Article <u>Mikhail Khomenko</u> · Dec 23, 2019 12m read

Open Exchange

Deploying a Simple IRIS-Based Web Application Using Amazon EKS

Last time we deployed a simple IRIS application to the Google Cloud. Now we ' re going to deploy the same project to Amazon Web Services using its Elastic Kubernetes Service (<u>EKS</u>).

We assume you 've already forked the IRIS project to your own private repository. It 's calledsername>/my-objectscript-rest-docker-template in this article. <rootrepodir> is its root directory.

Before getting started, install the <u>AWS</u> command-line interface and, for Kubernetes cluster creation, <u>eksctl</u>, a simple CLI utility. For AWS you can try to use <u>aws2</u>, but you ' II need to set aws2 usage in kube config file as described <u>here</u>.

AWS EKS

Like AWS resources in general, EKS is not free. But you can create a free-tier account to play with AWS features. Keep in mind, though, that not everything you want to play with is included in the free tier. So, to manage your current budget and understand the financial issues, read this and this.

We ' II assume you already have an AWS account and root access to it, and that you don ' t use this root access but have <u>created a user</u> with admin permissions. You ' II need to put the access key and secret key of this user into the AWS credentials file under the [dev] profile (or whatever you choose to name the profile):

\$ cat Aaws/credentials [dev] awsaccesskeyid = ABCDEFGHIJKLMNOPQRST awssecretaccesskey = 1234567890ABCDEFGHIJKLMNOPQRSTUVWXYZ1234

We 're going to create resources in the AWS "eu-west-1" region, but you should chooser descent to your location and replace "eu-west-1" by your region everywhere below in the text.

By the way, all needed files (.circleci/, eks/, k8s/) are also stored here to simplify copying and pasting.

All required EKS resources will be created from scratch. You ' II find the mazon EKS Workshop to be a good resource to get an initial impression.

Now let 's check our access to AWS (we 've used a dummy account here): \$ export AWSPROFILE=dev

\$ aws sts get-caller-identity { "Account": "012345678910", "UserId": "ABCDEFGHIJKLMNOPQRSTU", "Arn": "arn:aws:iam::012345678910:user/FirstName.LastName"

}

\$ eksctl version

[] version.Info{BuiltAt:"", GitCommit:"", GitTag:"0.10.2"}

We could run "eksctl create cluster <cluster<u>n</u>ame> --region eu-west-1 " now, relying on the fact that all the default settings are good for us, or we can manage our own settings by creating a configuration file and using it.

The latter is preferable because it allows you to store such a file in a version control system (VCS). Examples of configurations can be found <u>here</u>. After reading about the different settings <u>here</u>, let 's try to create our own configuration:

\$ mkdir <root<u>repod</u>ir>/eks; cd <root<u>repod</u>ir>/eks

\$ cat cluster.yaml

apiVersion: eksctl.io/v1alpha5 kind: ClusterConfig

metadata:

name: dev-cluster region: eu-west-1 version: '1.14'

vpc:

cidr: 10.42.0.0/16 nat: gateway: Single clusterEndpoints: publicAccess: true privateAccess: false

nodeGroups:

name: ng-1 amiFamily: AmazonLinux2 ami: ami-059c6874350e63ca9 # AMI is specific for a region instanceType: t2.medium desiredCapacity: 1 minSize: 1 maxSize: 1

Worker nodes won't have an access FROM the Internet# But have an access TO the Internet through NAT-gatewayprivateNetworking: true

We don't need to SSH to nodes for demo ssh: allow: false

Labels are Kubernetes labels, shown when 'kubectl get nodes --show-labels' labels:

role: eks-demo

Tags are AWS tags, shown in 'Tags' tab on AWS console'

tags:

role: eks-demo

CloudWatch logging is disabled by default to save money

Mentioned here just to show a way to manage it

#cloudWatch:

clusterLogging:

enableTypes: []

Note that "nodeGroups.desiredCapacity = 1" would make no sense in a production environment, but it 's fine for our demo.

Also note that AMI images could differ between regions. Look for "amazon-eks-node-1.14" and choose one of the latest:

aws Servie	ces 🗸	Resource Groups	~ %					4	✓ Ireland ✓	Support 👻
New EC2 Experience Tell us what you think		Launch Actions *							Δ	0 + G
EC2 Dashboard New	^	Public images 👻 🔍	search : amazon-eks-node-1.14 💿	Add filter					② K < 1 to	5 of 5 > >
Events	4	Name	AMI Name	-	AMI ID	Source	Owner	Visibility	Status	Creation Date
Reports		a	amazon-eks-node-1.14-custom		ami-0f1064df170e77d49	987388283595/	987388283595	Public	available	September 30, 3
Limite		a	amazon-eks-node-1.14-v20190822		ami-0a7313ee6c5ea3183	amazon/amazon	amazon	Public	available	August 31, 2019
Limits		a	amazon-eks-node-1.14-v20190906		ami-0497f6feb9d494baf	amazon/amazon	amazon	Public	available	September 6, 2
▼ INSTANCES		a	amazon-eks-node-1.14-v20190927		ami-059c6874350e63ca9	amazon/amazon	amazon	Public	available	September 28, :
Instances		a	amazon-eks-node-1.14-v20191119		ami-02dca57ad67c7bf57	amazon/amazon	amazon	Public	available	November 19, 2
Instance Types										
Launch Templates New Take graphical screenshot Spot Requests	L	Select an AMI above								
Savings Plans										
Reserved Instances										
Dedicated Hosts										
Scheduled Instances										
Capacity Reservations										
▼ IMAGES AMIs Bundle Tasks										
ELASTIC BLOCK										

Now let 's create the cluster—the control plane and worker nodes: \$ eksctl create cluster -f cluster.yaml

By the way, when you no longer need a cluster, you can use the following to delete it: \$ eksctl delete cluster --name dev-cluster --region eu-west-1 --wait

Creating a cluster takes about 15 minutes. During this time you can look at the eksctl output:



You can also view <u>the CloudFormation console</u>, which will have two stacks. You can drill down into each one and look at the Resources tab to see exactly what will be created, and at the Events tab to check the current state of the resources creation.

CloudFormation ×	CloudFormation > Stacks						
Stacks	Stacks (2)		C Delete Update Stack actions ▼ Create stack Active ▼ ● View nested				
Exports	Q Filter by stack name						
Designer	Stack name	Status	Created time	Ť			
CloudFormation registry	O eksctl-dev-cluster-nodegroup-ng-1	ev-cluster-nodegroup-ng-1 OCREATE_IN_PROGRESS 2019-11-27 15:26:55 UTC+0200 EKS nodes (AI					
Resource types Take graphical screenshot	O eksctl-dev-cluster-cluster	CREATE_COMPLETE	2019-11-27 15:15:11 UTC+0200 EKS cluster (dedicated VPC: true, dedicated VPC)	at			

The cluster was successfully created, although you can see in the eksctl output that we had difficulties connecting to it: "unable to use kubectl with the EKS cluster".

Let's fix this by installing the <u>aws-iam-authenticator</u> (IAM) and using it to create a kube context: \$ which aws-iam-authenticator /usr/local/bin/aws-iam-authenticator

```
$ aws eks update-kubeconfig --name dev-cluster --region eu-west-1
```

\$ kubectl get nodes NAME STATUS ROLES AGE VERSION ip-10-42-140-98.eu-west-1.compute.internal Ready <none> 1m v1.14.7-eks-1861c5

It should work now, but we created a cluster with a user who has administrator rights. For the regular deployment process from CircleCI, it 's better toreate a special AWS user, named, in this case, CircleCI, with only programmatic access and the following policies attached:



The first policy is built into AWS, so you only need to select it. The second one should be created on your own. <u>Here</u> is a creation process description. The policy 'AmazonEKSDescribePolicy ' should look like:

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": [
               "eks:DescribeCluster",
               "eks:ListClusters"
        ],
        "Resource": "*"
        }
    ]
}
```

After user creation, save the user 's access key and secret access key — we 'll need them soon.

We also want to give this user rights within the Kubernetes cluster itself, as described in this<u>article</u>. In short, after creating the EKS cluster, only the IAM user, creator, has access to it. To add our CircleCl user, we need to replace default empty "mapUsers" section in the cluster AWS authentication settings (configmap aws-auth, 'data' section) by the following lines using <u>kubectl edit</u> (use your own Account Id instead of '01234567890 '):

\$ kubectl -n kube-system edit configmap aws-auth

•••

...

data:

mapUsers: | - userarn: arn:aws:iam::01234567890:user/CircleCl username: circle-ci groups: - system:masters

We ' II use the Kubernetes manifests fronthe earlier article (see the "Google Cloud Prerequisites" section) with one change: in the deployment image field we use placeholders. We ' II store those manifests in the rootrepodir>/k8s directory. Note that the deployment file was renamed to deployment.tpl:

\$ cat <root<u>repod</u>ir>/k8s/deployment.tpl

• • •

spec:

containers:

- image: DOCKERREPONAME/iris-rest:DOCKERIMAGETAG

• • •

CircleCI

The deployment process on the CircleCI side is similar to the process used for GKE:

- Pull the repository
- Build the Docker image
- Authenticate with Amazon Cloud
- Upload the image to Amazon Elastic Container Registry (ECR)
- Run the container based on this image in AWS EKS

Like last time, we ' II take advantage of already created and tested CircleCI configuration templatesorbs.

- <u>aws-ecr</u> orb for building an image and pushing it to ECR
- aws-eks orb for AWS authentication
- <u>kubernetes orb</u> for Kubernetes manifests deployment

Our deployment configuration looks like this: \$ cat <rootrepodir>/.circleci/config.yml version: 2.1 orbs:

```
aws-ecr: circleci/aws-ecr@6.5.0
 aws-eks: circleci/aws-eks@0.2.6
 kubernetes: circleci/kubernetes@0.10.1
jobs:
 deploy-application:
 executor: aws-eks/python3
 parameters:
 cluster-name:
 description: |
 Name of the EKS cluster
 type: string
 aws-region:
 description: |
 AWS region
 type: string
 account-url:
 description: |
 Docker AWS ECR repository url
 type: string
 tag:
 description: |
 Docker image tag
 type: string
 steps:
 - checkout
 - run:
 name: Replace placeholders with values in deployment template
 command:
 cat k8s/deployment.tpl |/
 sed "s|DOCKERREPONAME| << parameters.account-url >> |" |/
 sed "s|DOCKERIMAGETAG|<< parameters.tag >>|" > k8s/deployment.yaml; /
 cat k8s/deployment.yaml
 - aws-eks/update-kubeconfig-with-authenticator:
 cluster-name: << parameters.cluster-name >>
 install-kubectl: true
 aws-region: << parameters.aws-region >>
 - kubernetes/create-or-update-resource:
 action-type: apply
 resource-file-path: "k8s/namespace.yaml"
 show-kubectl-command: true
 - kubernetes/create-or-update-resource:
 action-type: apply
```

resource-file-path: "k8s/deployment.yaml" show-kubectl-command: true get-rollout-status: true resource-name: deployment/iris-rest namespace: iris - kubernetes/create-or-update-resource: action-type: apply resource-file-path: "k8s/service.yaml" show-kubectl-command: true namespace: iris workflows: main: jobs: - aws-ecr/build-and-push-image: aws-access-key-id: AWSACCESSKEYID aws-secret-access-key: AWSSECRETACCESSKEY region: AWSREGION account-url: AWSECRACCOUNTURL repo: iris-rest create-repo: true dockerfile: Dockerfile-zpm path: . tag: \${CIRCLESHA1} - deploy-application: cluster-name: dev-cluster aws-region: eu-west-1 account-url: \${AWSECRACCOUNTURL} tag: \${CIRCLESHA1} requires:

- aws-ecr/build-and-push-image

The <u>Workflows</u> section contains a list of jobs, each of which can be either called from an orb, such as <u>aws-ecr/build-and-push-image</u>, or defined directly in the configuration using "deploy-application".

The following code means that the deploy-application job will be called only after the aws-ecr/build-and-push-image job finishes:

requires:

- aws-ecr/build-and-push-image

The Jobs section contains a description of the deploy-application job, with a list of steps defined, including:

- run, to run a script that dynamically sets the Docker-image repository and tag
- aws-eks/update-kubeconfig-with-authenticator, which uses <u>aws-iam-authenticator</u> to set up a connection to Kubernetes
- <u>kubernetes/create-or-update-resource</u>, which is used several times as a way to run " kubectl apply " from CircleCI

We use variables and they, of course, should be defined in CircleCI on the "Environment variables" tab:

	Settings »	View my-objectscrip	View my-objectscript-rest-docker-template »		
-Gp	PROJECT SETTINGS Overview	Environment Variables			
1.	BUILD SETTINGS	Environment Variables for //my-object	script-rest-docker-template Import Variab	les Add Variable	
Take graphi	Environment Variables al screenshot	Add environment variables to the job. You can add sensitive of	data (e.g. API keys) here, rather than placing them ir	the repository.	
	Advanced Settings	Name	Value	Remove	
**	NOTIFICATIONS	AWS_ACCESS_KEY_ID	xxxxHTP5	×	
	Chat Notifications	AWS_ECR_ACCOUNT_URL	xxxx.com	×	
~0	Status Badges	AWS_REGION	xxxxst-1	×	
Ŷo		AWS_SECRET_ACCESS_KEY	xxxxdRQD	×	
	PERMISSIONS				
	Checkout SSH keys				

The following table shows the meaning of the variables used:

AWSACCESSKEYID	Access key of CircleC
AWS <u>S</u> ECRET <u>A</u> CCESS <u>K</u> EY	Secret key of CircleCl
AWS <u>R</u> EGION	eu-west-1, in this case
AWSECRACCOUNTURL	URL of the <u>AWS ECR</u> as 01234567890.dkr.ec where ' 01234567890

Here 's how we trigger the deployment process: \$ git add .circleci/ eks/ k8s/ \$ git commit -m "AWS EKS deployment " \$ git push This will show the two jobs in this workflow:

	Workflows »	» my-objectscript-rest-docker-template » master » 08e9fa5c-46e9-4778-a3	15-aaf044d896	da
ifp.		master / main <pre> Fix readme </pre>	26 sec ago	O0:26→ b9dfaf0
11.	Cancel			
Takegraphica	2 jobs in this work	low		
*	💿 aws-ecr/buil… 🖄	00:23 deploy-application		
¢ŝ				

Both jobs are clickable, and this allows you to see details of the steps taken. Deployment takes several minutes. Once it completes, we can check the status of the Kubernetes resources and of the IRIS application itself:

\$ kubectl -n iris get pods -w # Ctrl+C to stop

\$ kubectl -n iris get service NAME TYPE CLUSTER-IP EXTERNAL-IP PORT(S) AGE iris-rest LoadBalancer 172.20.190.211 a3de52988147a11eaaaff02ca6b647c2-66 3499201.eu-west-1.elb.amazonaws.com 52773:32573/TCP 15s

Allow several minutes to propagate the DNS-record. Until then you ' II receive a " Could not resolve host " error when running curl:

\$ curl -XPOST -H "Content-Type: application/json" -u <u>system</u>:SYS a3de5298814 7a11eaaaff02ca6b647c2-663499201.eu-

west-1.elb.amazonaws.com:52773/person/ -d '{"Name":"John Dou"}'

\$ curl -XGET -u <u>system</u>:SYS a3de52988147a11eaaaff02ca6b647c2-663499201. eu-west-1.elb.amazonaws.com:52773/person/all [{"Name":"John Dou"},]

Wrapping up

At first glance, deployment to AWS EKS looks more complex than to GKE, but it 's not really much different. If your organization uses AWS, you now know how to add Kubernetes to your stack.

Note that the EKS API was recently extended to support <u>managed groups</u>. These allow you to deploy the control plane and the data plane as a whole, and they look promising. Moreover, <u>Fargate</u>, the AWS serverless compute engine for containers, is now available.

Finally, a quick note about AWS ECR: don 't forget to set up <u>bifecycle policy</u> for your images.

<u>#AWS #Best Practices #Cloud #Containerization #DevOps #Docker #Kubernetes #InterSystems IRIS #Open</u> <u>Exchange</u> <u>Check the related application on InterSystems Open Exchange</u>

Source

URL: https://community.intersystems.com/post/deploying-simple-iris-based-web-application-using-amazon-eks