



Darpan Patel

Set Up Kubernetes Deployment

```
aws [Alt+S] United States (N. Virginia) Darpankumar-patel-IAM-Admin @ nextwork-alias-darpankumar-patel
[+] Building 10.9s (10/10) FINISHED docker:default
=> [internal] load build definition from Dockerfile 0.0s
=> => transferring dockerfile: 269B 0.0s
=> [internal] load metadata for docker.io/library/python:3. 0.4s
=> [internal] load .dockerignore 0.0s
=> => transferring context: 2B 0.0s
=> [1/5] FROM docker.io/library/python:3.9-alpine@sha256:39 1.7s
=> => resolve docker.io/library/python:3.9-alpine@sha256:39 0.0s
=> => sha256:3981e169a5c90cadeb071e8d40110.29KB / 10.29KB 0.0s
=> => sha256:6fff75d017c59dc0211cc72756d75 1.73KB / 1.73KB 0.0s
=> => sha256:30af12278e0e9852b18a2a3830e41 3.00KB / 3.00KB 0.0s
=> => sha256:fa07684b16b82247c3539e986a65ff 3.80MB / 3.80MB 0.1s
=> => sha256:09ebce43b37a228ea2d33853ee 460.21KB / 460.21KB 0.1s
=> => sha256:4ef74e18634db41cf0117f2ecb3 14.98MB / 14.98MB 0.5s
=> => extracting sha256:1697494014092247c3539e986a65ff 0.2s
=> => sha256:427eab949a336b38da329b251cd7f0a879 248B / 248B 0.2s
=> => extracting sha256:09ebce43b37a228ea2d33853ee3bd99fab 0.1s
=> => extracting sha256:4ef74e18634db41cf0117f2ecb07321e7 0.8s
=> => extracting sha256:427eab949a336b38da329b251cd7f0a879e 0.0s
=> [internal] load build context 0.1s
=> => transferring context: 42.42kB 0.0s
=> [6/9] WORKDIR /app 0.1s
```

i-09ca4f679d44abc94 (nextwork-eks-instance)
PublicIPs: 54.86.48.236 PrivateIPs: 172.31.95.165

CloudShell Feedback © 2025, Amazon Web Services, Inc. or its affiliates. Privacy Terms Cookie preferences



Darpan Patel

Introducing Today's Project!

In this project, I will clone a backend application from GitHub, build a Docker image of it, push that image to Amazon ECR, and prepare it for deployment with Kubernetes** because this is a crucial step in modern cloud-native development. It helps ensure that applications are containerized, version-controlled, and ready to be deployed in a scalable and reliable Kubernetes environment using AWS services. This workflow mirrors real-world DevOps pipelines and strengthens my skills in containerization, CI/CD, and cloud orchestration.

Tools and concepts

I used Amazon EKS, Git, Docker, and Amazon ECR to build, containerize, store, and deploy the backend application efficiently. Key steps include cloning the backend repository from GitHub, exploring and understanding the backend code, building a Docker image of the app, creating an Amazon ECR repository to securely store the container image, pushing the Docker image to ECR, and finally using Amazon EKS (Elastic Kubernetes Service) to deploy the containerized backend app in a scalable and consistent environment. This workflow ensures seamless integration, version control, and reliable deployment of the app.

Project reflection

This project took me approximately 1 hour to complete. The most challenging part was setting up and configuring the Amazon ECR repository and ensuring the Docker image pushed correctly with all permissions in place. My favourite part was exploring the backend code and understanding how the Flask app connects with the Hacker News API to fetch and process data, which gave me a clear picture of how the whole system works end-to-end.



Darpan Patel

Something new that I learnt from this experience was how to seamlessly integrate

Docker container images with Amazon ECR and then use those images in a Kubernetes environment like Amazon EKS. I also gained a deeper understanding of

What I'm deploying

To set up today's project, I launched a Kubernetes cluster. Steps I took to do this included launching an EC2 instance as my control hub, installing eksctl, and configuring AWS credentials. Then, I created an IAM role with AdministratorAccess and attached it to my EC2 instance to give it the necessary permissions. Finally, I used eksctl to create an EKS cluster with 3 t2.micro nodes. This setup gave me a fully managed Kubernetes environment where I could deploy and manage containerized applications efficiently.

I'm deploying an app's backend

In this step, I will connect to my EC2 instance again and install Git because I need Git to clone the backend application repository from GitHub. Once Git is installed and configured, I will clone the backend code so that I have the full application source locally on my EC2 instance, which is essential for building and deploying the backend to the Kubernetes cluster.



Darpan Patel

```
[ec2-user@ip-172-31-95-165 ~]$ ls
nextwork-flask-backend
```

Building a container image

Once I cloned the backend code, my next step is to build a container image of the backend. This is because Kubernetes deploys applications using container images, which package the app's code, dependencies, and environment settings into a single, portable unit. Building the container image ensures that the backend can run consistently and reliably across different environments, whether in development, testing, or production. It also enables Kubernetes to easily create, scale, and manage multiple instances of the backend as needed.

When I tried to build a Docker image of the backend, I ran into a permissions error because the `ec2-user` on my instance didn't have the necessary rights to run Docker commands directly. Docker requires root-level privileges to create containers and build images, and since Docker was installed for the root user, running Docker commands without `sudo` caused a permission denied error. This meant I either had to use `sudo` before Docker commands or update the user permissions to allow `ec2user` to run Docker commands without elevated privileges.

To solve the permissions error, I added the `ec2-user` to the Docker group by running the command `sudo usermod -a -G docker ec2-user`. The Docker group is a special user group in Linux that grants its members permission to run Docker commands without needing root (or sudo) privileges. By adding my



Darpan Patel

user to this group, I was able to execute Docker commands directly, avoiding the permissions error.

```
aws [Search] [Alt+S] United States (N. Virginia) Darpankumar-patel-IAM-Admin @ nextwork-alias-darpankumar-patel

[ec2-user@ip-172-31-95-165 nextwork-flask-backend]$ docker build -t nextwork-flask-backend .
[+] Building 10.9s (10/10) FINISHED docker:default
=> [internal] load build definition from Dockerfile 0.0s
=> => transferring dockerfile: 269B 0.0s
=> [internal] load metadata for docker.io/library/python:3.9-alpine@sha256:3881e169a5c90cadeb971e9d40110.29kB / 10.29kB 0.0s
=> [internal] load .dockerignore 0.0s
=> => transferring context: 2B 0.0s
=> [1/5] FROM docker.io/library/python:3.9-alpine@sha256:3881e169a5c90cadeb971e9d40110.29kB / 10.29kB 0.0s
=> => resolving docker.io/library/python:3.9-alpine@sha256:3881e169a5c90cadeb971e9d40110.29kB / 10.29kB 0.0s
=> => sha256:6ff759b17c58dc0211cc72756d79 1.73kB / 1.73kB 0.0s
=> => sha256:5bbf01257b0ee9b52b14a2a383bb41 5.08kB / 5.08kB 0.0s
=> => sha256:fe07694b16b82247c3539ed9eae65ff 3.80MB / 3.80MB 0.1s
=> => sha256:09ebce43b37a228ea2d338536e 460.21kB / 460.21kB 0.1s
=> => sha256:4ef74e18634db41cf0f117f2ecbd07321e7 14.89MB / 14.89MB 0.5s
=> => extracting sha256:fe07694b16b82247c3539ed9eae65ff 3.80MB / 3.80MB 0.2s
=> => sha256:427eab949a336b38da329b251cd7f0a879a 248B / 248B 0.2s
=> => extracting sha256:09ebce43b37a228ea2d338536e3bd98fab 460.21kB / 460.21kB 0.1s
=> => sha256:4ef74e18634db41cf0f117f2ecbd07321e7 14.89MB / 14.89MB 0.5s
=> => extracting sha256:427eab949a336b38da329b251cd7f0a879a 248B / 248B 0.2s
=> [internal] load build context 0.1s
=> => transferring context: 42.42kB 0.0s
=> [2/5] WORKDIR /app 0.1s
```

i-09ca4f679d44abc94 (nextwork-eks-instance)
PublicIPs: 54.86.48.236 PrivateIPs: 172.31.95.165

CloudShell Feedback © 2025, Amazon Web Services, Inc. or its affiliates. Privacy Terms Cookie preferences

Container Registry

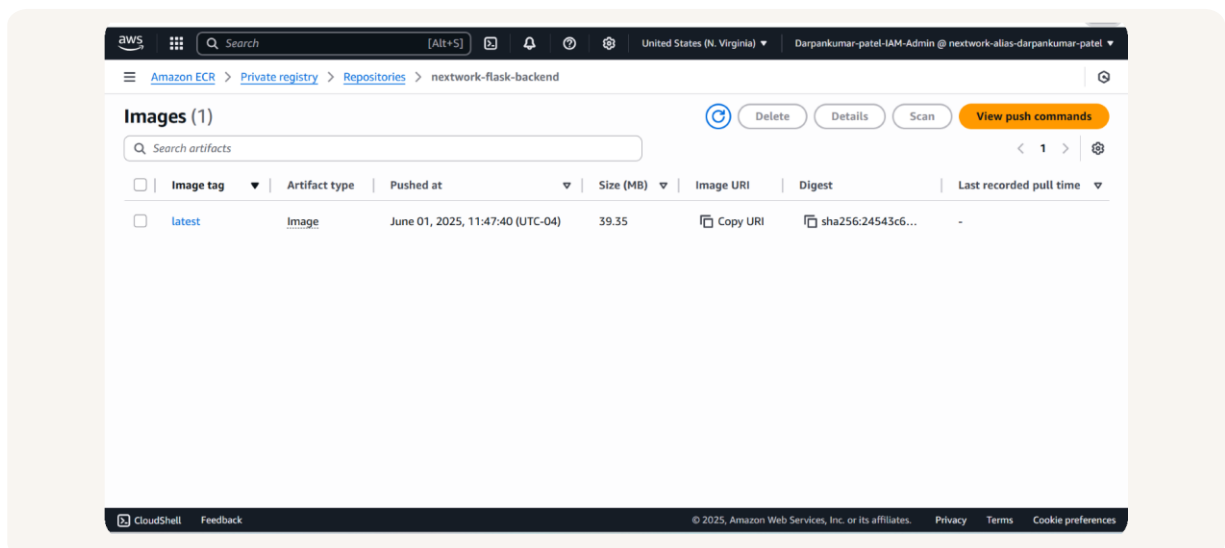
I'm using Amazon ECR in this project to store and manage the Docker container image of the backend application so that it can be pulled and deployed by Amazon EKS. ECR is a good choice for the job because it's tightly integrated with other AWS services like EKS and IAM, making authentication and image access seamless and secure. It also supports features like vulnerability scanning, version control, and high availability, which are essential for deploying scalable and secure containerized applications in a production-ready environment.

Container registries like Amazon ECR are great for Kubernetes deployment because they provide a centralized and secure location to store and manage



Darpan Patel

container images. This allows Kubernetes to pull the correct image version automatically during deployment, ensuring consistency across development, testing, and production environments. ECR also supports image tagging, making it easier to manage versions like `latest` or `v1.0`. Using a registry eliminates the need to preload images manually on each cluster node, simplifying scaling and updates. ECR integrates with AWS IAM for access control and can automatically scan images for vulnerabilities, enhancing security. It also streamlines automation in CI/CD pipelines by allowing seamless image pushing and pulling. Overall, container registries like ECR help maintain a reliable, secure, and efficient container deployment workflow with Kubernetes.



EXTRA: Backend Explained

After reviewing the app's backend code, I've learnt that the backend is a lightweight Flask application designed to serve as an API that takes a user's search topic, queries the Hacker News Search API for related content, processes the results, and returns them in a clean JSON format. It relies on a few key dependencies to handle HTTP requests, routing, and API creation, and the entire backend is packaged into a Docker container for consistent



Darpan Patel

deployment. This setup ensures the backend can efficiently fetch, process, and serve data to support the frontend or other services in a scalable and reliable way.

Unpacking three key backend files

The requirements.txt file lists all the Python dependencies that the backend application needs to run properly. It specifies the exact versions of libraries like Flask

(for building the web app), Flask-RESTx (for creating the API), Requests (for making

HTTP calls to the Hacker News API), and Werkzeug (which helps with routing in Flask). This file ensures that when the app is deployed or the Docker image is built, all necessary packages are installed consistently across environments.

placeholder

The Dockerfile gives Docker instructions on how to build a container image for the backend application. Key commands in this Dockerfile include: `FROM python:3.9alpine` which sets up a lightweight Python 3.9 environment as the base image. WORKDIR /app` to set the working directory inside the container. COPY requirements.txt requirements.txt` to copy the dependencies list into the container. RUN pip3 install -r requirements.txt` to install all necessary Python packages. COPY .`