ECS - CASE STUDY

Andraz Brodnik brodul

@brodul

AS PERFORMED BY THE FRANK DANIELS COMIC OPERA COMPANY

THE IDOL'S EYE

COMIC OPERA IN THREE ACTS

LIBRETTO BY HARRY B. SMITH

MUSIC BY VICTOR HERBERT

Fairy Tales

Waltz Song

50 cts

EDWARD SCHUBERTH & CO.
(J.F.H.MEYER.)

NEW YORK
11 EAST 22ND STREET

LONDON
2 STAR YARD, CAREY ST.

LEIPZIG C. DIECKMANN

BASIC TERMS

- ALB => Application Load Balancer
- ECR => Elastic Container Registry (Dockerhub)
- CodeCommit (Github)
- CodeDeploy (handles different deployment strategies)
- CodePipeline
- Deployment strategies: Recreate, Blue/Green and Canary

ECS (FARGATE)

- Elastic Container Service
- Cluster
- Tasks (has containers)
- Services (ensures that Tasks are there if they crash)
- TaskDefinitions (blueprint for Tasks)

CUSTOMER

- Small startup
- Price performance
- Culture
- Focus on the product
- Infra as Code?
- KISS
- 2 challenges

CHALLENGE 1

C1: ECS PREFECT

- Python framework for data workflow orchestration
- You have flows that have connected task
- Can use ECS as execution platform
- Flows run once a day/week
- Very memory hungry

C1: ECS PREFECT

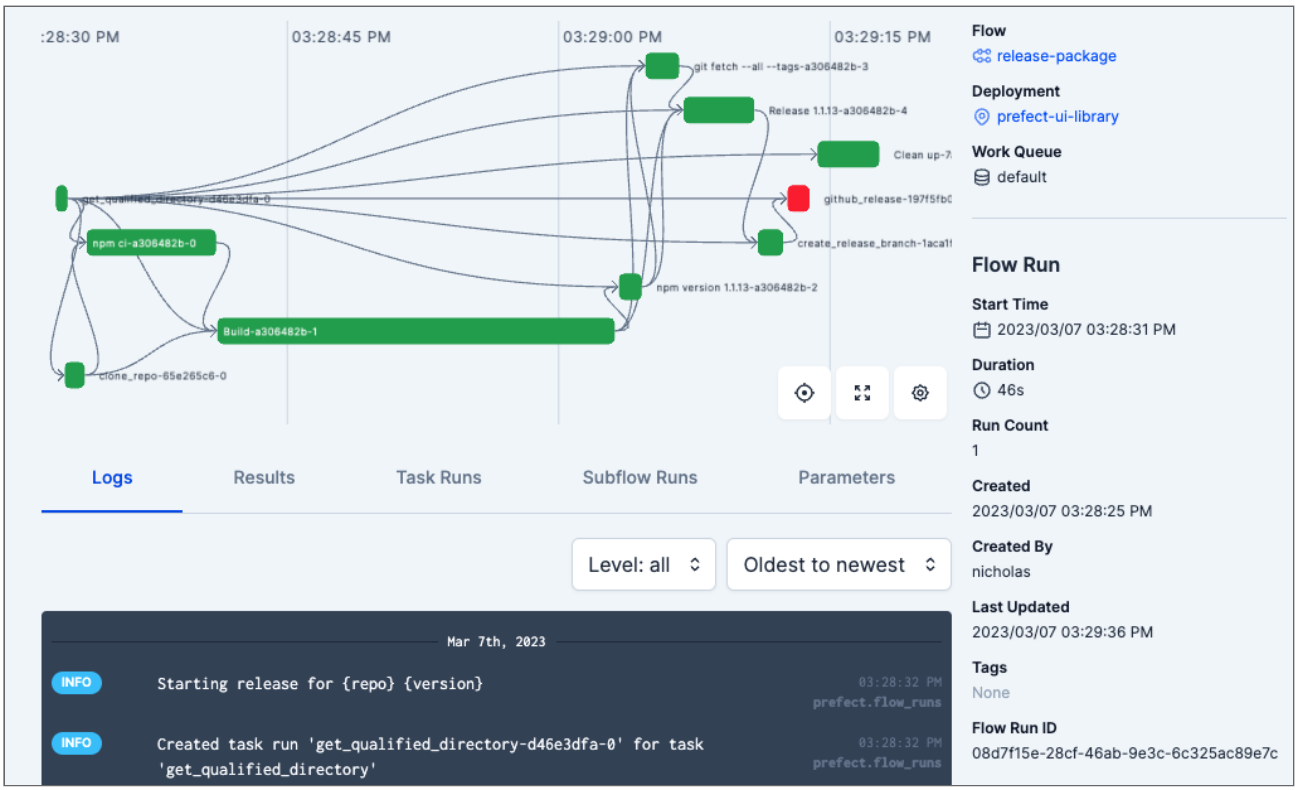

C1: ECS PREFECT

https://prefecthq.github.io/prefect-aws/ecs_guide/

C1: ECS

- After some research:
- 1 service with 1 task => worker that spawns more tasks on demand
- the tasks that compute the flows should be FARGATE or FARGATE_SPOT
- "lets do it quick in the console"

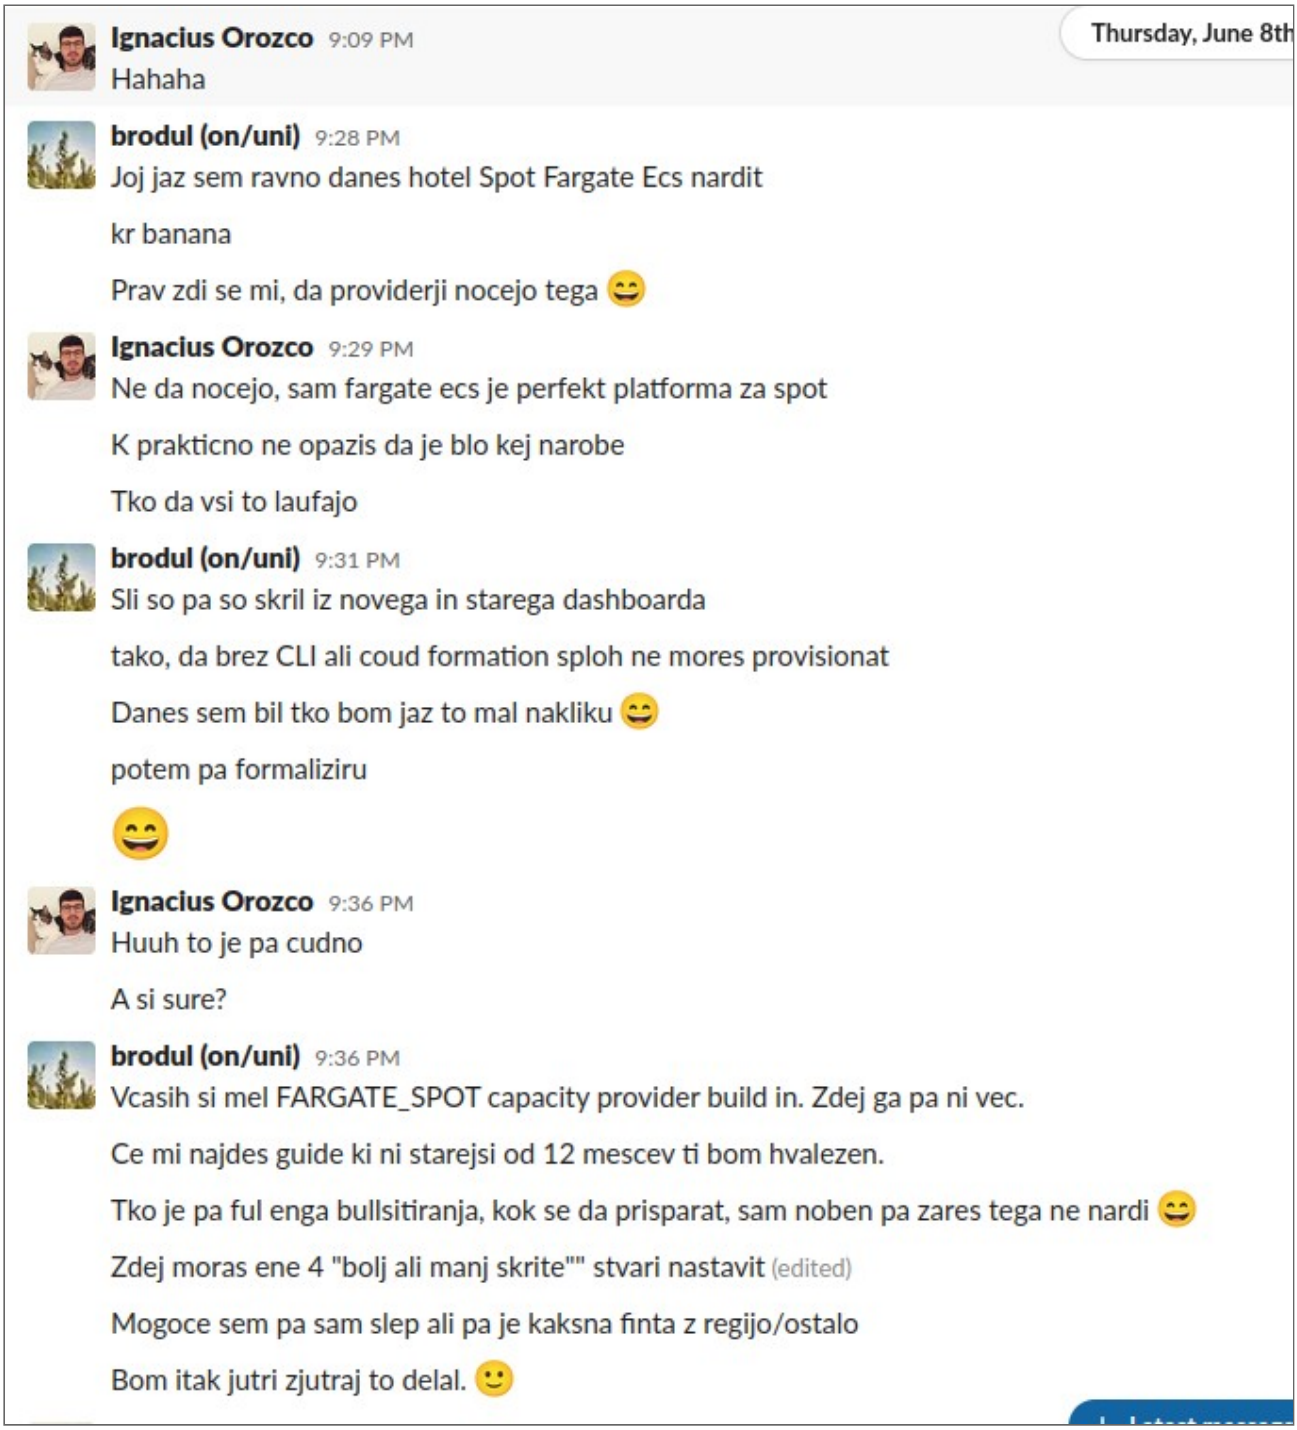
Thursday, June 8th
Ignacius Orozco 9:09 PM
Hahaha
brodul (on/uni) 9:28 PM
Joj jaz sem ravno danes hotel Spot Fargate Ecs nardit
kr banana
Prav zdi se mi, da providerji nocejo tega
Ignacius Orozco 9:29 PM
Ne da nocejo, sam fargate ecs je perfekt platforma za spot
K prakticno ne opazis da je blo kej narobe
Tko da vsi to laufajo
brodul (on/uni) 9:31 PM
Sli so pa so skril iz novega in starega dashboarda
tako, da brez CLI ali coud formation sploh ne mores provisionat
Danes sem bil tko bom jaz to mal nakliku
potem pa formaliziru
Ignacius Orozco 9:36 PM
Huuh to je pa cudno
A si sure?
brodul (on/uni) 9:36 PM
Vcasih si mel FARGATE_SPOT capacity provider build in. Zdej ga pa ni vec.
Ce mi najdes guide ki ni starejsi od 12 mescev ti bom hvalezen.
Tko je pa ful enga bullsitiranja, kok se da prisparat, sam noben pa zares tega ne nardi
Zdej moras ene 4 "bolj ali manj skrite"" stvari nastavit (edited)
Mogoce sem pa sam slep ali pa je kaksna finta z regijo/ostalo
Bom itak jutri zjutraj to delal.

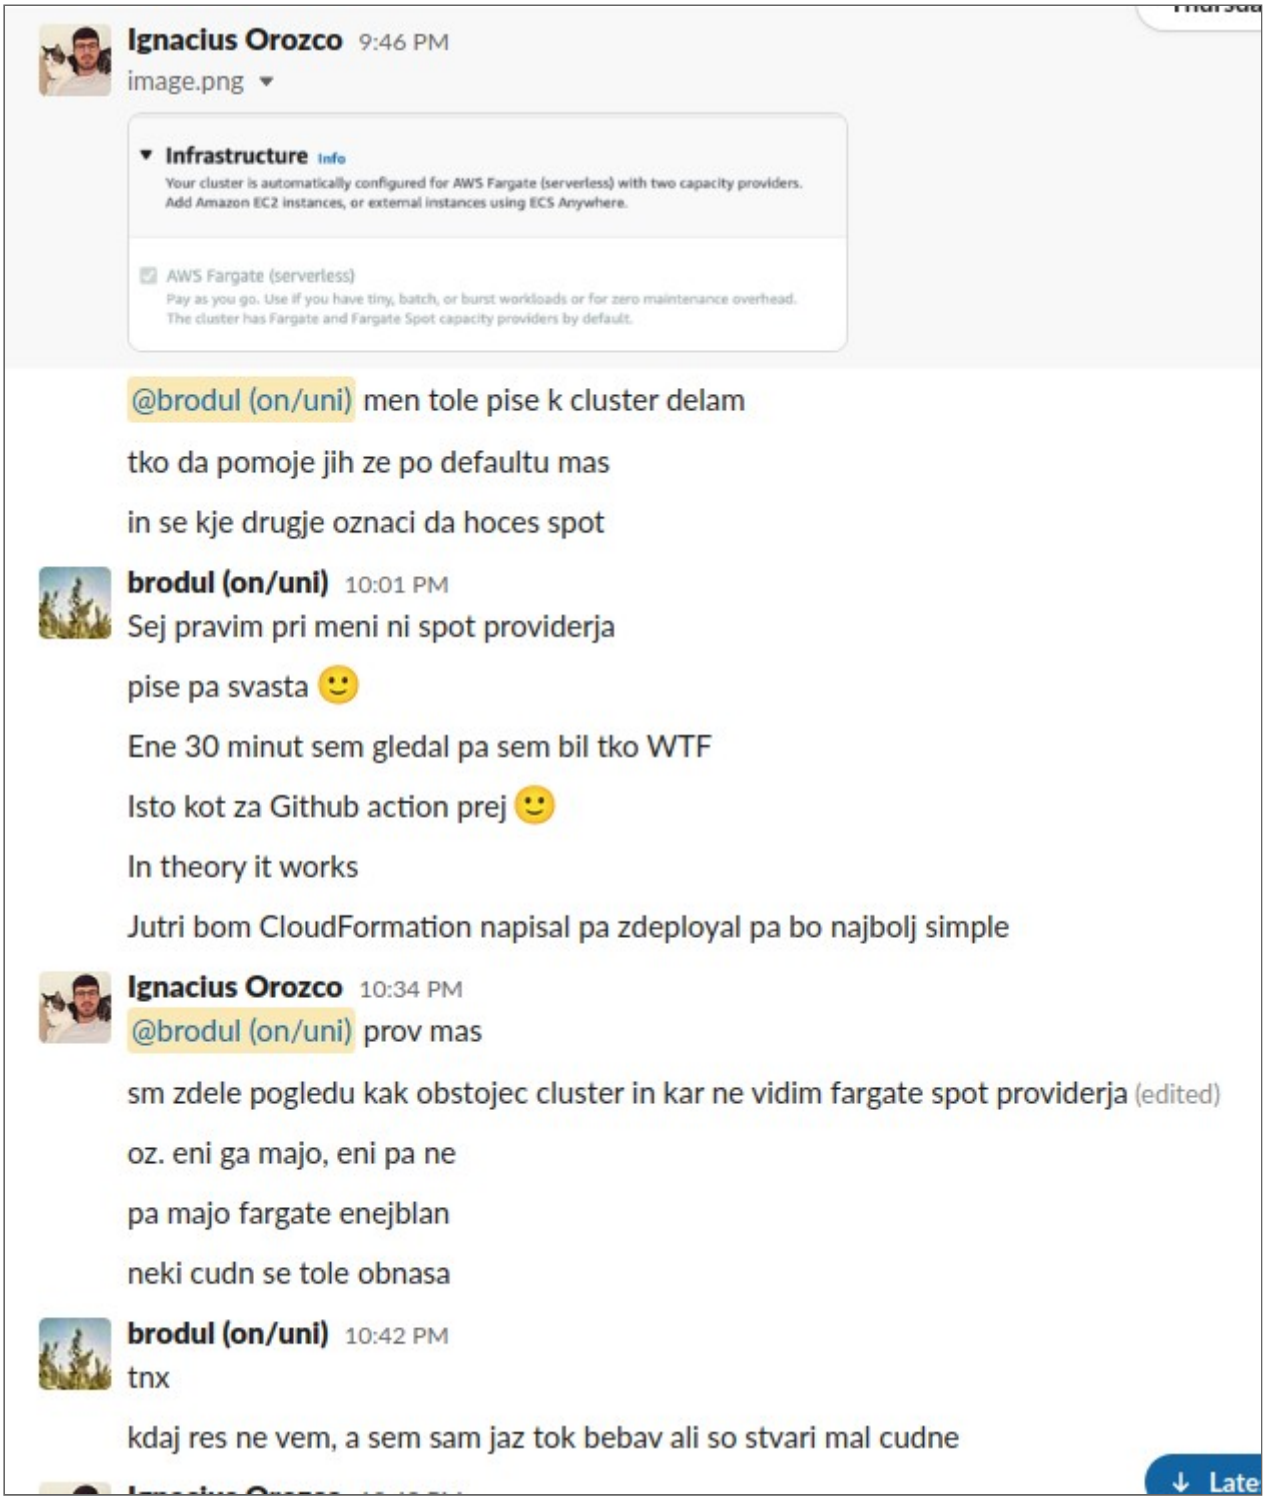
Ignacius Orozco 9:46 PM
image.png
Infrastructure Info
Your cluster is automatically configured for AWS Fargate (serverless) with two capacity providers.
Add Amazon EC2 instances, or external instances using ECS Anywhere.
AWS Fargate (serverless)
Pay as you go. Use if you have tiny, batch, or burst workloads or for zero maintenance overhead.
The cluster has Fargate and Fargate Spot capacity providers by default.
@brodul (on/uni) men tole pise k cluster delam
tko da pomoje jih ze po defaultu mas
in se kje drugje oznaci da hoces spot
brodul (on/uni) 10:01 PM
Sej pravim pri meni ni spot providerja
pise pa svasta
Ene 30 minut sem gledal pa sem bil tko WTF
Isto kot za Github action prej
In theory it works
Jutri bom CloudFormation napisal pa zdeployal pa bo najbolj simple
Ignacius Orozco 10:34 PM
@brodul (on/uni) prov mas
sm zdele pogledu kak obstojec cluster in kar ne vidim fargate spot providerja (edited)
oz. eni ga majo, eni pa ne
pa majo fargate enejblan
neki cudn se tole obnasa
brodul (on/uni) 10:42 PM
tnx
kdaj res ne vem, a sem sam jaz tok bebav ali so stvari mal cudne

Ignacius Orozco 10:43 PM

sm sel pogledat se v star ECS interface

pa ga tam tut ni

tko da sm mal mindfucked

brodul (on/uni) 10:44 PM

ja, jaz sem isto mislil da je samo nov UI issue

pa ga na starem tudi ni

Za Github action vsaj vem da sem hittal ze ene 4 odprte buge

😄

Ignacius Orozco 10:58 PM

nja ne vem

sm zdele se mal raziskoval pa si ne znam razlozit

sporoc ce bos kej ugotovu

9 replies

brodul (on/uni) 14 days ago
Mal mam feeling, da default capacity providerja ne pobere in je treba bit expliciten tudi na servicu

Ignacius Orozco 14 days ago
sam to je preko API, kaj pa v vmesniku?

brodul (on/uni) 14 days ago
men ni ratal

brodul (on/uni) 14 days ago
PAc ce hocem naredit FARGATE_SPOT CapacityProviderStrategy prek vmesnika mi rece, da se ne zme zacet z FARGATE

brodul (on/uni) 14 days ago
form validation

brodul (on/uni) 14 days ago
😄

brodul (on/uni) 14 days ago
Tudi nekega tutoriala nisem nasel, ki nebi bil outdatan

brodul (on/uni) 14 days ago
I guess

brodul (on/uni) 14 days ago
Neki majo sfukan

C1: SETUP DONE

- Used existing CloudFormation templates provided by Prefect community
- Configured Github actions Workflows to address specifics
- Docker image
- other
- Some modification to code
- ENV configuration
- AWS client initialization

"AWS offers you unlimited scaling ..."

" ... until you hit an AWS service qouta"

FARGATE ECS

Supported Configurations

CPU	Memory Values
0.25 vCPU	0.5 GB, 1 GB, and 2 GB
0.5 vCPU	Min. 1 GB and Max. 4 GB, in 1 GB increments
1 vCPU	Min. 2 GB and Max. 8 GB, in 1 GB increments
2 vCPU	Min. 4 GB and Max. 16 GB, in 1 GB increments
4 vCPU	Min. 8 GB and Max. 30 GB, in 1 GB increments
8 vCPU	Min. 16 GB and Max. 60 GB, in 4 GB increments
16 vCPU	Min. 32 GB and Max. 120 GB, in 8 GB increments

64 vCPU -> 2000 vCPU

AWSServiceRoleForServiceQuotas/CaseManager (Role) Fri Jun 16 2023 10:30:47 GMT+0200 (Central European Summer Time)	Limit increase request 1 Service: AWS Fargate Region: US East (Northern Virginia) Limit name: Fargate Spot vCPU resource count New limit value: 2000.0 ---------- Use case description: This support case was created by Service Quotas

andraz[redacted]
AM)

Tue Jun 20 2023
11:16:22 GMT+0200
(Central European
Summer Time)

Hello,

that's great news. But there is an error.
I checked the qouta in the Console and it's still at 64 vCPU

Best Andraz

Attachments
Selection_024.jpg

aws

Amazon Web
Services

Tue Jun 20 2023
04:00:48 GMT+0200
(Central European
Summer Time)

Translate

Was this response helpful? Click here to rate:
★★★★★

Hello there,

Thank you for your patience and detailed use case.

I am glad to inform your request has been approved for limit increase as follows :

Service: AWS Fargate
Region: US East (Northern Virginia)
Limit name: Fargate Spot vCPU resource count
New limit value: 2000.0

Please allow up to 30 minutes for the new quota to take effect and become available for use.

aws

Amazon Web Services

Wed Jun 21 2023 02:42:43 GMT+0200 (Central European Summer Time)

Translate

Was this response helpful? Click here to rate:
★★☆☆☆ Below Average

Hello there,

I'm really sorry to know that the limits shows same as before.

I've checked for you and could see that spot vCPU quota has been increased to 2000. However, I'd like to confirm that there's no confusion between Amazon Elastic Container Service (Amazon ECS) and AWS Fargate quotas .

In order to process your request, confirm if you would like to opt in to or out of Amazon Elastic Container Service (Amazon ECS) Fargate, Amazon Elastic Kubernetes Service (Amazon EKS) Fargate, or both.

andraz [redacted]
AM)

Wed Jun 21 2023
16:45:39 GMT+0200
(Central European
Summer Time)

Hello,

I clicked through all the regions and I see that there is a qouta increase in eu-central-1 Frankfurt of 4,000, but we requested a qouta increase in US East (Northern Virginia).

I can see that the limit in Northern Virginia is 64 vCPU (please see the image attached).

We would like to use ECS with Fargate Spot Capacity Provider. I checked all the ECS qoutas and we don't reach the limit. But we do reach the limit on the vCPU Fargate Spot qouta (please see the image attached).
We are currently interested in using ECS and not EKS.

I checked the link that you provided and figured that I requested the qouta increase decribed in the documentation.

Can you double check? In which region was the qouta increased? What is the reason that I can't see it on my end?

Attachments
Selection_030.jpg

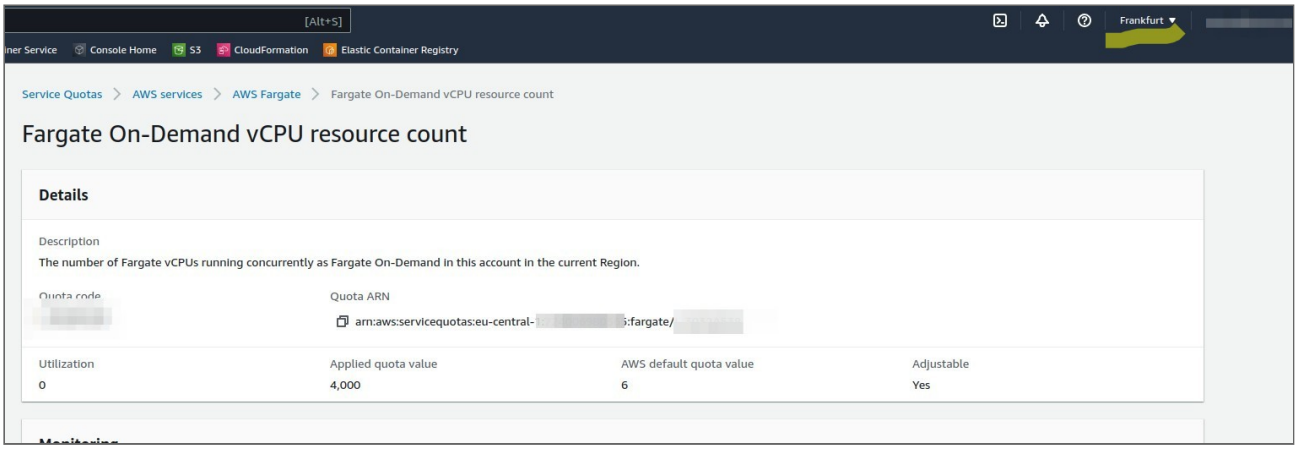
[Alt+S]
Frankfurt
iner Service
Console Home
S3
CloudFormation
Elastic Container Registry
Service Quotas
AWS services
AWS Fargate
Fargate On-Demand vCPU resource count
Fargate On-Demand vCPU resource count
Details
Description
The number of Fargate vCPUs running concurrently as Fargate On-Demand in this account in the current Region.
Quota code
Quota ARN
arn:aws:servicequotas:eu-central-
:fargate/
Utilization
0
Applied quota value
4,000
AWS default quota value
6
Adjustable
Yes

aws

Amazon Web Services

Mon Jun 26 2023 22:50:25 GMT+0200 (Central European Summer Time)

Translate

Was this response helpful? Click here to rate:

Hello there,

Sebastian here from the billing and accounts team. Thank you very much for your time and patience while we worked on your request.

I'm happy to inform you that after relentlessly pushing to expedite the several validations needed, and advocating in your behalf as strongly as possible, I have been able to approve your limit increase request as follows:

Limit increase request 1
Service: AWS Fargate
Region: US East (Northern Virginia)
Limit name: Fargate Spot vCPU resource count
New limit value: 2000

Learnings:

- Never open a support ticket on free tier support
- Qouta increases can take time. In our case 10 days

RATE LIMITING

```
Submission failed.
botocore.errorfactory.ClientException:
An error occurred (ClientException) when
calling
the RegisterTaskDefinition operation:
Too many concurrent attempts to create
a new revision of the specified family.
```

- 1 req/s Register / UnRegister Task Definition calls
- 10 req/s Task run call => 100/s Task runs total

PREFECT-AWS CODE

```
def _register_task_definition(
  self, ecs_client: _ECSClient,
task_definition: dict) -> str:
  """
  Register a new task definition with AWS.
  """
  # TODO: Consider including a global cache
  #       for this task definition since
  #       registration of task definitions
  #       is frequently rate limited
  task_definition_request =
copy.deepcopy(task_definition)

  # [...]

  return response["taskDefinition"]
["taskDefinitionArn"]
```

FARGATE SPOT VCPU

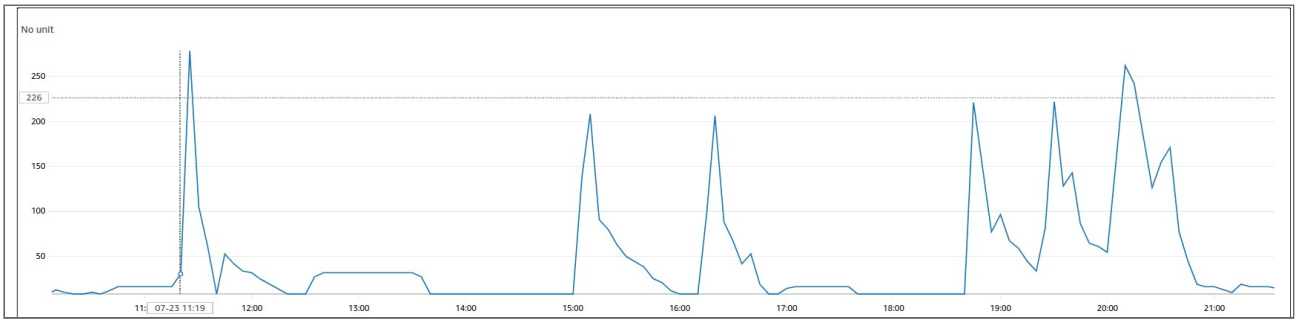

CHALLENGE 2

C2: RESTFUL PYTHON API

- Serves data from RDS
- Container
- Blue/green
- Potential to scale

C2: POTENTIAL SOLUTIONS

- Kubernetes - AWS EKS
- AWS AppRunner
- AWS ECS
- ...

C2: JUST FOLLOW THE TUTORIAL

Tutorial: Create a pipeline with an Amazon ECR source and ECS-to-CodeDeploy deployment

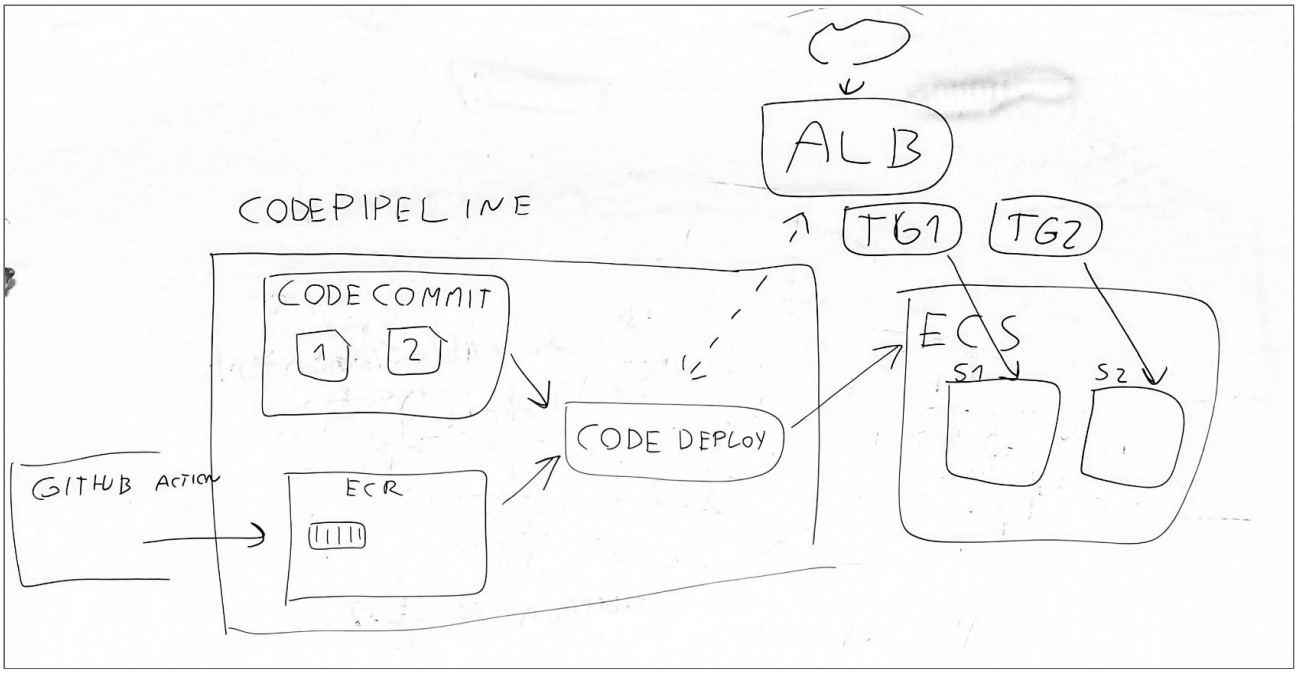
ALB
TG1
TG2
CODEPIPELINE
CODE COMMIT
1
2
ECS
S1
S2
CODE DEPLOY
GITHUB ACTION
ECR

APPSPEC.YAML

```
version: 0.0
Resources:
  - TargetService:
      Type: AWS::ECS::Service
      Properties:
        TaskDefinition: <TASK_DEFINITION>
        LoadBalancerInfo:
          ContainerName: "sample-website"
          ContainerPort: 80
```

TASKDEF.JSON

```
{
    "executionRoleArn":
"arn:aws:iam::account_ID:role/ecsTaskExecutionR
ole",
    "containerDefinitions": [
        {
            "name": "sample-website",
            "image": "<IMAGE1_NAME>",
            "essential": true,
            "portMappings": [
                {
                    "hostPort": 80,
                    "protocol": "tcp",
                    "containerPort": 80
                }
            ]
        }
    ],
    "requiresCompatibilities": [
        "FARGATE"
    ],
    "networkMode": "awsvpc",
    "cpu": "256",
    "memory": "512",
    "family": "ecs-demo"
}
```

C2: IAC TERRAFORM

```
brodul:infrastructure/ (api_deploy) $ cloc modules
      13 text files.
      13 unique files.
       0 files ignored.

github.com/AlDanial/cloc v 1.96  T=0.05 s (262.5 files/s, 16130.6 lines/s)
-------------------------------------------------------------------------------
Language                     files          blank        comment           code
-------------------------------------------------------------------------------
HCL                             13            115             66            618
-------------------------------------------------------------------------------
SUM:                            13            115             66            618
-------------------------------------------------------------------------------
```

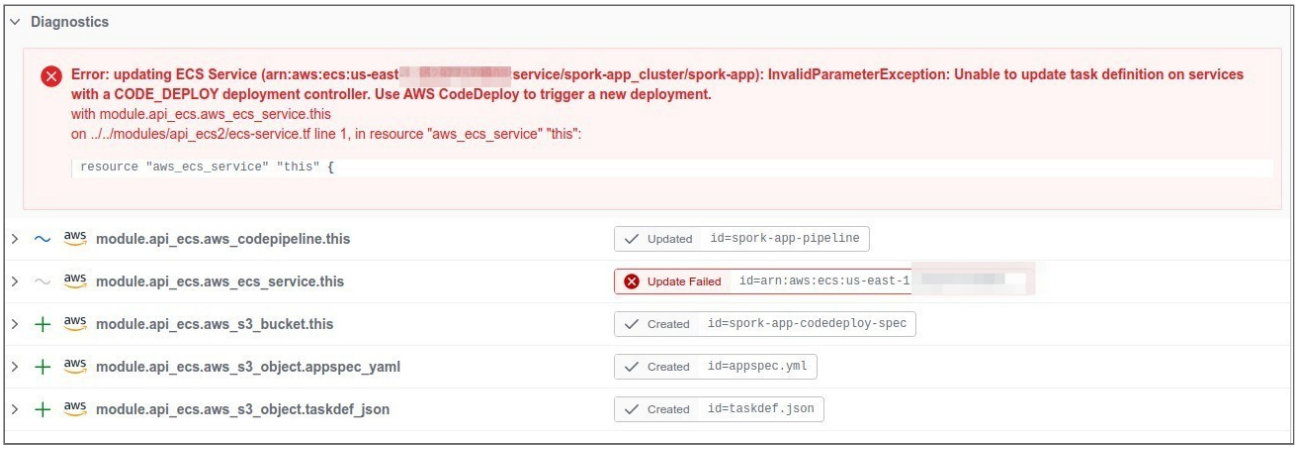
Diagnostics
Error: updating ECS Service (arn:aws:ecs:us-east service/spork-app_cluster/spork-app): InvalidParameterException: Unable to update task definition on services with a CODE_DEPLOY deployment controller. Use AWS CodeDeploy to trigger a new deployment.
with module.api_ecs.aws_ecs_service.this
on ../../modules/api_ecs2/ecs-service.tf line 1, in resource "aws_ecs_service" "this":
resource "aws_ecs_service" "this" {
module.api_ecs.aws_codepipeline.this
Updated id=spork-app-pipeline
module.api_ecs.aws_ecs_service.this
Update Failed id=arn:aws:ecs:us-east-1
module.api_ecs.aws_s3_bucket.this
Created id=spork-app-codedeploy-spec
module.api_ecs.aws_s3_object.appspec_yaml
Created id=appspec.yml
module.api_ecs.aws_s3_object.taskdef_json
Created id=taskdef.json

Error: creating ECS Service (spork-app): InvalidParameterException: TaskDefinition is delete_in_progress
with module.api_ecs.aws_ecs_service.this
on ../../modules/api_ecs2/ecs-service.tf line 1, in resource "aws_ecs_service" "this":

```
resource "aws_ecs_service" "this" {
```

Learnings:

- Don't think just follow the tutorial
- Some of the objects are immutable sometimes (managed by CodeDeploy)
- Sometimes it's hard to delete Task Definitions (or to find potential orphans)

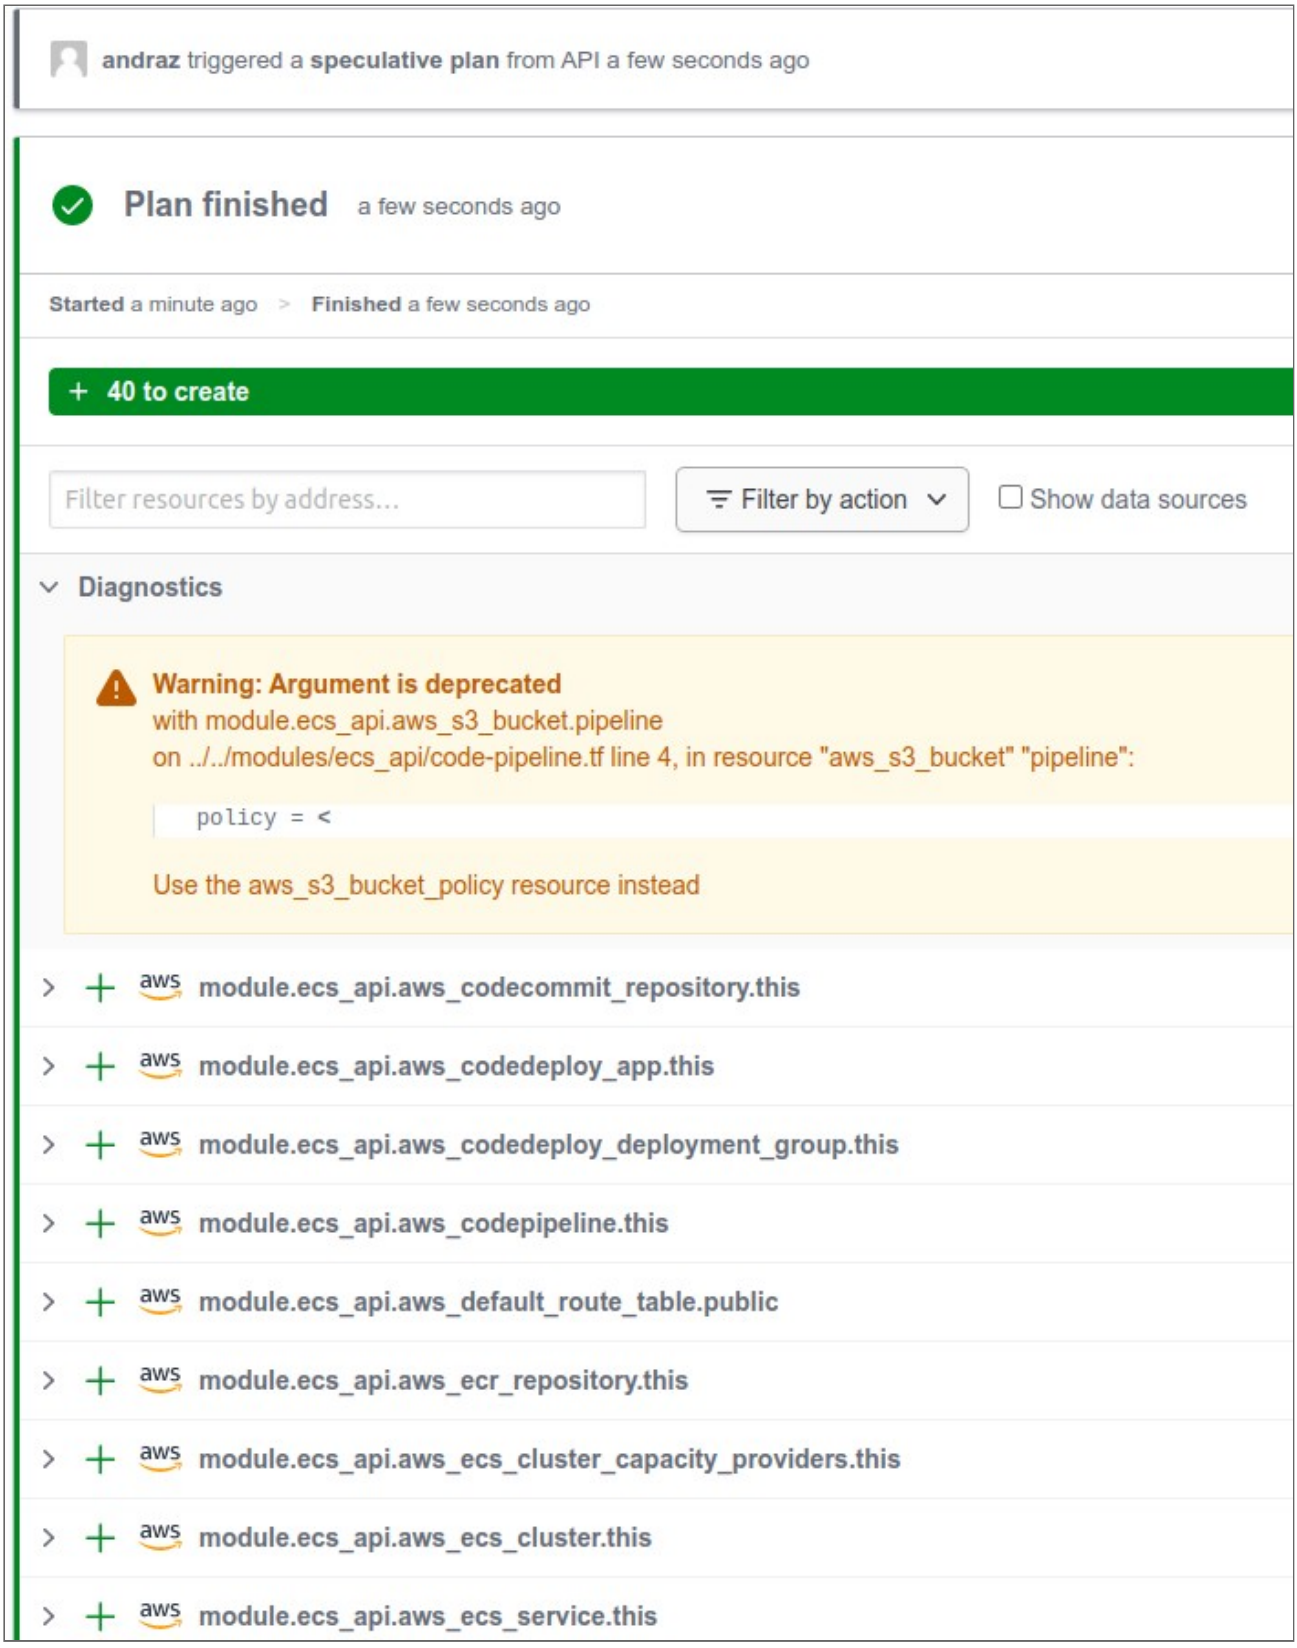
andraz triggered a speculative plan from API a few seconds ago
Plan finished a few seconds ago
Started a minute ago > Finished a few seconds ago
+ 40 to create
Filter resources by address...
Filter by action
Show data sources
Diagnostics
Warning: Argument is deprecated
with module.ecs_api.aws_s3_bucket.pipeline
on ../../modules/ecs_api/code-pipeline.tf line 4, in resource "aws_s3_bucket" "pipeline":
policy = <
Use the aws_s3_bucket_policy resource instead
module.ecs_api.aws_codecommit_repository.this
module.ecs_api.aws_codedeploy_app.this
module.ecs_api.aws_codedeploy_deployment_group.this
module.ecs_api.aws_codepipeline.this
module.ecs_api.aws_default_route_table.public
module.ecs_api.aws_ecr_repository.this
module.ecs_api.aws_ecs_cluster_capacity_providers.this
module.ecs_api.aws_ecs_cluster.this
module.ecs_api.aws_ecs_service.this

CURRENT ISSUES

- When there is a new ECR image CodePipeline doesn't trigger
- App specific (needs adoption)

Troubleshooting CodePipeline

WTH

Beeeeerrr
@brodul

Speaker notes