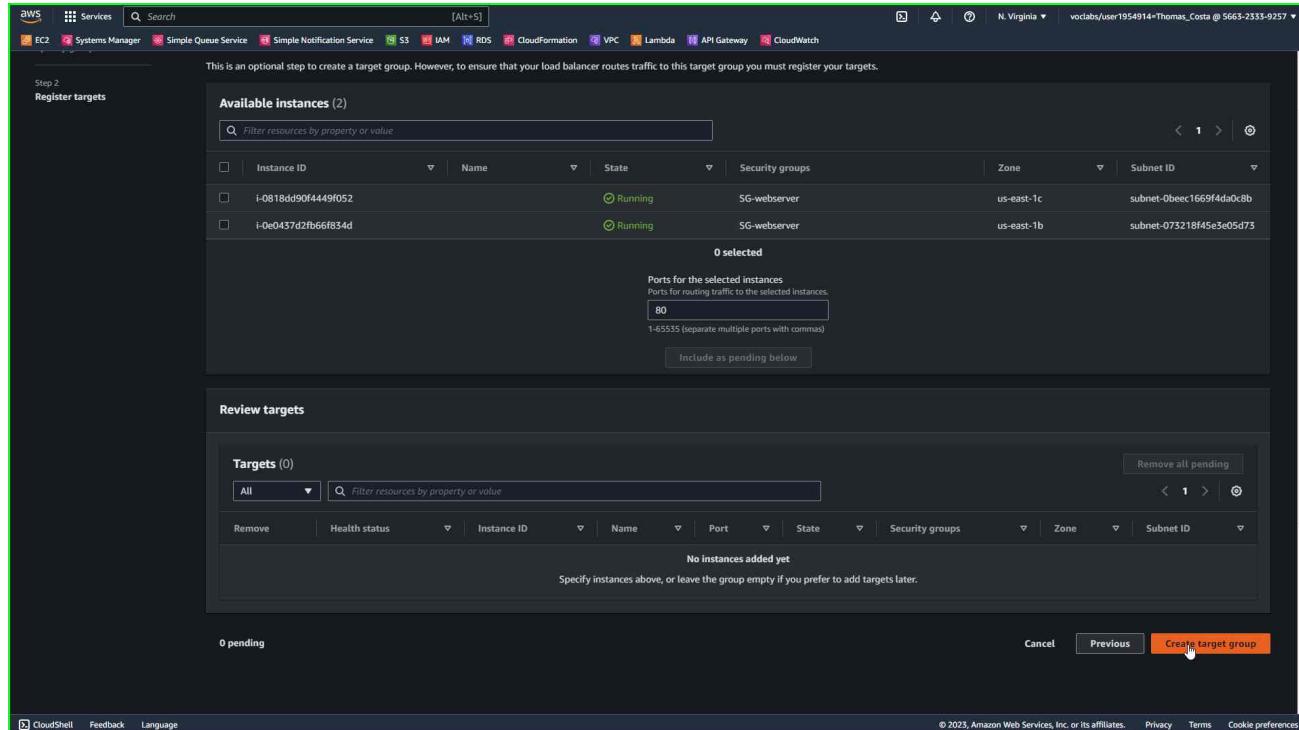
 Certain default attributes will be applied to your target group. You can view and edit them after creating the target group.

Tags - optional

Consider adding tags to your target group. Tags enable you to categorize your AWS resources so you can more easily manage them.

Implantando um Application Load Balancer (ALB) na AWS com EC2

Nesta tela clique somente em Create target group:



This is an optional step to create a target group. However, to ensure that your load balancer routes traffic to this target group you must register your targets.

Available instances (2)

Instance ID	Name	State	Security groups	Zone	Subnet ID
i-0818dd90f4449f052	SG-webserver	Running	us-east-1c	subnet-0beec1669f4da0c8b	
i-0e0437d2fb66f834d	SG-webserver	Running	us-east-1b	subnet-073218f45e3e05d73	

Ports for the selected instances
Ports for routing traffic to the selected instances.
80
1-65535 (separate multiple ports with commas)

Review targets

Targets (0)

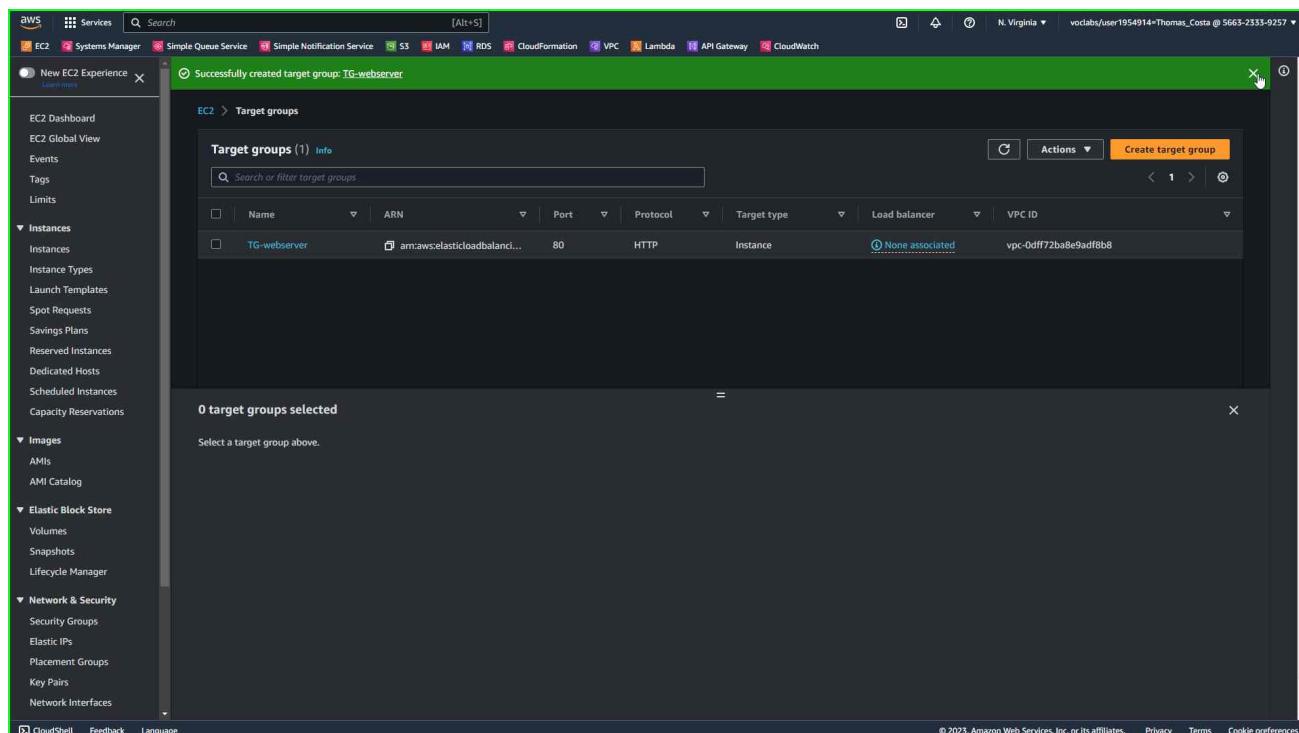
Instance ID	Name	Port	State	Security groups	Zone	Subnet ID
No instances added yet						

Specify instances above, or leave the group empty if you prefer to add targets later.

0 pending

Cancel Previous **Create target group**

Target group foi criado com sucesso:



Target groups (1) Info

Name	ARN	Port	Protocol	Target type	Load balancer	VPC ID
TG-webserver	arn:aws:elasticloadbalancing:us-east-1:566323339257:targetgroup/TG-webserver/0dff72ba8e9adfb8b8	80	HTTP	Instance	None associated	vpc-0dff72ba8e9adfb8b8

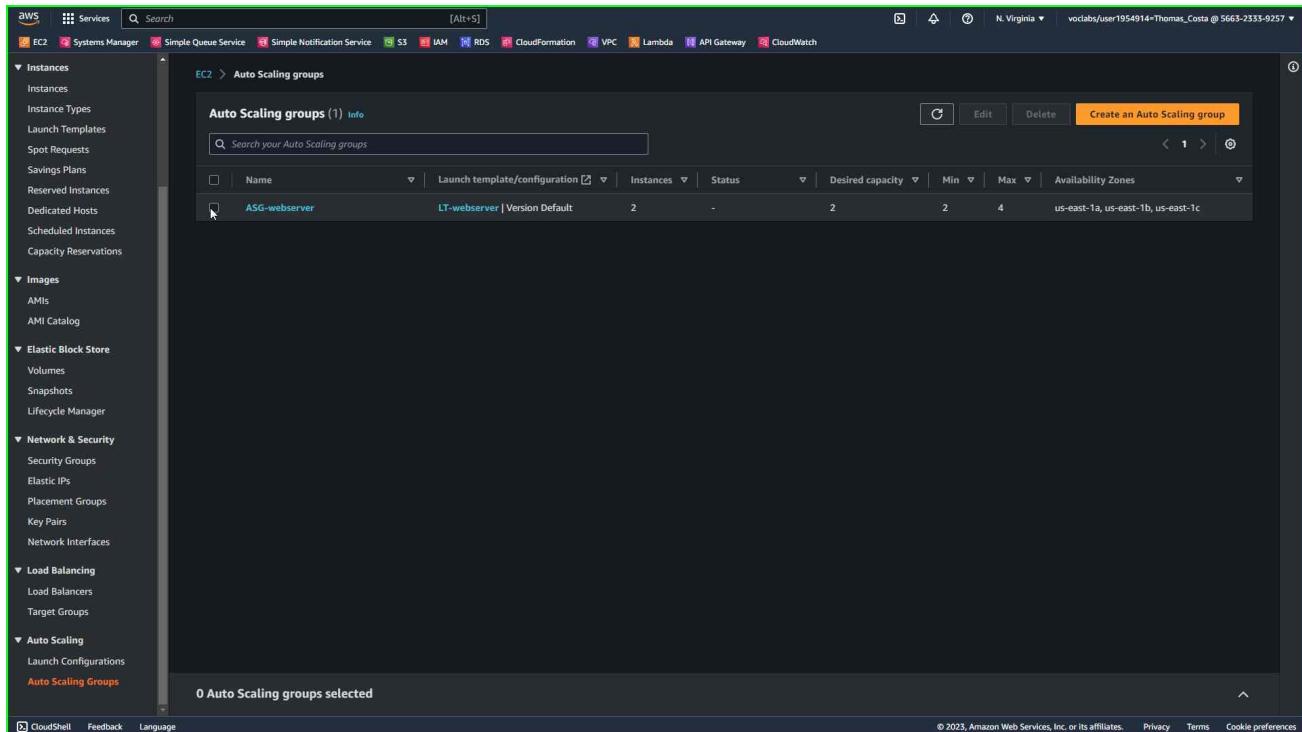
0 target groups selected

Select a target group above.

CloudShell Feedback Language © 2023, Amazon Web Services, Inc. or its affiliates. Privacy Terms Cookie preferences

Implantando um Application Load Balancer (ALB) na AWS com EC2

Vamos voltar para a tela do Auto Scaling groups selecionar o ASG-webserver:

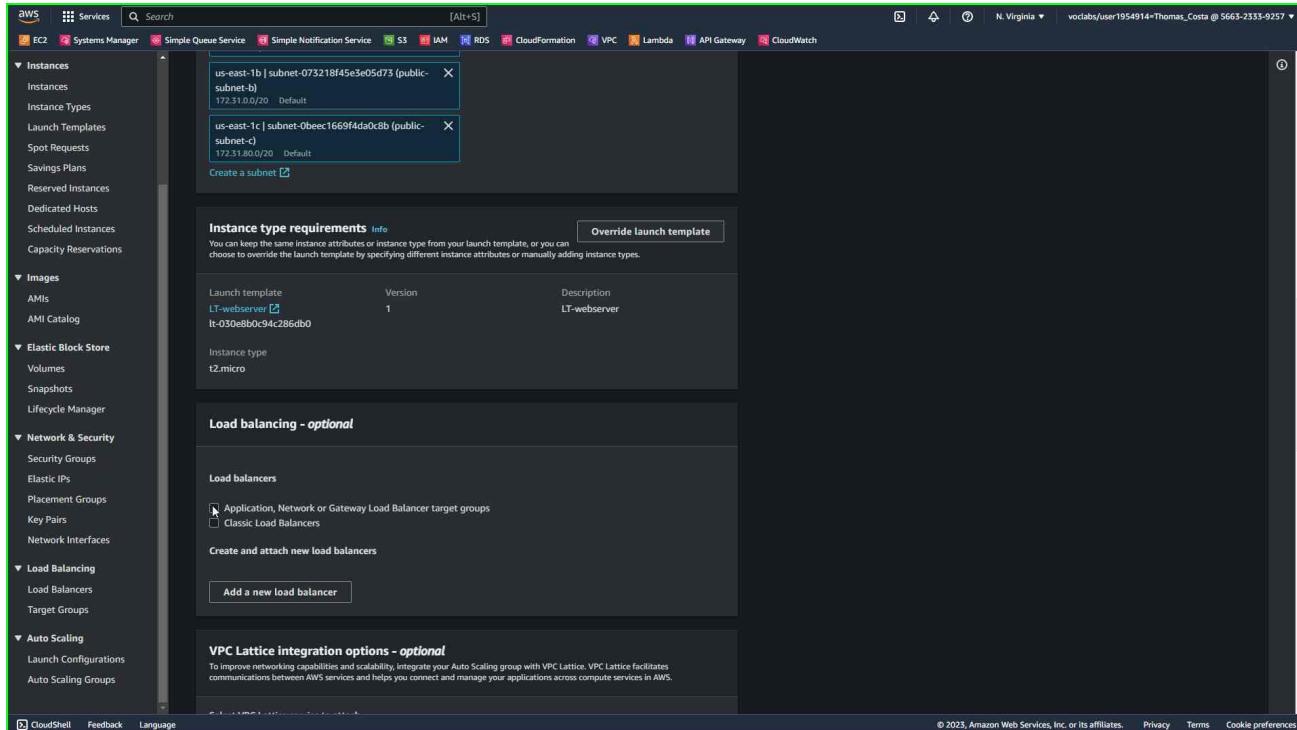


The screenshot shows the AWS Management Console with the EC2 service selected. The left sidebar is expanded to show the 'Auto Scaling' section, with 'Auto Scaling Groups' selected. The main content area displays a table titled 'Auto Scaling groups (1)'. The table has one row for 'ASG-webserver', which is associated with the 'LT-webserver' launch template and is set to a desired capacity of 2. The table includes columns for Name, Launch template/configuration, Instances, Status, Desired capacity, Min, Max, and Availability Zones (us-east-1a, us-east-1b, us-east-1c). The bottom of the page shows a message '0 Auto Scaling groups selected'.

Name	Launch template/configuration	Instances	Status	Desired capacity	Min	Max	Availability Zones
ASG-webserver	LT-webserver Version Default	2	-	2	2	4	us-east-1a, us-east-1b, us-east-1c

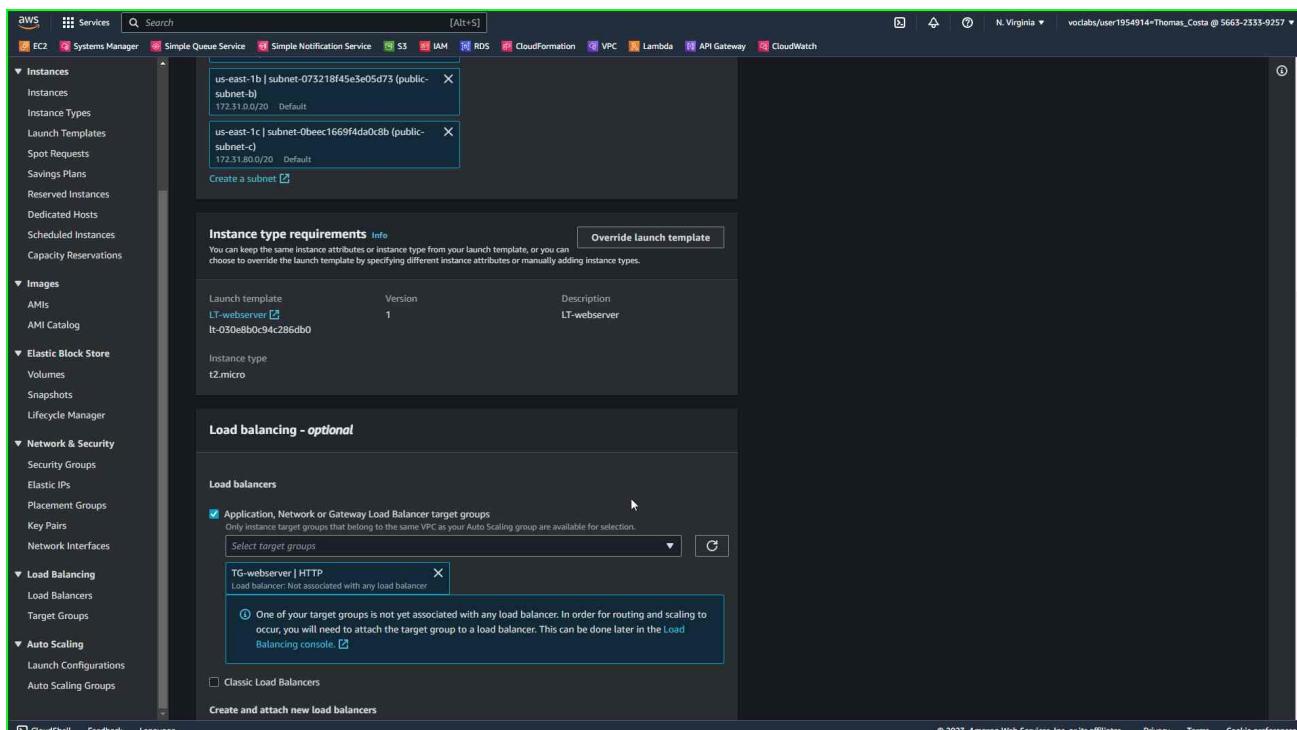
Implantando um Application Load Balancer (ALB) na AWS com EC2

Selecionar a opção Application, Network or Gateway Load Balancer target groups:



The screenshot shows the AWS EC2 Instances configuration page. In the 'Load balancing - optional' section, the checkbox for 'Application, Network or Gateway Load Balancer target groups' is checked. Other options like 'Classic Load Balancers' are available but not selected. The page also shows a list of subnets and instance type requirements.

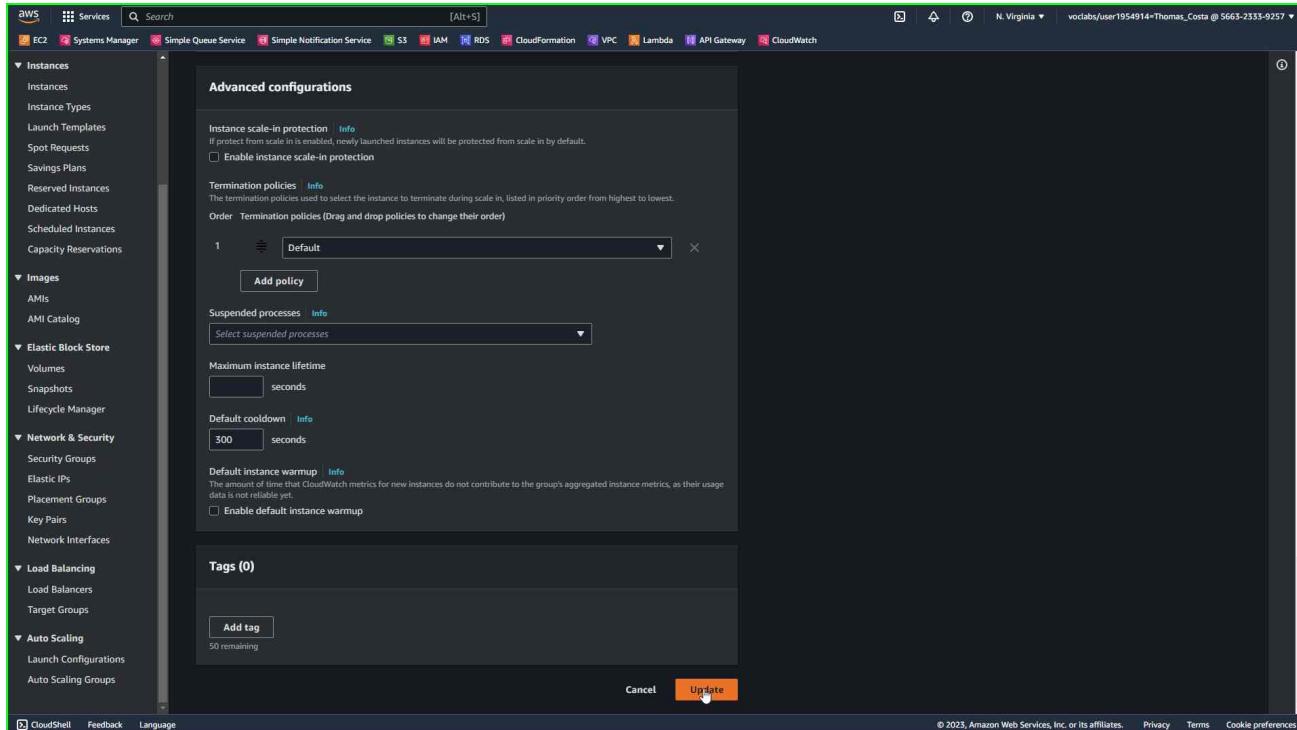
Selecionar o Target group TG-webserver:



The screenshot shows the AWS EC2 Instances configuration page. In the 'Select target groups' dropdown, 'TG-webserver | HTTP' is listed. A note at the bottom of the dropdown indicates that the target group is not yet associated with any load balancer and needs to be attached later. Other target groups like 'Classic Load Balancers' are also listed but not selected.

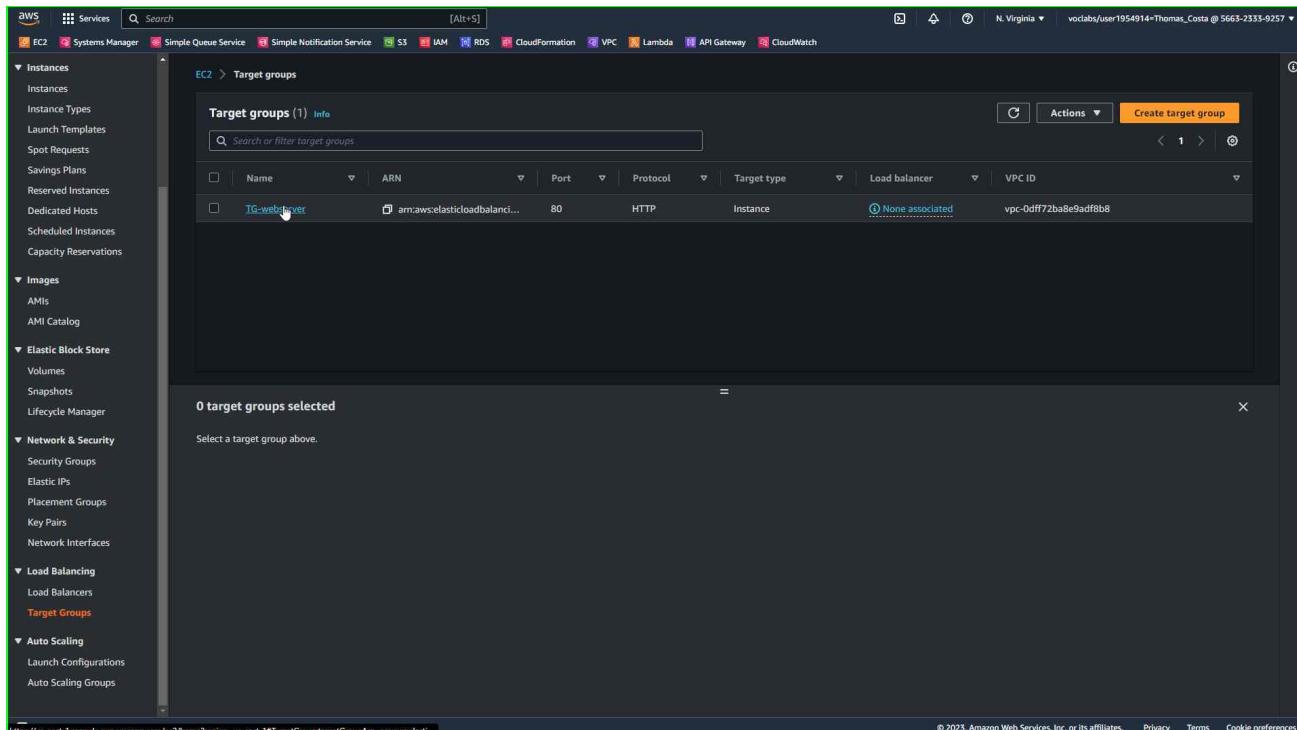
Implantando um Application Load Balancer (ALB) na AWS com EC2

Clique em Update:



The screenshot shows the AWS EC2 Instances configuration page. The left sidebar is expanded to show various EC2-related services. The main panel is titled 'Advanced configurations' and contains several sections: 'Instance scale-in protection' (with an info link and a checkbox for 'Enable instance scale-in protection'), 'Termination policies' (with an info link and a dropdown menu showing 'Default'), 'Suspended processes' (with an info link and a dropdown menu showing 'Select suspended processes'), 'Default cooldown' (with an info link and a dropdown menu showing '300 seconds'), and 'Default instance warmup' (with an info link and a checkbox for 'Enable default instance warmup'). At the bottom of the panel, there is a 'Tags (0)' section with an 'Add tag' button and a note that '50 remaining'. At the very bottom right of the panel, there is an 'Update' button highlighted with a red box.

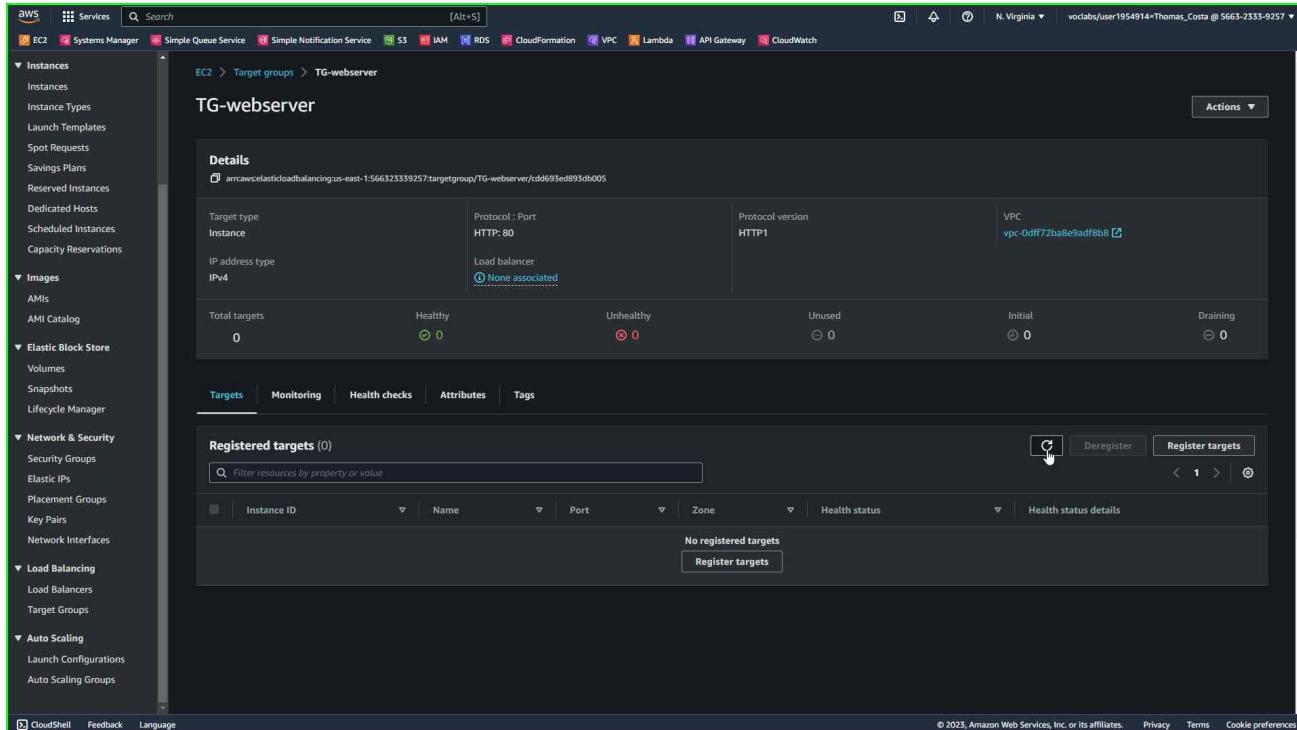
Volte a tela do Target groups e selecione o TG-webserver:



The screenshot shows the AWS EC2 Target groups page. The left sidebar is expanded to show various EC2-related services, with 'Target Groups' selected. The main panel is titled 'Target groups (1) Info' and shows a table with one row: 'TG-webserver' (Name), 'arn:aws:elasticloadbalanc...', '80' (Port), 'HTTP' (Protocol), 'Instance' (Target type), 'None associated' (Load balancer), and 'vpc-0dff72ba8e9adff8b8' (VPC ID). At the bottom of the panel, there is a note '0 target groups selected' and a message 'Select a target group above.' At the very bottom right of the panel, there is an 'Update' button highlighted with a red box.

Implantando um Application Load Balancer (ALB) na AWS com EC2

Clique no botão mostrado na imagem abaixo:

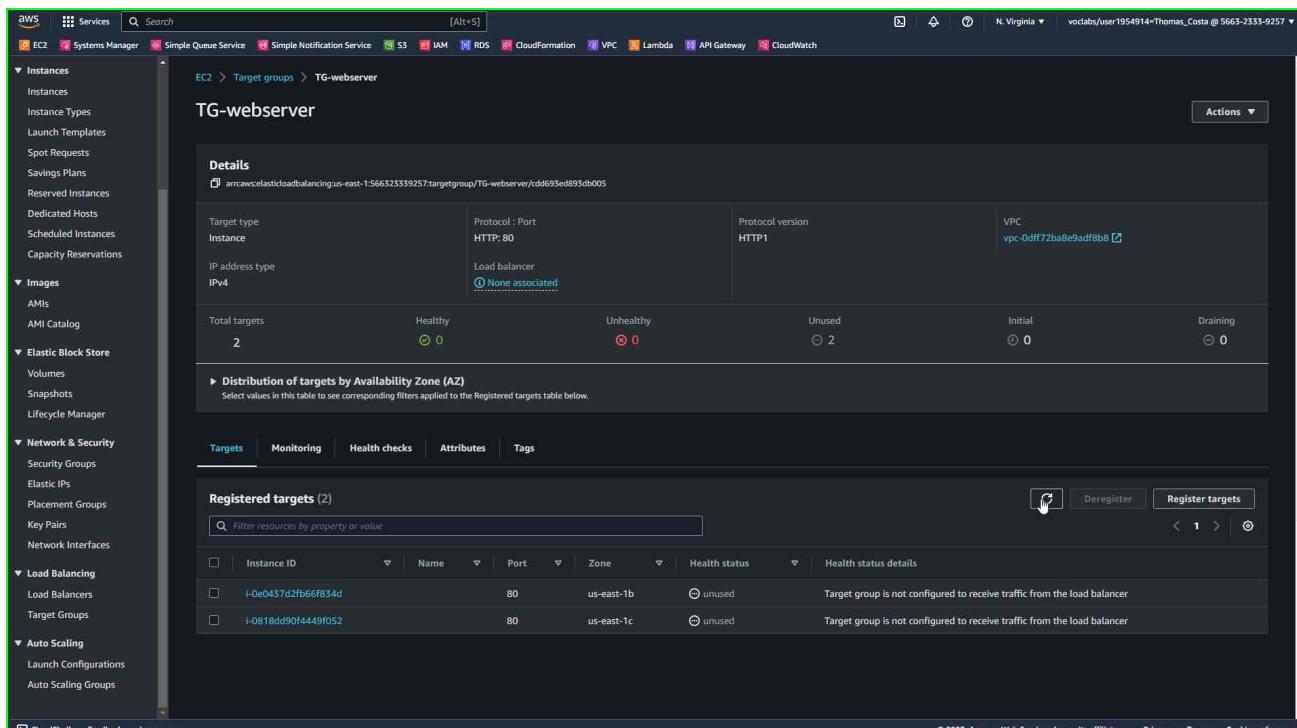


The screenshot shows the AWS EC2 Target Groups console. The left sidebar is collapsed. The main area shows the 'TG-webserver' target group. The 'Targets' tab is selected, displaying the following table:

Total targets	Healthy	Unhealthy	Unused	Initial	Draining
0	0	0	0	0	0

Below the table, there are tabs for 'Monitoring', 'Health checks', 'Attributes', and 'Tags'. The 'Health status' column shows 'Unused' for all targets. The 'Actions' dropdown menu is open, and the 'Register targets' button is highlighted with a green box.

As instâncias EC2 estão atribuídas corretamente para o Target group mas com Health Status em unused:



The screenshot shows the AWS EC2 Target Groups console. The left sidebar is collapsed. The main area shows the 'TG-webserver' target group. The 'Targets' tab is selected, displaying the following table:

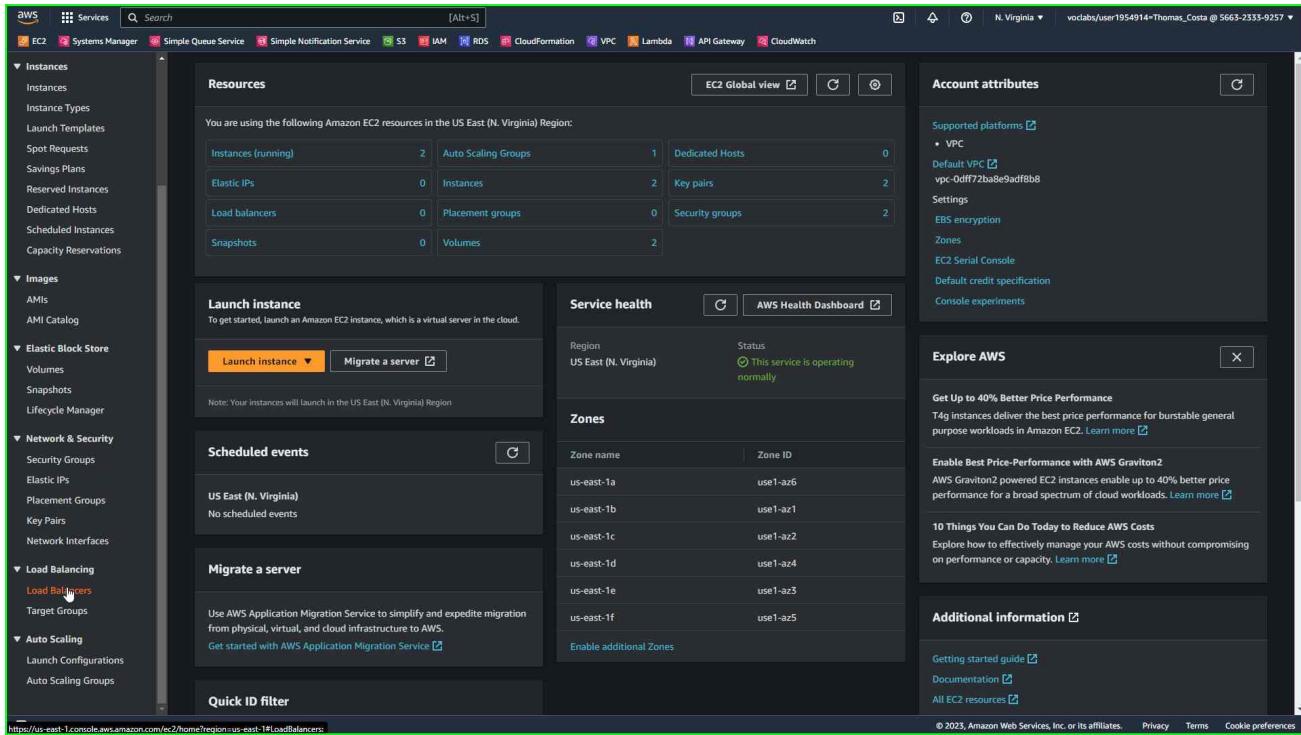
Total targets	Healthy	Unhealthy	Unused	Initial	Draining
2	0	0	2	0	0

Below the table, there is a section titled 'Distribution of targets by Availability Zone (AZ)' with a note: 'Select values in this table to see corresponding filters applied to the Registered targets table below.' The 'Health status' column shows 'Unused' for both targets. The 'Actions' dropdown menu is open, and the 'Register targets' button is highlighted with a green box.

Implantando um Application Load Balancer (ALB) na AWS com EC2

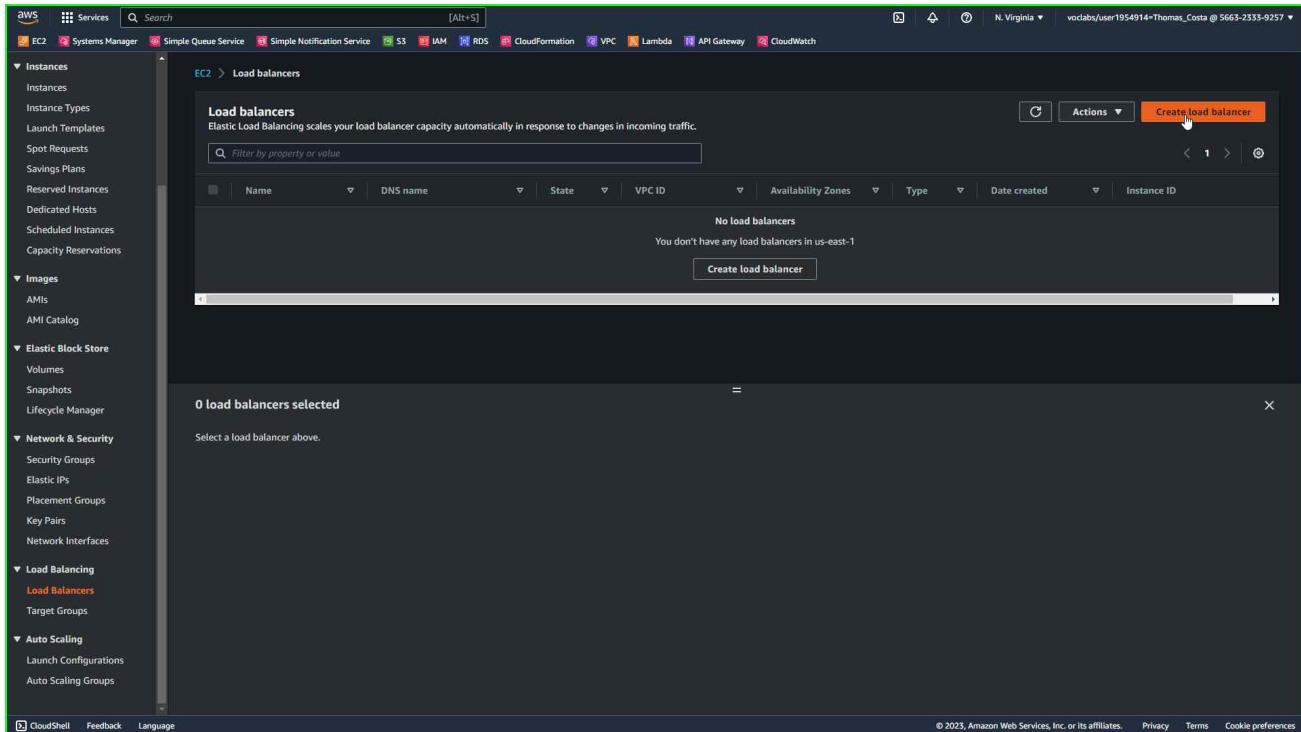
Parte 4 – Criando o Application Load Balancer (ALB)

Na tela principal do EC2, selecionar a opção Load Balancers:



The screenshot shows the AWS EC2 console dashboard. The sidebar on the left is expanded, showing various EC2 services: Instances, Images, Elastic Block Store, Network & Security, Load Balancing, and Auto Scaling. The 'Load Balancing' section is currently selected, with 'Load Balancers' highlighted in orange. The main content area displays EC2 resources in the US East (N. Virginia) Region, including Instances (running: 2), Auto Scaling Groups (1), Dedicated Hosts (0), Elastic IPs (0), Instances (2), Key pairs (2), Load balancers (0), Placement groups (0), Security groups (2), Snapshots (0), and Volumes (2). Below this, there are sections for 'Launch instance', 'Scheduled events', 'Migrate a server', and a 'Quick ID filter'. On the right, there is an 'Account attributes' sidebar with options like 'Supported platforms', 'Default VPC', and 'Settings', and an 'Explore AWS' sidebar with links to price performance, migration, and cost reduction guides.

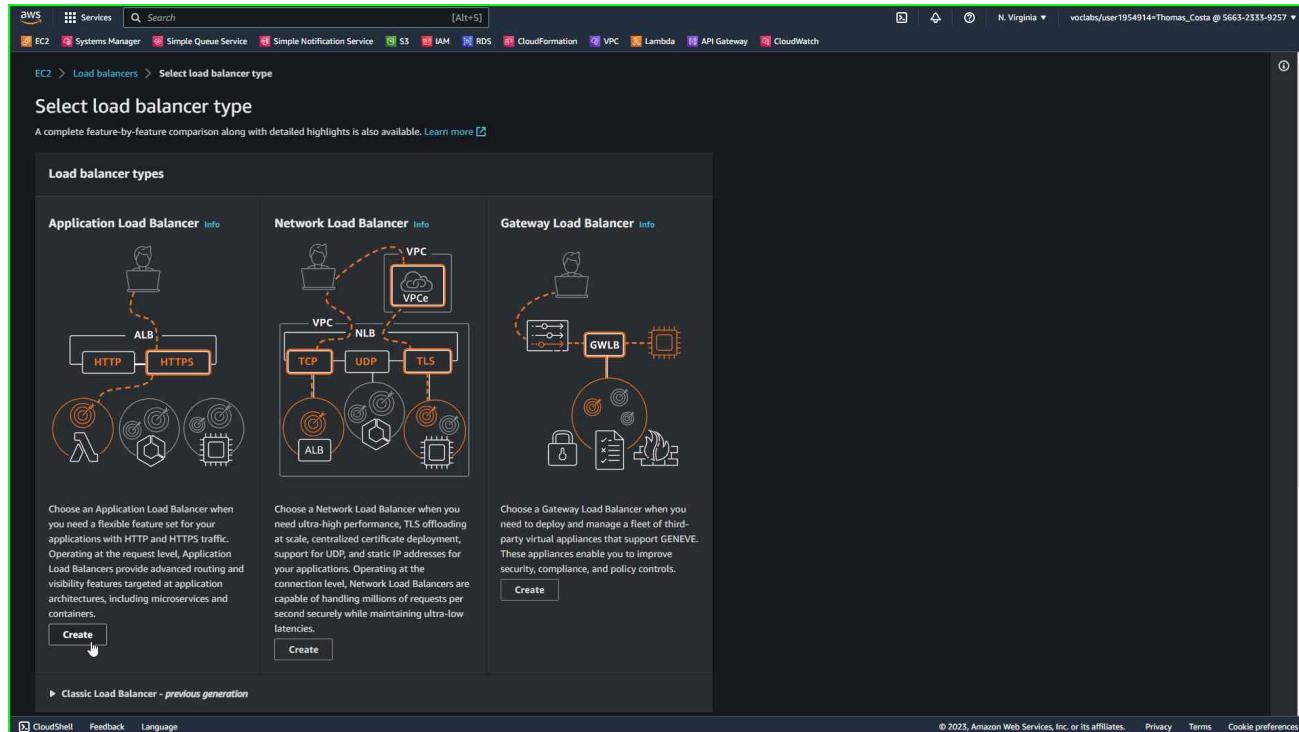
Clique no botão Create load balancer:



The screenshot shows the 'Load balancers' creation page within the EC2 console. The sidebar on the left is the same as the previous screenshot, with 'Load Balancing' selected. The main area is titled 'Load balancers' and contains a sub-header: 'Elastic Load Balancing scales your load balancer capacity automatically in response to changes in incoming traffic.' Below this is a search bar and a table header with columns: Name, DNS name, State, VPC ID, Availability Zones, Type, Date created, and Instance ID. A message 'No load balancers' is displayed, followed by a note: 'You don't have any load balancers in us-east-1'. A 'Create load balancer' button is located at the bottom of this section. A modal window titled '0 load balancers selected' is open, with the message 'Select a load balancer above.' and a close button 'X'.

Implantando um Application Load Balancer (ALB) na AWS com EC2

Clique no botão Create do item **Application Load Balancer**:



Load balancer types

Application Load Balancer Info

Choose an Application Load Balancer when you need a flexible feature set for your applications with HTTP and HTTPS traffic. Operating at the request level, Application Load Balancers provide advanced routing and visibility features targeted at application architectures, including microservices and containers.

Create

Network Load Balancer Info

Choose a Network Load Balancer when you need ultra-high performance, TLS offloading at scale, centralized certificate deployment, support for UDP, and static IP addresses for your applications. Operating at the connection level, Network Load Balancers are capable of handling millions of requests per second securely while maintaining ultra-low latencies.

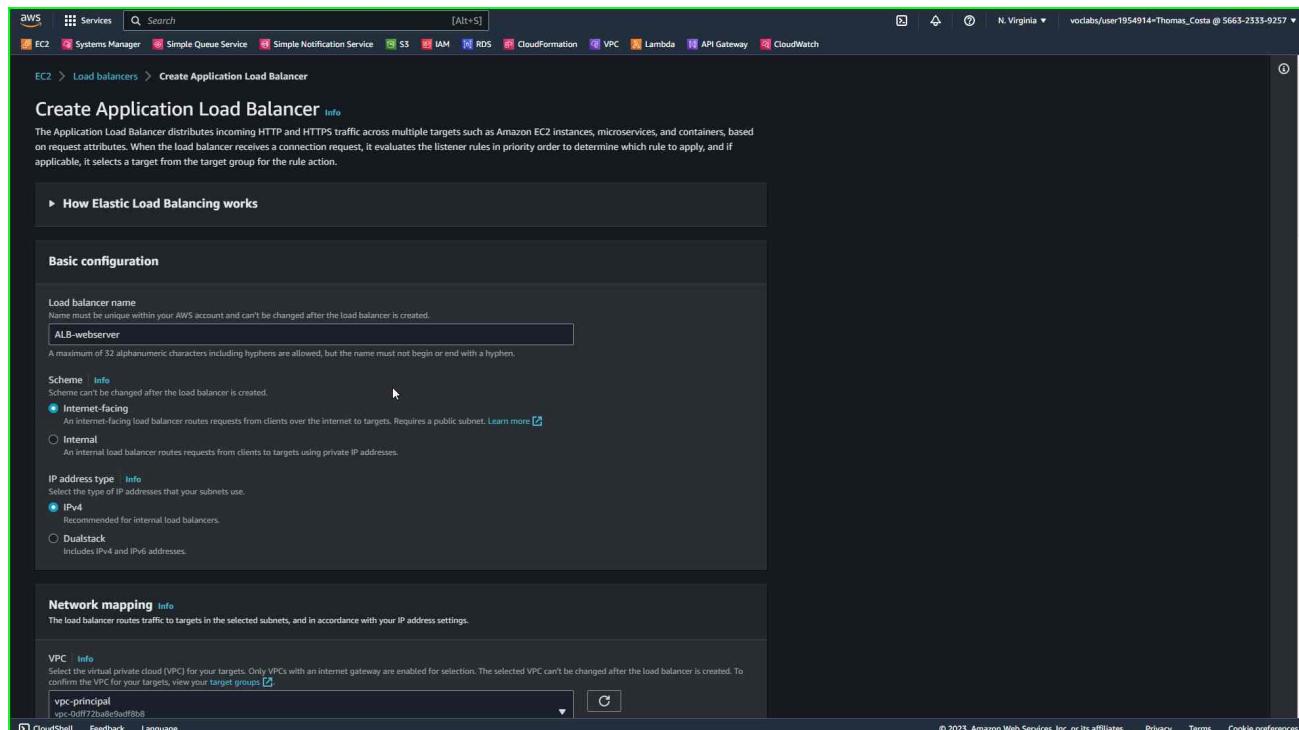
Create

Gateway Load Balancer Info

Choose a Gateway Load Balancer when you need to deploy and manage a fleet of third-party virtual appliances that support GENEVE. These appliances enable you to improve security, compliance, and policy controls.

Create

Coloque o nome do ALB de **ALB-webserver**:



Load balancer name

Name must be unique within your AWS account and can't be changed after the load balancer is created.

ALB-webserver

Scheme Info

Scheme can't be changed after the load balancer is created.

Internet-facing
An internet-facing load balancer routes requests from clients over the internet to targets. Requires a public subnet. [Learn more](#)

Internal
An internal load balancer routes requests from clients to targets using private IP addresses.

IP address type Info

Select the type of IP addresses that your subnets use.

IPv4
Recommended for internal load balancers.

Dualstack
Includes IPv4 and IPv6 addresses.

Network mapping Info

The load balancer routes traffic to targets in the selected subnets, and in accordance with your IP address settings.

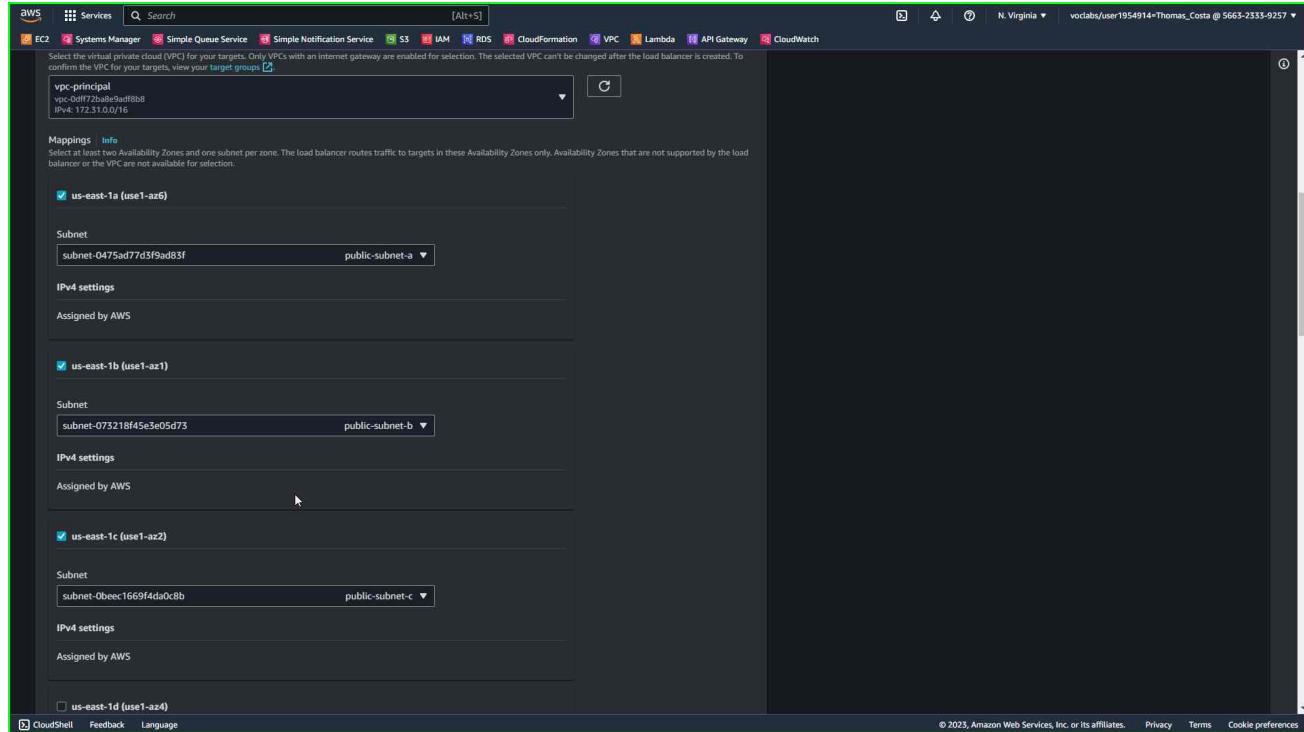
VPC Info

Select the virtual private cloud (VPC) for your targets. Only VPCs with an internet gateway are enabled for selection. The selected VPC can't be changed after the load balancer is created. To confirm the VPC for your targets, view your target groups.

vpc-principal
vpc-01ff177ba8a98d98b8

Implantando um Application Load Balancer (ALB) na AWS com EC2

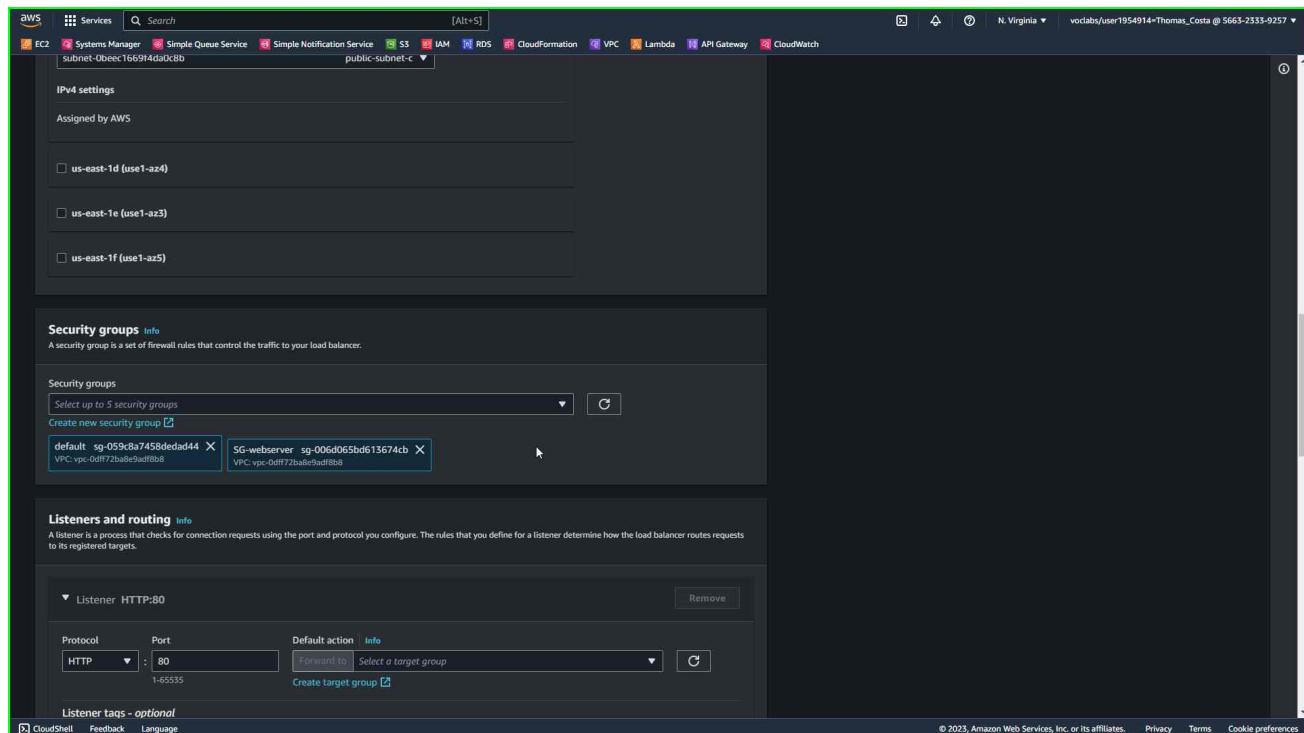
Selecione 3 AZ conforme imagem abaixo:



The screenshot shows the AWS VPC console with the following details:

- Virtual Private Cloud (VPC):** vpc-0df72ba8e9ad98b8
- Subnets:**
 - us-east-1a (use1-az6): subnet-0475ad77d3f9ad83f (public-subnet-a)
 - us-east-1b (use1-az1): subnet-073218f45e3e05d73 (public-subnet-b)
 - us-east-1c (use1-az2): subnet-0beecc1669f4da0c8b (public-subnet-c)
 - us-east-1d (use1-az4): (unchecked)
 - us-east-1e (use1-az3): (unchecked)
 - us-east-1f (use1-az5): (unchecked)

Selecione o Security Group SG-webserver:

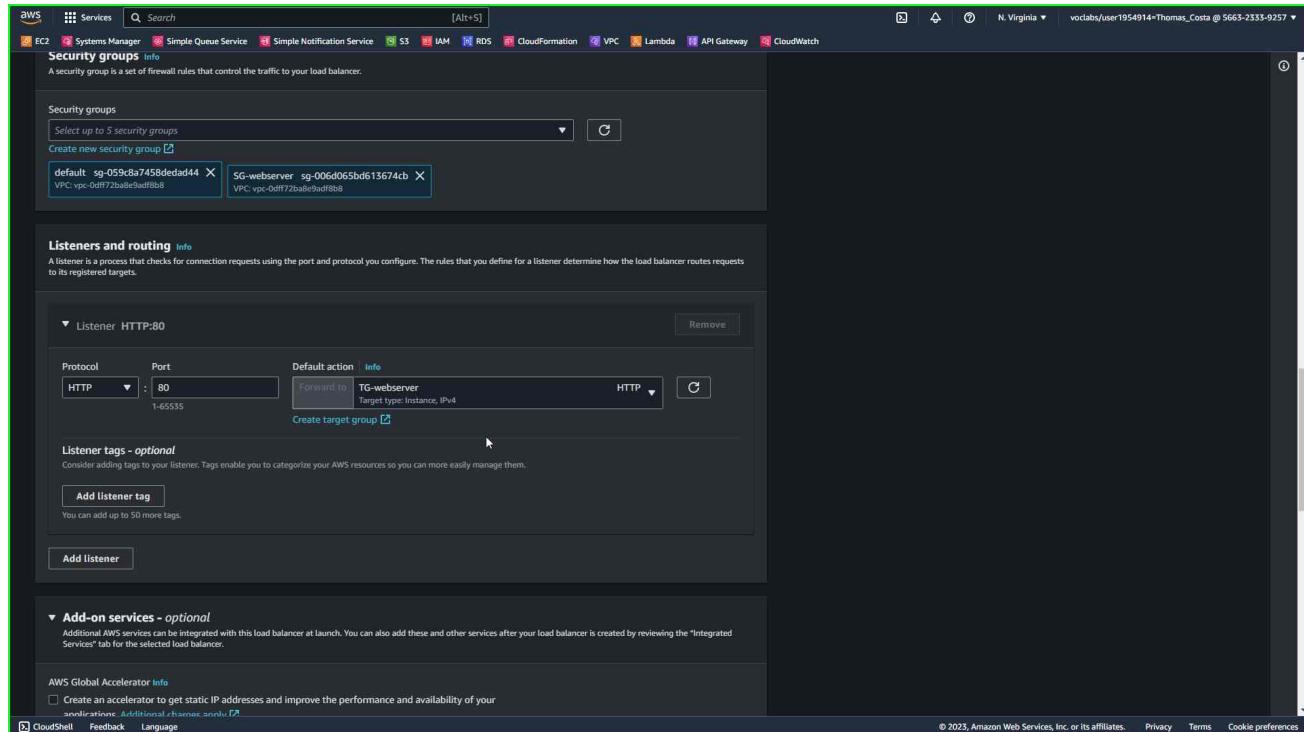


The screenshot shows the AWS Application Load Balancer (ALB) configuration page with the following details:

- Security groups:** SG-webserver (selected)
- Listeners and routing:**
 - Listener HTTP:80:** Protocol: HTTP, Port: 80, Default action: Select a target group (Forward to: SG-webserver)

Implantando um Application Load Balancer (ALB) na AWS com EC2

Selecione o Target group TG-webserver:



Security groups [Info](#)
A security group is a set of firewall rules that control the traffic to your load balancer.

Security groups
Select up to 5 security groups
Create new security group [\[+\]](#)

default sg-059c8a7458dedad44 X
VPC: vpc-0dff72ba8e9adfb8b
SG-webserver sg-006d065bd613674cb X
VPC: vpc-0dff72ba8e9adfb8b

Listeners and routing [Info](#)
A listener is a process that checks for connection requests using the port and protocol you configure. The rules that you define for a listener determine how the load balancer routes requests to its registered targets.

▼ Listener HTTP:80
Protocol: HTTP Port: 80 Default action: [Info](#) Forward to: TG-webserver Target type: Instance, IPv4
HTTP [\[+\]](#)
Create target group [\[+\]](#)

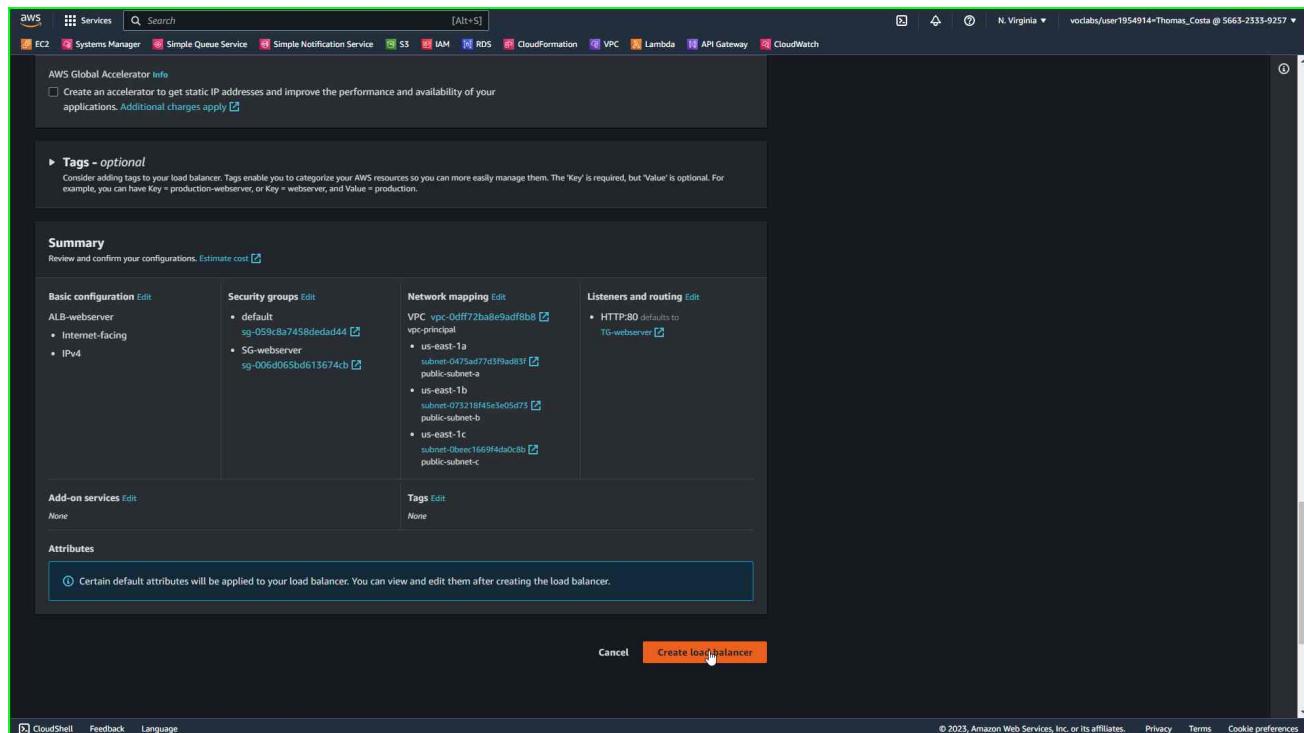
Listener tags - optional
Consider adding tags to your listener. Tags enable you to categorize your AWS resources so you can more easily manage them.
Add listener tag
You can add up to 50 more tags.

Add listener [\[+\]](#)

▼ Add-on services - optional
Additional AWS services can be integrated with this load balancer at launch. You can also add these and other services after your load balancer is created by reviewing the "Integrated Services" tab for the selected load balancer.

AWS Global Accelerator [Info](#)
Create an accelerator to get static IP addresses and improve the performance and availability of your applications. [Additional charges apply](#) [\[+\]](#)

Clique no botão Create load balancer:



AWS Global Accelerator [Info](#)
Create an accelerator to get static IP addresses and improve the performance and availability of your applications. [Additional charges apply](#) [\[+\]](#)

Tags - optional
Consider adding tags to your load balancer. Tags enable you to categorize your AWS resources so you can more easily manage them. The 'Key' is required, but 'Value' is optional. For example, you can have Key = production-webserver, or Key = webserver, and Value = production.

Summary
Review and confirm your configurations. [Estimate cost](#) [\[+\]](#)

Basic configuration [Edit](#)
ALB-webserver
• Internet-facing
• IPv4

Security groups [Edit](#)
• default
sg-059c8a7458dedad44 [\[+\]](#)
• SG-webserver
sg-006d065bd613674cb [\[+\]](#)

Network mapping [Edit](#)
VPC: vpc-0dff72ba8e9adfb8b
vpc-principal
• us-east-1a
subnet-0475ad277d3f9ad83f [\[+\]](#)
public-subnet-a
• us-east-1b
subnet-07221845e3e05d7f [\[+\]](#)
public-subnet-b
• us-east-1c
subnet-08e8c1669f4db0c8b [\[+\]](#)
public-subnet-c

Listeners and routing [Edit](#)
• HTTP:80 defaults to
TG-webserver [\[+\]](#)

Add-on services [Edit](#)
None

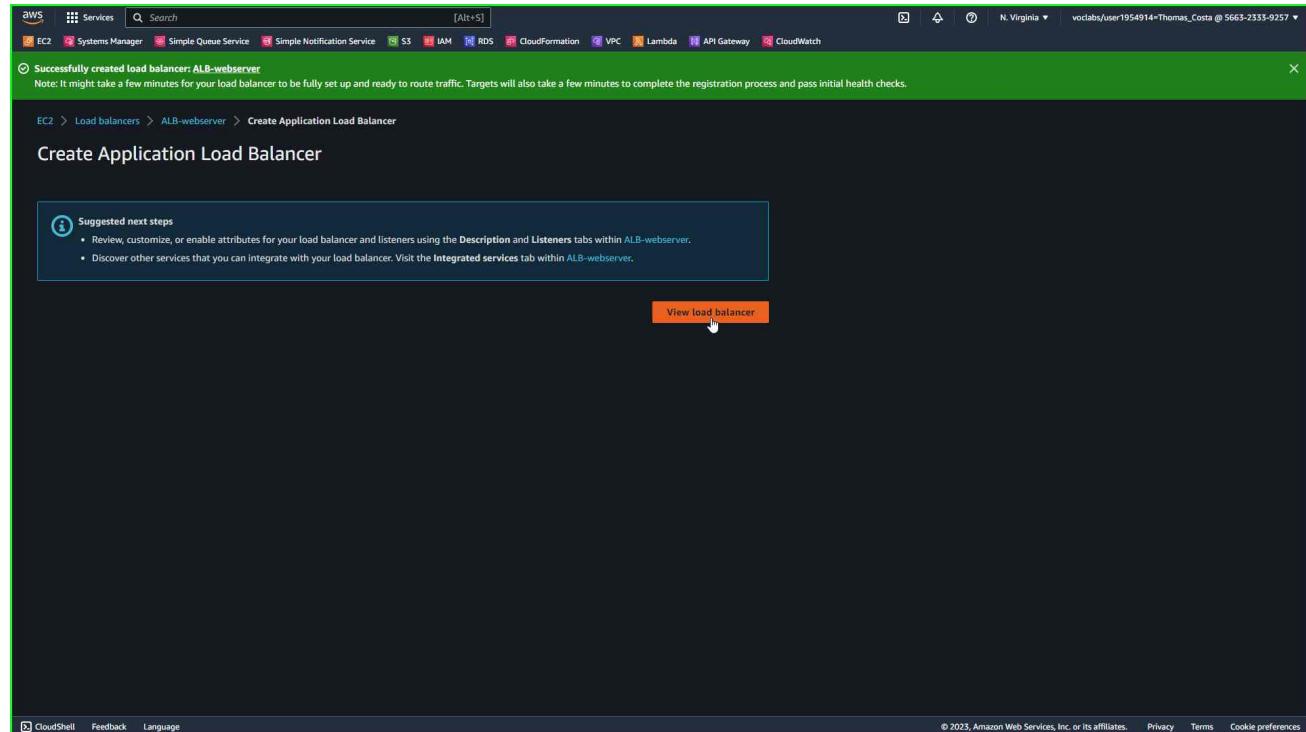
Tags [Edit](#)
None

Attributes
Certain default attributes will be applied to your load balancer. You can view and edit them after creating the load balancer.

[Cancel](#) [Create load balancer](#)

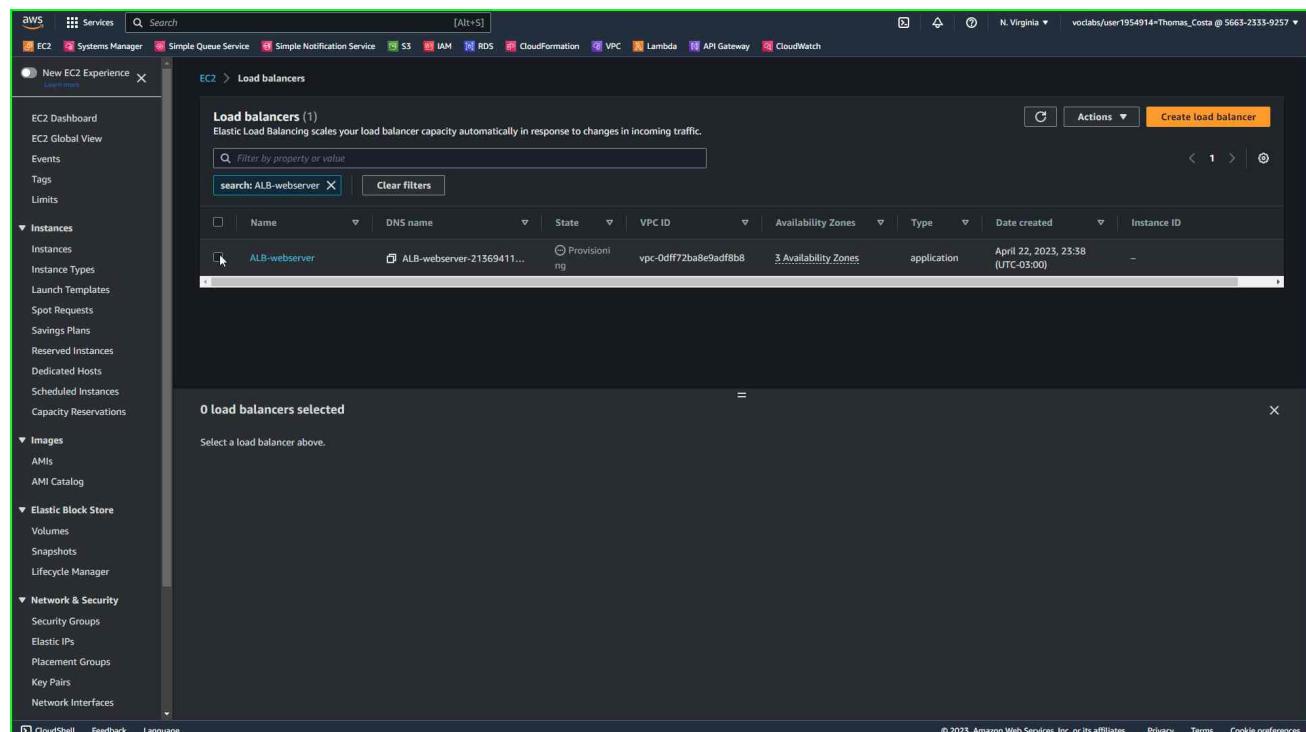
Implantando um Application Load Balancer (ALB) na AWS com EC2

ALB criado com sucesso:



The screenshot shows the AWS CloudFront 'Create Application Load Balancer' page. At the top, a green banner displays the message: 'Successfully created load balancer: ALB-webserver'. Below this, a note states: 'Note: It might take a few minutes for your load balancer to be fully set up and ready to route traffic. Targets will also take a few minutes to complete the registration process and pass initial health checks.' The main content area is titled 'Create Application Load Balancer'. A 'Suggested next steps' box contains two items: 'Review, customize, or enable attributes for your load balancer and listeners using the Description and Listeners tabs within ALB-webserver.' and 'Discover other services that you can integrate with your load balancer. Visit the Integrated services tab within ALB-webserver.' At the bottom right of this box is a 'View load balancer' button, which is highlighted with a mouse cursor. The bottom of the page includes standard AWS navigation links like CloudShell, Feedback, Language, and a copyright notice: '© 2023, Amazon Web Services, Inc. or its affiliates. Privacy Terms Cookie preferences'.

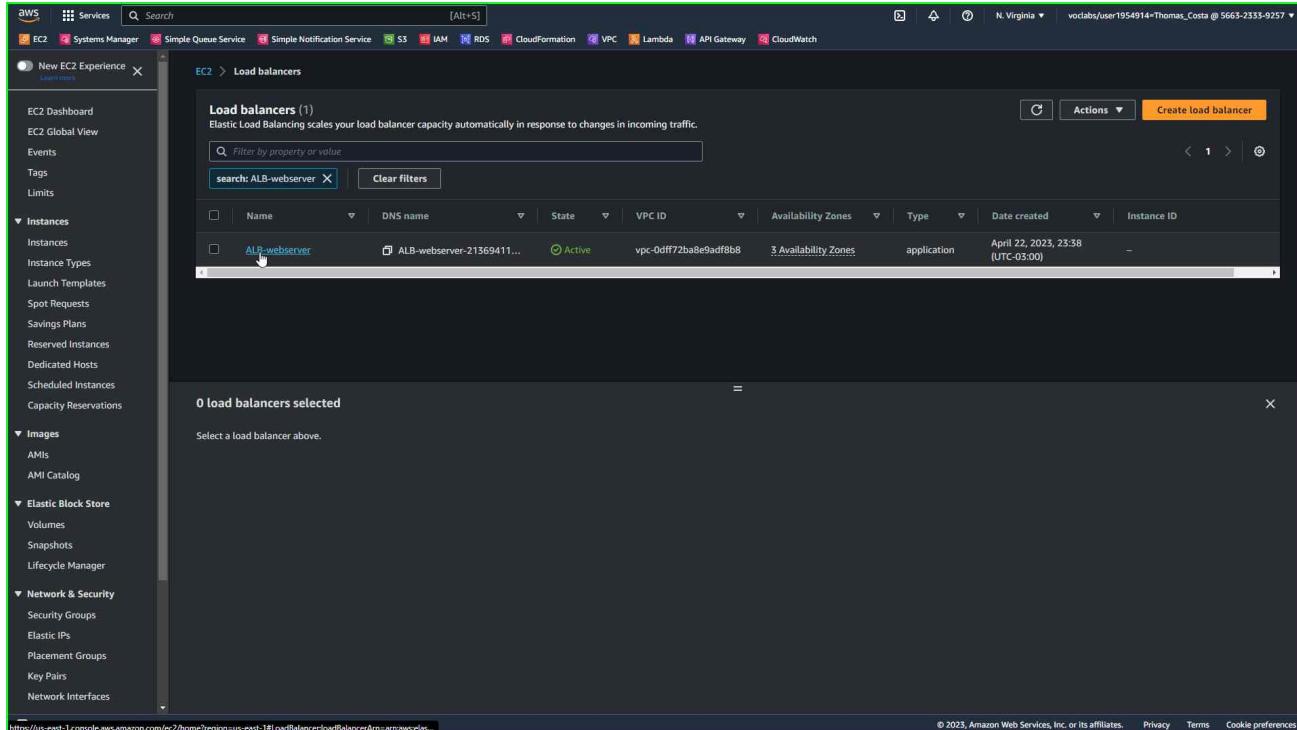
Aguardar o balanceador sair do status Provisioning:



The screenshot shows the AWS CloudFront 'Load balancers' page. The left sidebar is the 'New EC2 Experience' dashboard, with sections for EC2 Dashboard, EC2 Global View, Events, Tags, Limits, Instances, Images, and Network & Security. The main content area is titled 'Load balancers (1)'. It displays a table with one row for 'ALB-webserver'. The table columns are: Name, DNS name, State, VPC ID, Availability Zones, Type, Date created, and Instance ID. The 'ALB-webserver' row shows: ALB-webserver, ALB-webserver-21369411..., Provisioning, vpc-0dff72ba8e9adfb8b, 3 Availability Zones, application, April 22, 2023, 23:38 (UTC-03:00), and an empty Instance ID field. Below the table, a message says '0 load balancers selected' and 'Select a load balancer above.' The bottom of the page includes standard AWS navigation links like CloudShell, Feedback, Language, and a copyright notice: '© 2023, Amazon Web Services, Inc. or its affiliates. Privacy Terms Cookie preferences'.

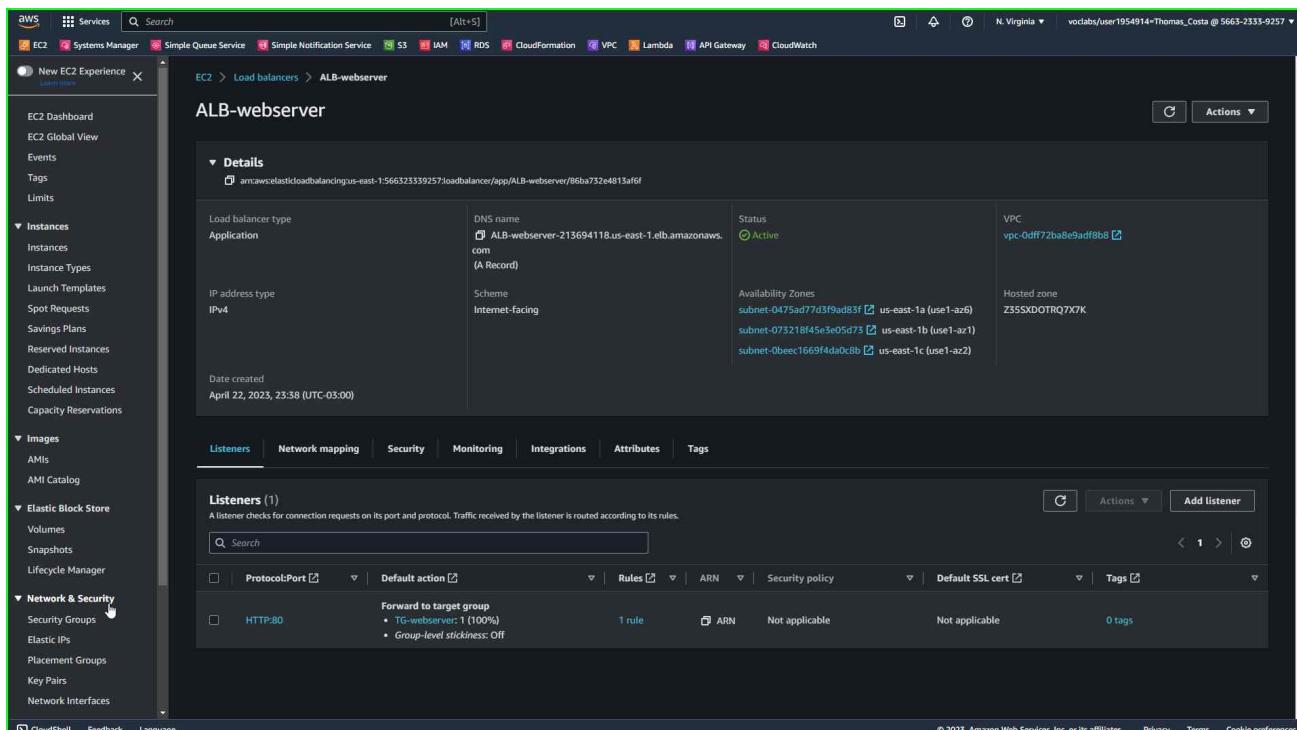
Implantando um Application Load Balancer (ALB) na AWS com EC2

Status do ALB em Active, pronto para ser utilizado:



The screenshot shows the AWS EC2 Load Balancers page. The left sidebar is expanded, showing various EC2-related options like Instances, Images, and Network & Security. The main content area is titled 'Load balancers (1)' and displays a table with one row. The row for 'ALB-webserver' shows the following details: Name (ALB-webserver), DNS name (ALB-webserver-213694118.us-east-1.elb.amazonaws.com), State (Active), VPC ID (vpc-0dff72ba8e9ad98b8), Availability Zones (3 Availability Zones), Type (application), Date created (April 22, 2023, 23:38 (UTC-03:00)), and Instance ID (empty). Below the table, a message says '0 load balancers selected' and 'Select a load balancer above.'

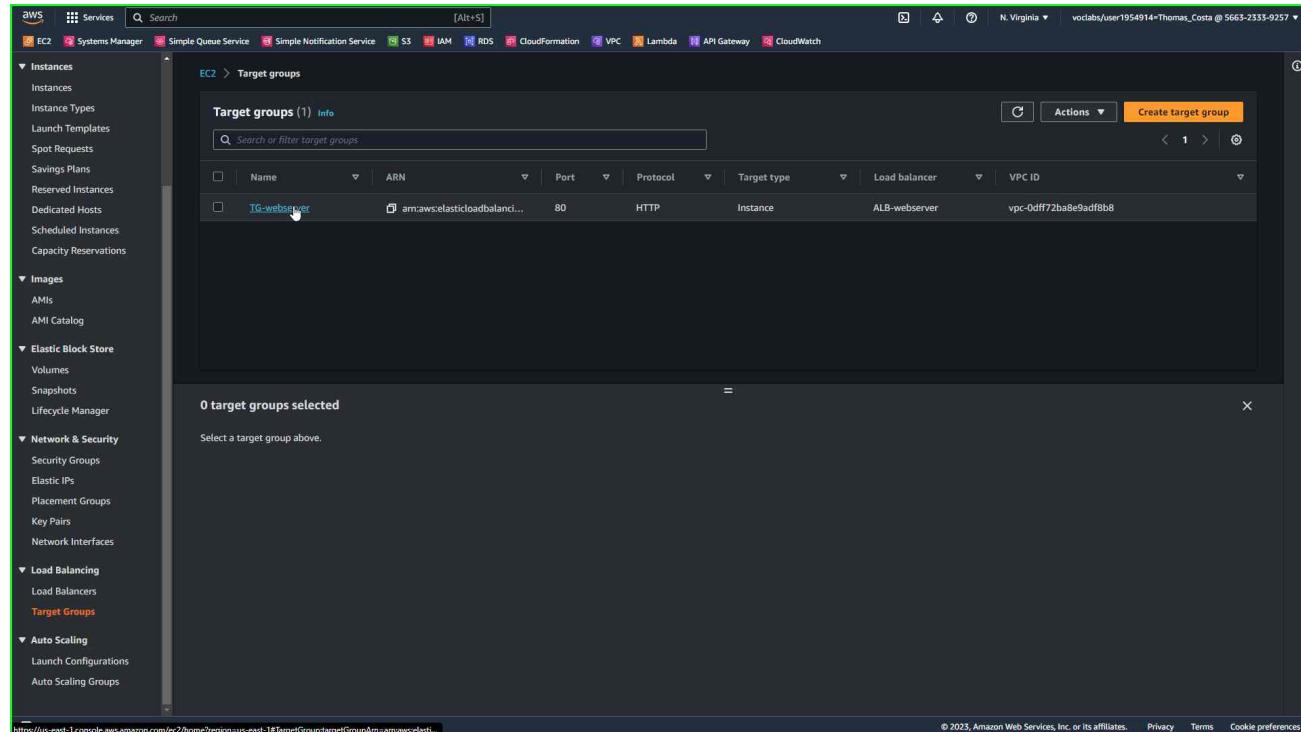
Selecione o ALB para ver as suas propriedades:



The screenshot shows the AWS EC2 Load Balancer properties page for 'ALB-webserver'. The left sidebar is expanded, showing various EC2-related options. The main content area is titled 'ALB-webserver' and has a 'Details' section. The 'Listeners' tab is selected, showing a table with one row for 'HTTP:80'. The row for 'HTTP:80' shows the following details: Protocol (HTTP), Port (80), Forward to target group (TG-webserver:1 (100%)), Rules (1 rule), ARN (empty), Security policy (Not applicable), Default SSL cert (empty), and Tags (0 tags). Below the table, a message says 'Listeners (1)' and 'A listener checks for connection requests on its port and protocol. Traffic received by the listener is routed according to its rules.'

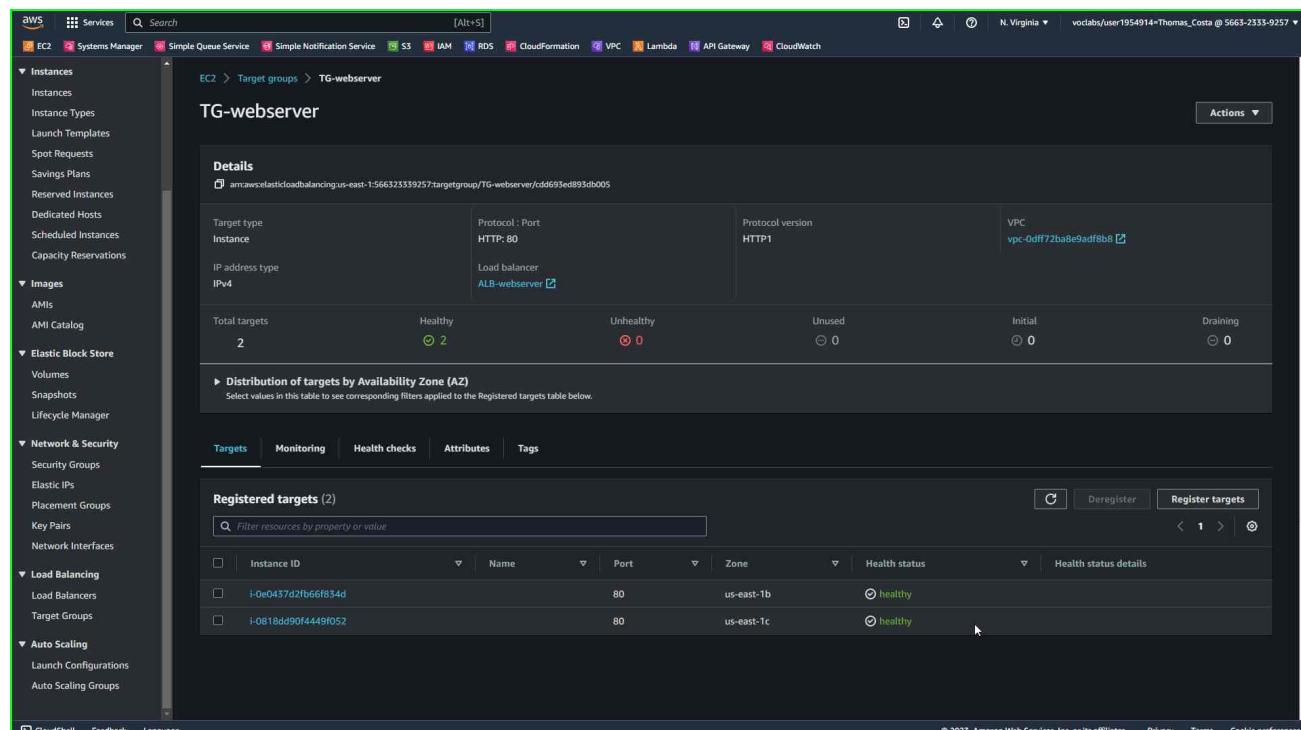
Implantando um Application Load Balancer (ALB) na AWS com EC2

Voltar a tela com a lista de Target groups e selecionar o que criamos nos itens anteriores:



The screenshot shows the AWS EC2 Target groups list. On the left, a sidebar navigation includes 'Instances', 'Images', 'Elastic Block Store', 'Network & Security', 'Load Balancing', 'Auto Scaling', and 'Target Groups'. The 'Target Groups' option is selected. The main pane displays a table with one row for 'TG-webserver'. The table columns are: Name, ARN, Port, Protocol, Target type, Load balancer, and VPC ID. The 'TG-webserver' row has the ARN 'arn:aws:elasticloadbalancing:us-east-1:1566323339257:targetgroup/cdd693ed893db005', Port '80', Protocol 'HTTP', Target type 'Instance', Load balancer 'ALB-webserver', and VPC ID 'vpc-0dff72ba8e9adfb8b8'. A modal window titled '0 target groups selected' is open, with the message 'Select a target group above.' and a close button 'X'.

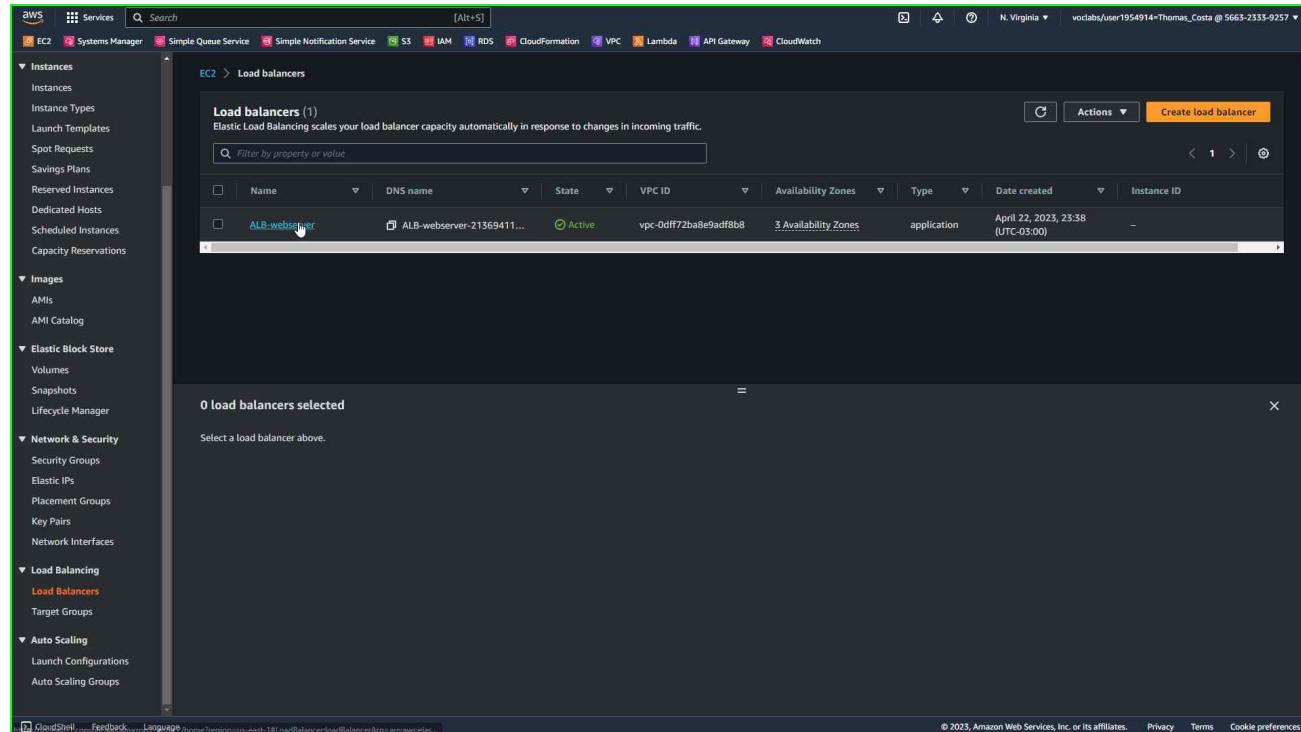
Balanceador reconhece as instâncias através do Target group conforme podemos ver o Health Status em healthy:



The screenshot shows the details for the 'TG-webserver' target group. The 'Details' section shows the ARN 'arn:aws:elasticloadbalancing:us-east-1:1566323339257:targetgroup/cdd693ed893db005', Target type 'Instance', Protocol 'HTTP: 80', Protocol version 'HTTP1', VPC 'vpc-0dff72ba8e9adfb8b8', and IP address type 'IPv4'. Below this, a table shows 'Total targets' as 2, with 2 'Healthy' and 0 'Unhealthy' and 'Unused'. The 'Targets' tab is selected, showing a table with two rows for instances 'i-0e0457d2fb66f834d' and 'i-0818dd90f4449f052', both in 'us-east-1b' and marked as 'healthy'. The 'Monitoring' tab is also visible.

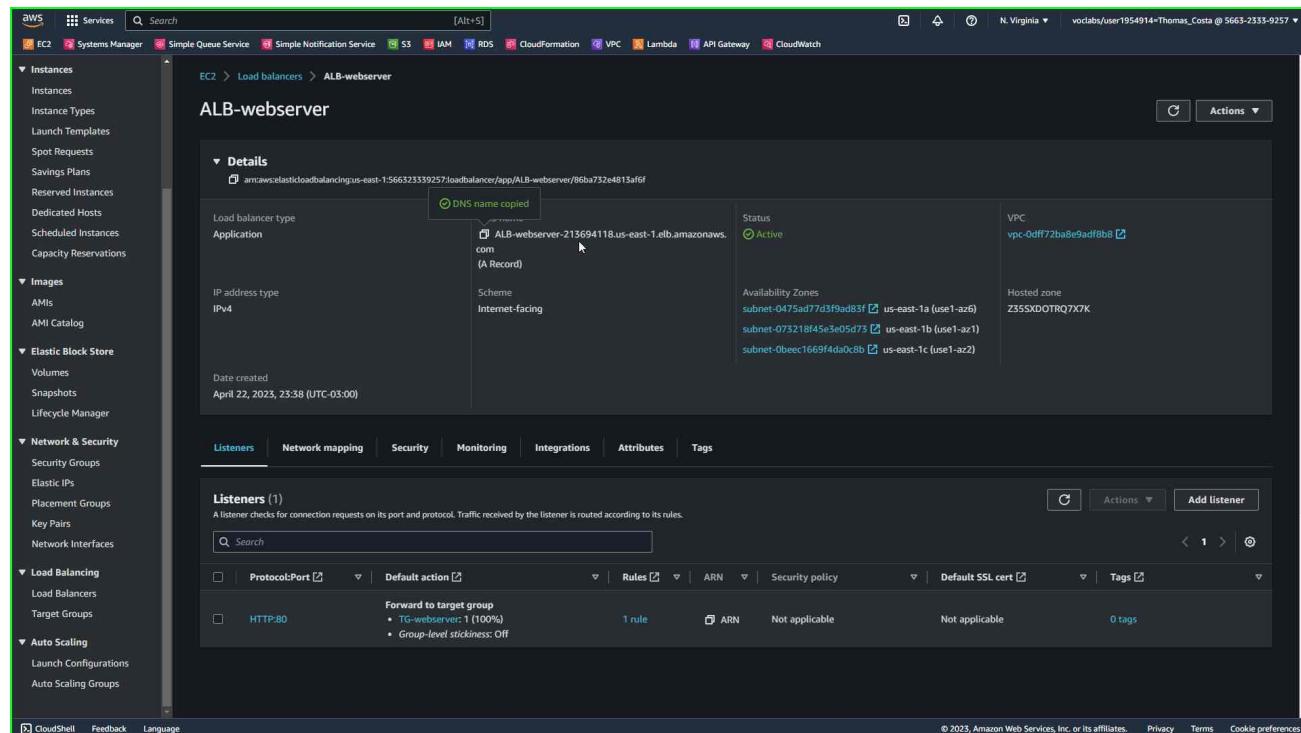
Implantando um Application Load Balancer (ALB) na AWS com EC2

Selecione novamente o平衡ador:



The screenshot shows the AWS EC2 Load Balancers page. The left sidebar is collapsed. The main content area shows a table with one row for 'ALB-webserver'. The table columns include Name, DNS name, State, VPC ID, Availability Zones, Type, Date created, and Instance ID. The 'ALB-webserver' row has a green 'Active' status and '3 Availability Zones'. The 'Actions' button is visible at the top right of the table. Below the table, a message says '0 load balancers selected' and 'Select a load balancer above.'

Copie o endereço público do平衡ador para testar no navegador:



The screenshot shows the AWS ALB Details page for 'ALB-webserver'. The left sidebar is collapsed. The main content area shows the 'Details' section with the DNS name 'ALB-webserver-213694118.us-east-1.elb.amazonaws.com' highlighted. Other details include Load balancer type: Application, Status: Active, VPC: 'vpc-0dff72ba8e9adff8b8', IP address type: IPv4, Scheme: Internet-facing, Availability Zones: 'us-east-1a (use1-az2)', 'us-east-1b (use1-az1)', and 'us-east-1c (use1-az2)'. The 'Listeners' tab is selected, showing a single listener for 'HTTP:80' with a 'Forward to target group' rule pointing to 'TG-webserver-1 (100%)'. The 'Listeners' table has columns for Protocol, Port, Default action, Rules, ARN, Security policy, Default SSL cert, and Tags. The 'Listeners' section also includes a search bar and an 'Add listener' button.

Implantando um Application Load Balancer (ALB) na AWS com EC2

Balanceador funcionando e acessando os EC2. Podemos verificar que a cada requisição a solicitação é enviada para uma instância:

Zona de Disponibilidade e IPs

Zona de Disponibilidade: us-east-1b

IP Público: 44.197.216.124

IP Privado: 172.31.14.81



Scripts do User Data

```
#!/bin/bash

sudo yum update -y

sudo yum install -y httpd

sudo systemctl start httpd

sudo systemctl enable httpd

EC2AZ=$(TOKEN=`curl -X PUT "http://169.254.169.254/latest/api/token" -H "X-aws-ec2-metadata-token-ttl-seconds: 21600"` && curl -H "X-aws-ec2-metadata-token: $TOKEN" -v http://169.254.169.254/latest/meta-data/placement/availability-zone)

EC2IPPUBLIC=$(TOKEN=`curl -X PUT "http://169.254.169.254/latest/api/token" -H "X-aws-ec2-metadata-token-ttl-seconds: 21600" && curl -H "X-aws-ec2-metadata-token: $TOKEN" -v http://169.254.169.254/latest/meta-data/public-ipv4`)

EC2IPPRIVATE=$(TOKEN=`curl -X PUT "http://169.254.169.254/latest/api/token" -H "X-aws-ec2-metadata-token-ttl-seconds: 21600" && curl -H "X-aws-ec2-metadata-token: $TOKEN" -v http://169.254.169.254/latest/meta-data/local-ipv4`)

echo '<html><head><title>Minha Página EC2</title></head><body><h1>Zona de Disponibilidade e IPs</h1><p><b>Zona de Disponibilidade:</b> AZID</p><p><b>IP Público:</b> IPPUBLIC</p><p><b>IP Privado:</b> IPPRIVATE</p></body></html>' > /var/www/html/index.txt

sudo sed "s/AZID/$EC2AZ/" /var/www/html/index.txt > /var/www/html/index1.txt

sudo sed "s/IPPUBLIC/$EC2IPPUBLIC/" /var/www/html/index1.txt > /var/www/html/index2.txt

sudo sed "s/IPPRIVATE/$EC2IPPRIVATE/" /var/www/html/index2.txt > /var/www/html/index.html
```