

AWS User Guide



DuploCloud

TABLE OF CONTENTS

INTRODUCTION	4
PRE-REQUISITES	4
AWS GUIDE	5
INFRASTRUCTURE	5
TENANT	6
QUICK START – A DOCKER MICROSERVICE	7
VIRTUAL MACHINES	9
RDS DATABASE	10
ELASTIC CACHE	11
S3 BUCKET	11
ELASTIC SEARCH	12
CONTAINERS	12
SECRETS MANAGER	16
DYNAMODB	16
SQS QUEUE	16
SNS TOPIC	16
KINESIS STREAM	17
LAMBDA	17
API GATEWAY	18
WORKSPACES	18
SERVICE DISCOVERY (CLOUDMAP)	19
IoT	19
CLOUD WATCH	20
WAF	20
USING AWS CONSOLE	23
FAQ	24
AWS USE CASES	26
DOCKER WEBAPP	26
S3 BACKED WEBAPP	28
SERVERLESS WEBAPP	35
PASSING SECRETS	36
MONITORING	36

FAULTS/ALARMS	37
ALERTING/NOTIFICATION	38
CENTRAL LOGGING	39
BACKUP & RESTORE	40
DISASTER RECOVERY	41
COST MANAGEMENT.....	41
BIGDATA/ETL.....	42
SECURITY & COMPLIANCE.....	51
CLOUD MIGRATION (SMS)	51
REMOTE WORKING	51
ADDITIONAL DEPLOYMENT FUNCTIONALITIES	51
ADVANCED FUNCTIONS	51
BYOH.....	52
ADMINISTRATOR'S GUIDE	53
CONFIGURATION SCOPES.....	53
ACCESS CONTROL	57
ACCESSING CLOUD SERVER	57
<i>VPN Access</i>	<i>57</i>
SHARING ENCRYPTED DB	59
SSH EC2 INSTANCE	63
CI/CD GUIDE – KATKIT: DUPLOCLOUD'S CI/CD COMPONENT	64
FIRST DEPLOYMENT.....	65
ENVIRONMENTS.....	65
LINK REPOSITORY	66
PHASES	66
KATKIT CONFIG.....	67
ADVANCED FUNCTIONS	68

Introduction

DuploCloud is a rules-based orchestration engine. It is the subject matter expert in digital form. It is not a DevOps tool that requires a cloud expert to operate, but instead it is the cloud expert itself. The software has a resume with over 500 skills that include AWS, Azure, Kubernetes, Security and IoT among others. It allows users to build applications via a simple declarative interface and not have to deal with low level infrastructure details like VPC, security groups, IAM, Jenkins setup, etc. These constructs are still in play, but the DuploCloud software abstracts it away for you by auto-generating the configuration based on application needs.

DuploCloud is single tenant software that installs in your cloud account. Users interface with the software via the browser UI and/or API calls and all data and configuration stays within your cloud account. All configurations that have been created and applied by the software are transparently available to be reviewed and edited in your cloud account. Essentially, the software auto-generates the same automation scripts that a human being would have written manually. Here is an explainer video about the product:

[DuploCloud Product Demo \(AWS\).mp4](#)

Pre-requisites

- **Sign-in Account:** DuploCloud works off Google and Office365 Oauth. To login to DuploCloud you either need a Google or Office365 account.
- **Docker Knowledge:** If you are deploying a Docker based microservice using DuploCloud then it is assumed that you are familiar with Docker.
 - You should have a docker image for your application that has been tested locally in your computer. Make sure it runs in detached mode (i.e., docker run -d option).
 - The Image should have been pushed to your repository. If the repository is private, then you will have to set the credentials for Docker hub in your DuploCloud account.
- **AWS SDK Familiarity:** If you are using AWS services like S3, DynamoDB, SQS, and SNS you must have a basic knowledge of how these services can be consumed. The interface to create these services via DuploCloud will be very declarative and self-explanatory. You do not need any access keys in your code to access these services. Use the AWS constructor that does not take credentials, but takes only region which must be US-West-2 unless otherwise specified by your enterprise administrator or during signup.

Terminologies

- **Infrastructure:** An infrastructure maps 1:1 with a VPC/VNET and can be in any region. Each infrastructure has a set of subnets spread across multiple availability zones. In AWS there is a NAT gateway for private subnets.
- **Tenant or Project:** Tenant or Project is a sandbox or unit of deployment. All resources in a given tenant are isolated from another tenant via Security groups and IAM policies (optionally VPC). For applications to be reachable outside the sandbox a Port mapping or ELB must be created.
- **User:** This is an individual with a user ID. Each user could have access to one more tenants/projects.
- **Host:** This is an EC2 instance or VM. This is where your application will run.
- **Service:** Service is your application code packaged as a single docker image and running as a set of one or more containers. It is specified as - image-name; replicas; env-variables; vol-mappings if any. DuploCloud also allows running applications that are not packaged as Docker images.
- **ELB:** A Service can be exposed outside of the tenant\project via an ELB and DNS name. ELB is defined as - Service name + container-port + External port + Internal-or-internet facing. Optionally, a wild card certificate can be chosen for SSL termination. You can choose to make it internal which will expose it only within your VPC to other applications.
- **DNS Name:** By default, when a Service is exposed via an ELB, DuploCloud will create a friendly DNS Name. A user can choose to edit this name. The domain name must have been configured in the system by the admin.
- **Docker Host or Fleet Host:** If a host is marked as part of the fleet, then DuploCloud will use it to deploy containers. If the user needs a host for development purposes, say a test machine, then they would mark it as not part of the pool or fleet.

AWS Guide

Infrastructure

A completely new isolated environment can be created for use cases from DuploCloud. Proceed to Administrator --> Infrastructure. Click on +Add button. It opens a form where it asks to enter basic information about the new environment.

Infrastructure > Admin > Infrastructure > Add Infrastructure

Name: demoinfra ✓

Region: us-west-1 x v

VPC CIDR: 10.231.0.0/16 ✓

Cloud: AWS x v

Availability Zones: 2 x v

Subnet CIDR Mask: 24 ✓

☒ Enable EKS

Cancel Create

Once you submit the form, DuploCloud creates a new VPC with 2 subnets (private, public) in each availability zone, necessary security groups, NAT Gateway, Internet Gateway, Route tables, etc. DuploCloud also configures the VPC peering with the master VPC which is initially configured in DuploCloud. Infrastructure creation takes around 10mins to complete. Once its complete, a new plan with the same infrastructure name is automatically created and populated with the details of the newly created infrastructure. This plan can then be used to create the tenants.

Plan Details

```
1 {
2   "Name": "default",
3   "HwProvider": 0,
4   "BlockBYOHosts": false,
5   "Images": [
6     {
7       "Name": "amazon-linux",
8       "ImageId": "ami-0d6621c01e8c2de2c",
9       "OS": "linux",
10      "Tags": [],
11      "Username": "ec2-user"
12    },
13    {
14      "Name": "eks-worker",
15      "ImageId": "ami-077e0e0754c311a47",
16      "OS": "linux",
17      "Tags": [],
18      "Username": "ec2-user"
19    }
20  ],
21 }
```

Close

Beta UI Tenant: Plans

Total 4 Show

NAME
DEFAULT
FINANCE
TESTEDIT
JOETEST

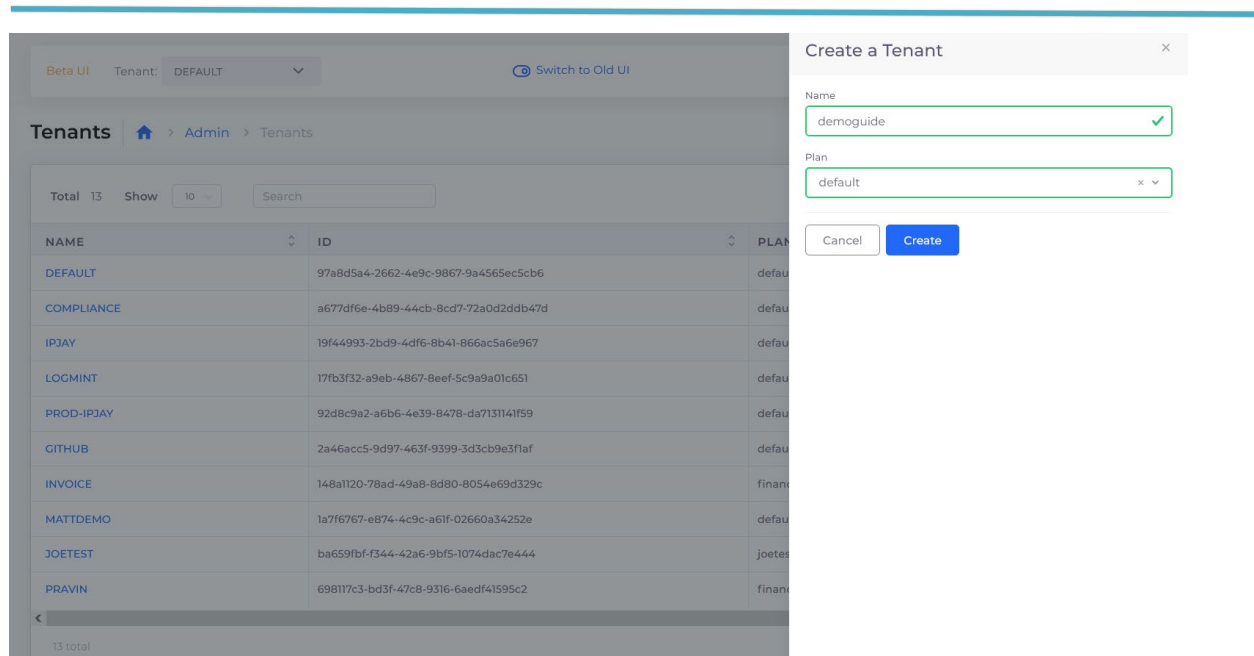
4 total

Diagnosics 16 Prasanna Administrator

TENANTS COUNT	ACTIONS
	🔗 ⓘ ⋮
	🔗 ⓘ ⋮
	🔗 ⓘ ⋮
	🔗 ⓘ ⋮

Tenant

Create a new Tenant from the already created plan. Give a min for the DuploCloud to create the Tenant specific SG and Instance profiles which acts as a security boundary.



Quick Start – A Docker Microservice

Deployment is a three-step process.

1. Create a Host (VM)
2. Deploy the service (Application)
3. Expose the application using LB

Watch this short video that shows how to package your code in a Docker image and deploy:

[DuploCloud Product Demo \(AWS\).mp4](#)

- **Create a Host** from the menu **DevOps → Hosts → EC2 → +Add** button above the table. Choose the desired instance type. The available instance types are set by your administrator or the plan you choose in DuploCloud. If you are not using this host for hosting containers, then set the pool as none. If you want a public IP for the host, then select the public subnet in the list of availability zones.

EC2 [Home](#) > [DevOps](#) > [Hosts](#) > [EC2](#) > Add Host

Add Host

Friendly Name: ✓ Instance Type:

☒ Advanced Options

Allocation Tags: ✓ Image Id:

Agent Platform: Availability Zone: Disk Size:

Enable Block EBS Optimization: Enable Hibernation:

Base64 Data: Tags: Volumes:

Network interfaces:

- **Deploy a Service (application)** from the menu **DevOps → Containers → EKS/ Native → +Add** sign above the table. Give a name for your services (no spaces); number of replicas; Docker image; volumes (if any). The number of replicas must be less than or equal to the number of hosts in the fleet.

Services [Home](#) > [DevOps](#) > [Containers](#) > [EKS/ Native](#) > [Services](#) > Add Service

Add Service

1 Basic Options [Minimal inputs to start service](#) > 2 Advanced Options [More options to configure service](#)

Service Name: Cloud: Platform:

Docker Image: Allocation Tag: Replicas:

Enable Host Network: Environment Variables:

Daemonset:

- Create an ELB from the menu **DevOps → Containers → EKS/ Native →** Select the desired service from the table and click on Load Balancer tab for Configuration. The URL suffix you specify under Health Check will be used by DuploCloud during a rolling upgrade i.e., when a service image is changed, then DuploCloud will take down one container replica at a time and bring up a new one. Once it's running, then the

DuploCloud agent running on the host will make a call to the URL and expects a 200 OK. If it does not get a 200 OK, then the upgrade is paused and the user needs to update with an image that fixes the issue.

The screenshot shows the DuploCloud interface with the 'Add Load Balancer Listener' dialog open. The main interface displays the 'nginxdemo' service under 'Services' > 'DevOps' > 'Containers' > 'EKS/Native' > 'Services'. The service is using the 'nginx:latest' image. The 'Load Balancers' tab is selected, showing 'No Load Balancers configured for this service.' and a 'Configure Load Balancer' link. The dialog on the right is titled 'Add Load Balancer Listener' and contains the following fields:

- Select Type: Classic
- Container port: 8080
- External port: 443
- Visibility: Public
- Application Mode: Docker Mode
- Health Check: /
- Backend Protocol: HTTP
- Certificates: poc.duplocloud.net

At the bottom of the dialog are 'Cancel' and 'Add' buttons.

- The DNS name for the service will be present in the Services Table. It takes about 5 to 10 minutes for the DNS name to start resolving.
- Update a Service from the menu **DevOps** → **Containers** → **EKS/Native** → click on Edit button under Actions column of the Services table or Select the desired service from the table and click on edit option under the Actions menu button.

Virtual Machines

You can create the EC2 instance from DuploCloud. Go to **DevOps** → **Hosts** → **EC2** → **+Add** button above the table. Choose the desired instance type. The available instance types are set by your administrator or the plan you choose in DuploCloud. If you are not using this host for hosting containers, then set the pool as none. If you want a public IP for the host, then select the public subnet in the list of availability zones.

EC2 [Home](#) > [DevOps](#) > [Hosts](#) > [EC2](#) > Add Host

Add Host

Friendly Name: ✓

Instance Type: x v

☒ Advanced Options

Allocation Tags: ✓

Image Id: x v

Agent Platform: x v

Availability Zone: x v

Disk Size: GB

Enable Block EBS Optimization: v

Enable Hibernation: x v

Base64 Data:

Tags:

Volumes:

Network interfaces:

RDS Database

Create a MySQL, MSFTSQL, PostGRES or Aurora DB from the menu **DevOps** → **Database** → **Relational Databases** → **+Add** button above the table.

RDS [Home](#) > [DevOps](#) > [Database](#) > [RDS](#)

Create a RDS

Rds Name: ✓

Create from Snapshot (Optional): v

User Name: ✓

User password: ✓

Rds Engine: x v

Rds Engine version (Optional): ✓

Database Size: x v

Storage size in GB: GB ✓

DB Parameter Group (Optional): v

Encryption Key (Optional): x v

Once the DB is created which takes about 5 to 10 mins then you can expand the created DB to fetch the endpoint and credentials. Use this endpoint and credentials to connect to the DB from your application running in the EC2 instance. Note that the DB is accessible only from inside the EC2 instance (including the containers running in it) in the current tenant. *The best way to pass DB endpoint, DBName and credentials to your application is through ENV variables of the Service.*

NAME	ENDPOINT	SIZE	STATUS	ENGINE	ACTIONS
DUPLODEMODB	duplodemodb.c1t3p7eyfk8n.us-west-2.rds.amazonaws.com:5432	db.t2.small	available	PostgreSQL	

Elastic Cache

Create a REDIS or Memcache from **DevOps** → **Database** → **Ecache** → **+Add** button above the table. *The best way to pass the cache endpoint to your application is through ENV variables of the Service.*

Create a Ecache

Type: REDIS

Name: test

Size: cachet2.small

Replicas: 1

Kms: duplo services-demoguide

Encryption At Transit: No Encryption

Cancel Create

S3 Bucket

Create S3 Bucket from **DevOps** → **Storage** → **S3** → **+Add** button above the table. Provide the name and select S3Bucket from the dropdown. The required permissions to access the S3 bucket from VM, Lambda functions and containers is provisioned automatically through Instance profiles and hence no Access key is required in the Application code. Application code should use the IAM role/Instance profile approach to connect to these services i.e., the AWS SDK constructor which only uses region should be used. Uploading the file and changing the bucket permissions can be done directly in the AWS console which can be accessed by clicking on the

>_ icon. Within this login the user will have enough permissions to configure the bucket, but no access or security level permissions will be given.

The image shows a user interface for creating an S3 bucket. On the right, a modal titled 'Create an S3 Bucket' is open, showing a 'Name' field with the value 'demos3' and a green checkmark, indicating it's valid. Below the field are 'Cancel' and 'Create' buttons. On the left, the main page is titled 'S3 Bucket' and shows a breadcrumb trail: Home > DevOps > Storage > S3 Bucket. Below the title, there's a table with columns 'NAME' and 'Total'. The table is currently empty, showing 'No data to display' and '0 total'.

Elastic Search

Create an Elasticsearch domain from **DevOps** → **Analytics** → **Elasticsearch** → **+Add** button above the table.

The image shows a form titled 'Create an Elasticsearch Domain' with a breadcrumb trail: Home > DevOps > Analytics > Elasticsearch > New Elasticsearch Domain. The form has two columns of fields. The left column contains: 'Domain Name' (demos), 'Elasticsearch Version' (7.10), 'Data Instance Size' (t2.small.elasticsearch), 'Data Instance Count' (1), 'Zone' (Zone-A), and 'Storage (in Gb)' (20). The right column contains: 'Master Instance Size' (t2.small.elasticsearch), 'Master Instance Count' (No Dedicated Master), 'Encryption Key' (No Encryption), and three checkboxes: 'Require SSL / HTTPS', 'Use Latest TLS Cipher', and 'Enable Node-to-Node Encryption'. All fields have green checkmarks indicating they are valid. At the top right of the form are 'Cancel' and 'Submit' buttons.

Containers

- **Built-In:** Any docker container can be deployed in the VM. Go to **DevOps** → **Containers** → **EKS/Native** → **+Add** button above the table. Give a name for your services (no spaces); number of replicas; Docker image; volumes (if any). The number of replicas must

be less than or equal to the number of hosts in the fleet. You can use AllocationTags to deploy the container in a specific set of hosts.

- **EKS:** DuploCloud supports EKS and AKS out of box. Kubernetes clusters are created during Infrastructure setup under **Administrator → Infrastructure**. The cluster is created in the same VPC as the Infrastructure. Typically, it takes 10 minutes for an Infrastructure with AKS/EKS cluster to be ready. Next, we will walk you through deploying an application within a Tenant in Kubernetes. The application will comprise of set of VMs, Deploymentset(pods) with an application load balancer. Pods can be deployed either through the DuploCloud Portal or through Kubectl(HelmCharts)
 - **Adding Tenants:** We map each tenant to a namespace in Kubernetes. For example, if a tenant is called analytics in DuploCloud, then there will be a namespace called duploservices-analytics in Kubernetes. All components of an application within this tenant are placed in this namespace. Since Nodes cannot be part of namespace in Kubernetes, we set a label called "tenantname" for all the nodes that are launched within the tenant. For example, a node launched in the analytics tenant will have label called "tenantname: duploservices-analytics". Any pods that are launched through the DuploCloud UI will have appropriate nodeselector to tie the pod to the nodes within the tenant. If you are deploying directly via kubectl then make sure your deployment has the appropriate nodeselector.
 - **Adding Hosts:** Once the tenant is created. You can go to **DevOps** and change to the particular tenant. Under the hosts menu, start by creating new Nodes(Hosts). Make sure the value of pool is set to "Eks Linux" (which will be the default).

- **Docker registry credentials & Kubernetes secrets:** These can be set **DevOps** → **Containers** → **EKS/Native** tabs. Docker registry credentials are passed to the Kubernetes cluster as `kubernetes.io/dockerconfigjson`. Other Kubernetes secrets can also be set here and referenced in your deployment later.
- **Adding Services:** When we say services, we don't mean a Kubernetes service, but we mean the DuploCloud term service (See terminology at the top). You can deploy DuploCloud services by clicking **+Add** button under the **EKS/Native** tab. Implicitly, DuploCloud converts these services either into a deployment set or a stateful set. If there are no volume mappings, then it is a deployment set, otherwise it is a stateful set. Most configs are self-explanatory, for example, Images, Replicas and environmental variables. The other advanced configuration can be provided under the other K8 Config section. The content of this section is a 1-1 mapping of the Kubernetes API. All types of configurations for deployments are stateful sets and are supported, while putting the appropriate JSON under the other K8 Config section. For example, to reference the secrets via config map, the JSON would look like the following.

```
{
  "Volumes": [
    {
      "name": "config-volume",
      "configMap": {
        "name": "game-config"
      }
    }
  ],
  "VolumesMounts": [
    {
      "name": "config-volume",
      "mountPath": "/etc/config"
    }
  ]
}
```

- **Kubectl token & config:** DuploCloud provides you JIT token(15min) to access the kubectl cluster. You can get the temporary token and the k8 cluster URL from the Infrastructure you created. Select the **Infrastructure** → **EKS** tab. Once you have the token and URL run these commands in your local.

```
kubectl config --kubeconfig=config-demo set-cluster EKS_CLUSTER -
-server=[EKS_API_URL] --insecure-skip-tls-verify
```

```
kubectl config --kubeconfig=config-demo set-credentials tempadmin
--token=[TOKEN]
```

```
kubectl config --kubeconfig=config-demo set-context EKS --
cluster=EKS_CLUSTER --user=tempadmin --namespace=duploservices-
[TENANTNAME]
```

```
export KUBECONFIG=config-demo
```

```
kubectl config use-context EKS
```

With these commands you have configured kubectl to point and access the k8 cluster. You can now run the commands to apply the deployment templates. `kubectl apply -f nginx.yaml`

```
apiVersion: apps/v1
kind: Deployment
metadata:
  name: nginx-deployment-g
  labels:
    app: nginx-deployment-g
spec:
  replicas: 1
  selector:
    matchLabels:
      app: nginx-deployment-g
  template:
    metadata:
      labels:
        app: nginx-deployment-g
    spec:
      nodeSelector:
        tenantname: "duploservices-stgeast1"
      containers:
        - name: nginx
          image: nginx:latest
          ports:
            - containerPort: 80
```

If you want longer duration tokens you can create them on your own and secure them out of band from DuploCloud.

- **Services display:** Once the deployment command runs successfully, go to the **services** tab of the Tenant. You can see those deployments. Now you can attach the load balancers for the services.

- **Load balancers, WAF and other workflows:** These all remain the same. See above on creation of these.
- **ECS Fargate:** Documentation TBD. Please contact the DuploCloud team for assistance.

Secrets Manager

Create any secrets under **DevOps → App Integration → Aws Secrets → +Add** button above the table.

The screenshot shows the AWS Secrets Manager console. On the left, there's a sidebar with 'Beta UI', 'Tenant: DEMOGUIDE', and a 'Switch' button. The main area shows the 'Aws Secrets' section with tabs for SNS, SQS, and Aws Secrets. Below the tabs is a table with columns for 'NAME' and a '0 total' count. On the right, a 'Create a Aws Secret' dialog is open. It has a 'Name' field with 'demoscret', a 'Secret value (Json or plain text)' field with 'JSON Key/Value pairs', and an 'Edit Json (key value pair)' field with '1'. There are 'Cancel' and 'Create' buttons at the bottom of the dialog.

DynamoDB

Create DynamoDB from under **DevOps → Database → DynamoDb tab → +Add** button above the table. The required permissions to access the DynamoDB from VM, Lambda functions and containers are provisioned automatically through Instance profiles and hence no Access key is required in the Application code. Application code should use the IAM role/Instance profile approach to connect to these services i.e., the AWS SDK constructor which only uses region should be used. More detailed configuration of the DynamoDB can be done directly in the AWS console which can be accessed by clicking on the Console >_ icon. Within this login the user will have enough permissions to configure the desired application specific details of the DynamoDB. But no access or security level permissions will be given.

SQS Queue

Create SQS Queue from under **DevOps → App Integration → SQS tab → +Add** button above the table. Refer to DynamoDB to know more about the permissions.

SNS Topic

Create SNS Topic from under **DevOps → App Integration → SNS tab → +Add** button above the table. Refer to DynamoDB to know more about the permissions.

Kinesis Stream

Create Kinesis Stream from under **DevOps → Analytics → Kinesis tab → +Add** button above the table. Refer to DynamoDB to know more about the permissions.

Lambda

Create Lambda function from under **DevOps → Serverless → Lambda tab → +Add** button. Give a name for the Lambda function and other values. This will create the lambda function. Click on AWS console to go to the AWS console for this function. Test the function. You can look at the tutorial above for the same or look at the AWS documentation.

The screenshot shows the 'Create a Lambda Function' interface in the AWS console. The breadcrumb navigation is 'DevOps > Serverless > Lambda Function > Create'. The interface is divided into two main sections. The left section contains input fields for 'Name' (filled with 'demolambda'), 'Description', 'Runtime' (a dropdown menu showing 'NodeJs-12'), 'Memory (MB)' (filled with '128'), 'Timeout (Seconds)' (filled with '2'), and 'Function Handler' (a text field with 'Function Handler Name'). The right section contains 'Environment Variables (Optional)' (a text area with a JSON snippet: `{ "Variables": {} }`), 'S3 Bucket' (a dropdown menu with 'Choose S3 Bucket'), and 'Function Package' (a text field with 'S3 Bucket Function Package'). At the top right of the form are 'Cancel' and 'Submit' buttons.

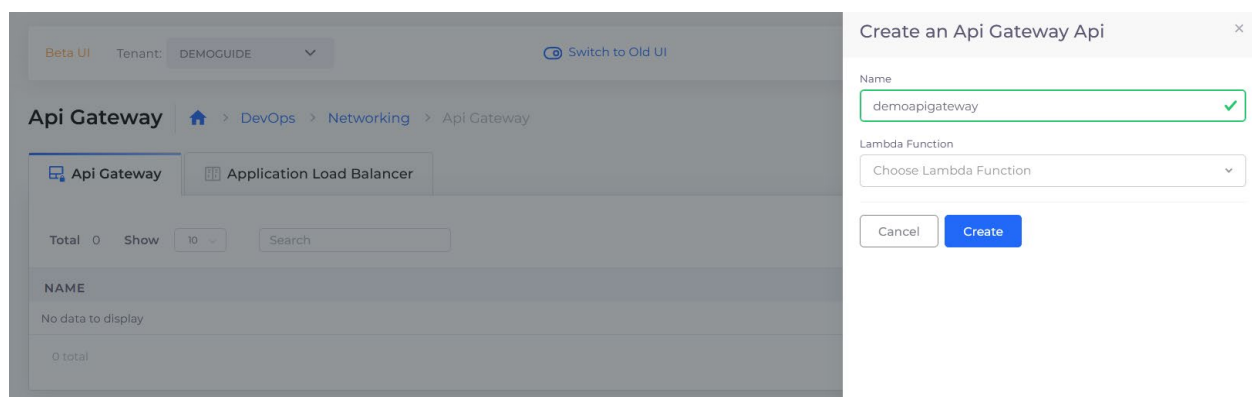
- **Updating Lambda Function and configuration:** To update the code for the Lambda function, create a new package with a different name and upload in S3 bucket again. Then select the lambda function in the table and click Edit under dropdown menu in Actions column. Make sure the right S3 bucket has been selected and provide the name of the function. To update a function configuration like timeout memory, etc. use the edit configuration button.
- **Integrating with other resources:** DuploCloud enables you to create a classic micro-services based architecture where your Lambda function can integrate with any other

resource within your tenant like S3, Dynamo, RDS or other Docker based microservices. DuploCloud will implicitly enable the Lambda function to communicate with other resources but block any communication outside the tenant (except ELB).

- **Triggers and Event Sources:** To setup a trigger or event source, the resource needs to be created via the DuploCloud Portal. Subsequently, one could trigger directly from that resource to the Lambda function in the AWS console menu of your Lambda function. Resources could be S3 buckets, API Gateway, DynamoDB, SNS etc. An example trigger via API Gateway was described in the tutorial.
- **Passing Secrets:** Passing secrets to a Lambda function can be done in much the same way as passing secrets to your Docker based service i.e., using environmental variables. For example, you can create a SQL database from the **Database → RDS** menu in DuploCloud, provide username and password, then in the Lambda menu give the same username and password. No secrets need to be stored anywhere outside like vault, git repo, etc.

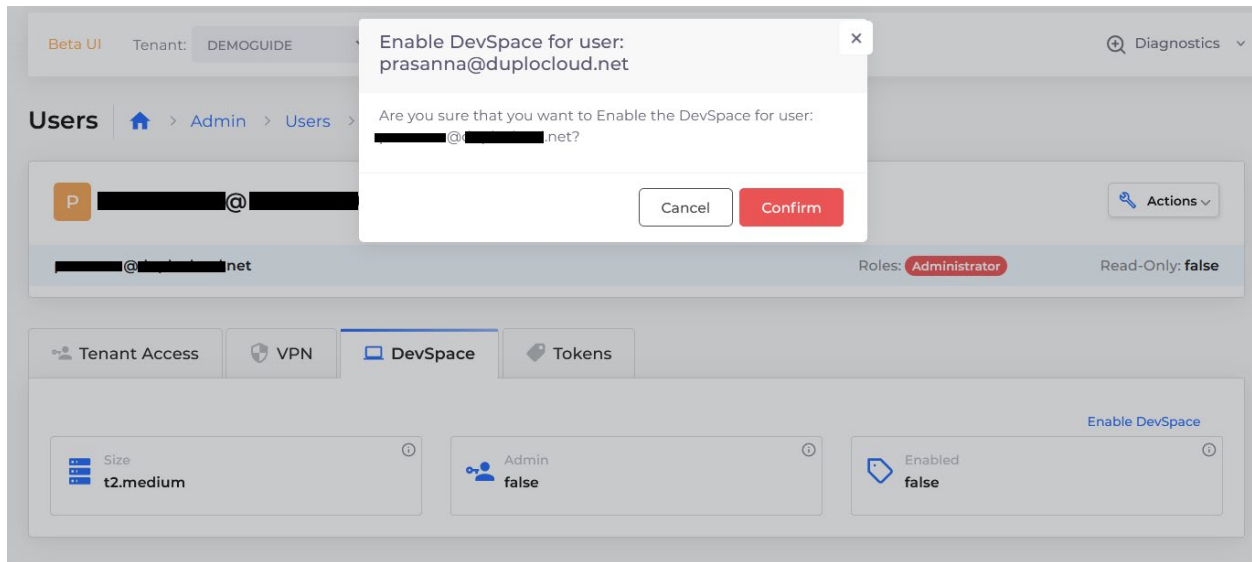
API Gateway

AWS ApiGateway RestApi is created from the DuploCloud Portal which will take care of creating the security policies to make the API Gateway accessible to other resources (like Lambda function) within the tenant. From the DuploCloud Portal only create the RestApi. All other configuration for the API (like defining methods, resources and pointing to lambda functions) should be done in the AWS console. The console for the API can be reached by selecting the API in **DevOps → Networking → API Gateway** table and then clicking on the Console button under the Actions menu. The above tutorial displays the integration of API Gateway and lambda function.

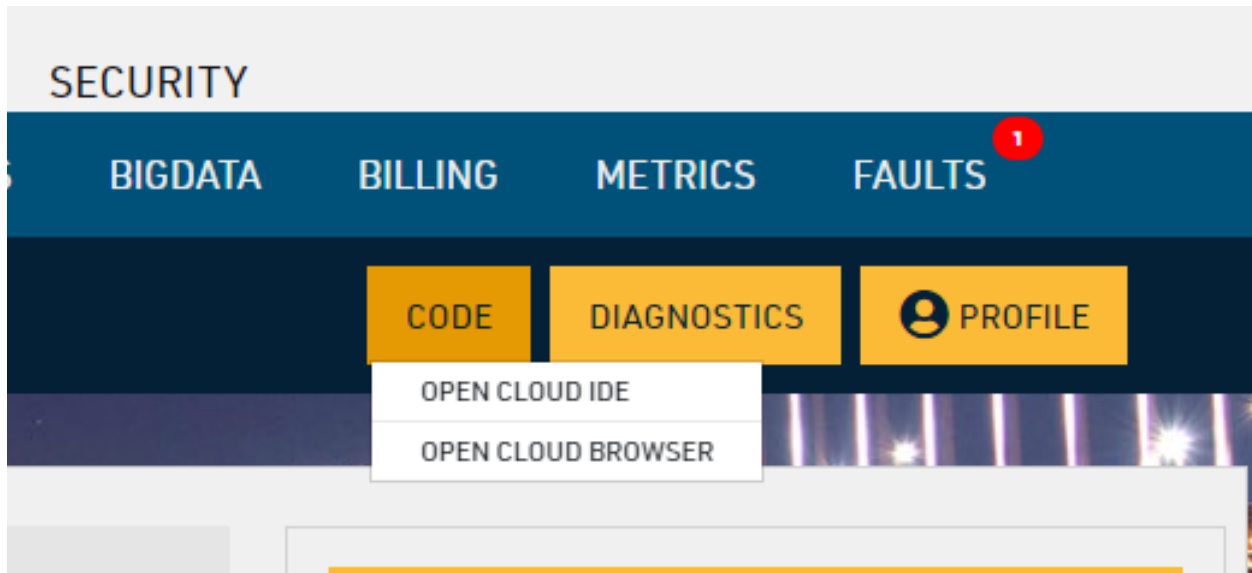


Workspaces

Individual workspaces/Devspaces can be created for the user in DuploCloud. These workspaces will have access to all the resources in the tenant which user has access. To create a workspace go to **Administrator → Users → Select user → Devspace** tab.



Once the workspace is created, it can be accessed from the Code next to Diagnostics.



Service Discovery (CloudMap)

Documentation TBD. Please contact DuploCloud team for assistance.

IoT

Documentation TBD. Please contact DuploCloud team for assistance.

Cloud Watch

By default, the metrics for a resource are available in the Metrics page in DuploCloud. But some metrics need agents to be installed in the system to collect the information, such as AWS Agent. DuploCloud provides a way to automatically install these agents on all the hosts whenever they are provisioned. For more information, refer to <https://portal.duplocloud.net/compliance/Implementation.html#agent-modules>

WAF

Creation of a Web Application Firewall is a onetime process. Create a WAF in the AWS Console, fetch the ID/ARN and update the plan in DuploCloud. Once the plan has the WAF it can be used attached to the LB.

Plans | [Admin](#) > [Plans](#) > Default

DEFAULT

[Actions](#)

default

Cloud: AWS

Region: Default Region

[Tenants](#)
[Images](#)
[Certificates](#)
[Config](#)
[WAF](#)
[KMS](#)
[Capabilities](#)

Total 2

Show 10

Search

+ Add

NAME	ID	ACTIONS
ApplicationWAF	arn:aws:wafv2:us-west-2:128329325849:regional/webacl/ApplicationWAF/db5e173c-2bbd-430f-9e16-10685bf0814b	:
owasp	ca81f828-1ea1-4a77-b413-734978d197c3	:

2 total

Region
Default Region

Tenant Count
11

VPC
10.20.0.0/16

Certificates
4

WAF
2

Beta UI

Tenant: DEMOGUIDE

Switch to Old UI

[Containers](#)
[Load Balancers](#)
[Configuration](#)
[LB Metrics](#)

DNS Name
 ngxindemo-demoguide.poc.duplocloud.net
[Edit](#)

LB Configuration
 Visibility: External
 Type: Classic
 DNS Name: DUPLO-demoguide-11474633081.us-west-2.elb.amazonaws.com
[Console](#)

Certificate
 arn:aws:acm:us-west-2:128329325849:certificate/39834e5e-0bc6-44b0-96fe-66e7dcd609e1

LB Listeners

PROTOCOL	EXT PORT	HOST PORT	CONTAINER	IS NATIVE	HEALTH
http	443	10100	8080	No	/

1 total

Set WAF for: ngxindemo

Load Balancer ARN
 WAF (Web ACL)
 Platform

☐ EnableAccessLogs
☐ Drop Invalid Headers

[Cancel](#)
[Update](#)

DuploCloud also provides a WAF Dashboard through which you can analyze the traffic that is coming in and the requests that are blocked. The Dashboard can be accessed from **Security → WAF**.



ADMINISTRATOR



DEVOPS



CI/CD



SECURITY



WAF



SIEM



Agents

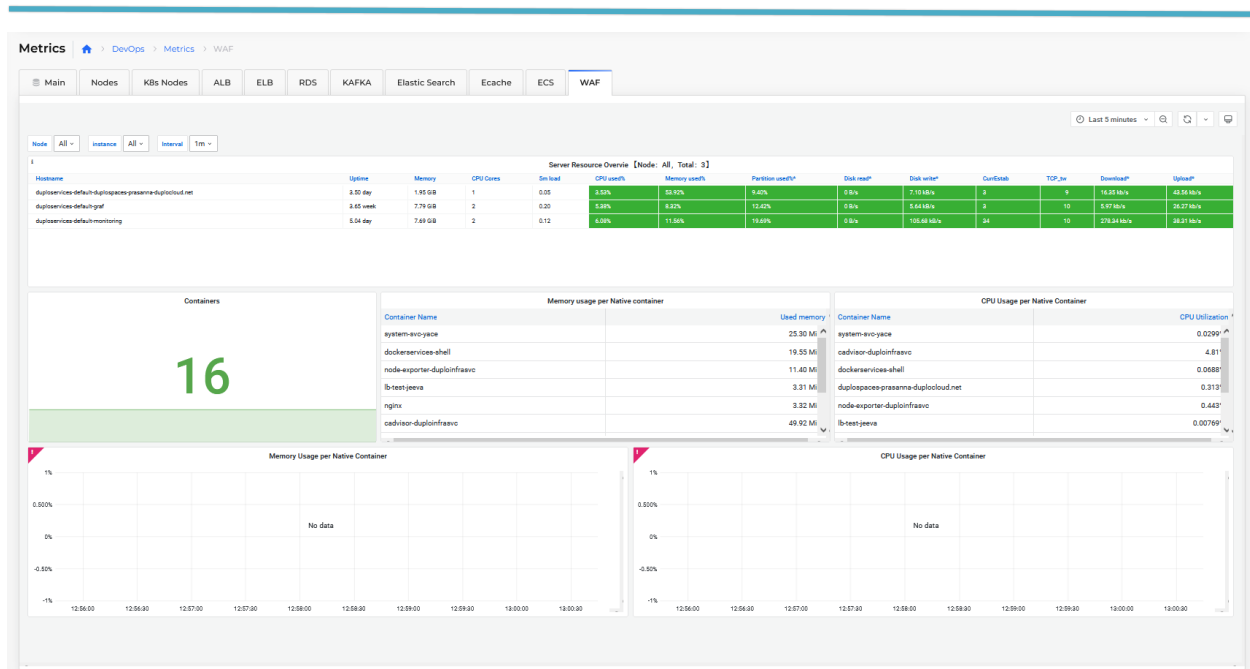


Faults



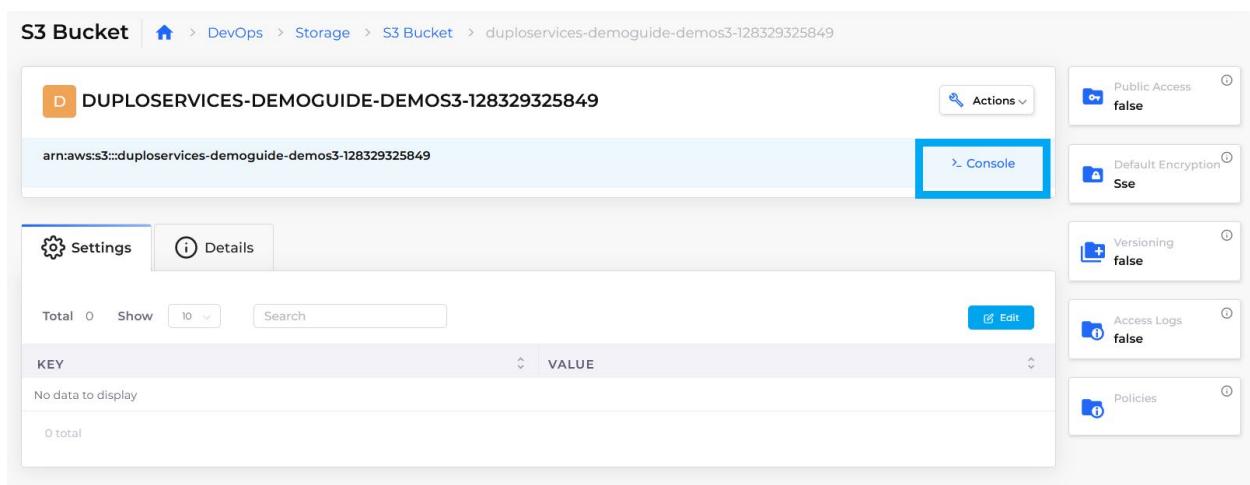
USER





Using AWS Console

DuploCloud users get AWS console access for advanced configurations of S3, Dynamo, SQS, SNS Topic, Kinesis stream, and API Gateway resources that were created through DuploCloud. ELB and EC2 parts of console is not supported yet. You can click on the AWS console icon to take you directly to that resource in AWS console with permissions scoped to the current tenant/project. You don't have permissions to create new resources directly from the AWS console. You need to do them in DuploCloud. You can do all operations on resources already created via DuploCloud. For example, you can create an S3 bucket from the DuploCloud UI and go to AWS console to add remove files, setup static web hosting etc. Similarly, you can create a DynamoDB in DuploCloud and use AWS console to add/remove entries in the table.



FAQ

- **How do I ssh into the host?**

Under each host you can click on Connection details under Action dropdown which will provide the key file and instructions to ssh.

- **My host is windows how do i RDP?**

Under host click on Connection details under Action dropdown, it will provide the password and instructions to rdp.

- **How do I get into the container where my code is running?**

Under the services Status TAB find the host where the container is running. Then ssh into the host (see instructions above) and then run `"sudo docker ps"` and get ID the container. Then run `"sudo docker exec -it containerid bash"`. You can tell which one is your container by using the image id. Don't forget the sudo in docker commands

- **I cannot connect to my service URL, how do I debug?**

Make sure the DNS name is resolving by running in your computer `"ping "`. Then we need to check if the application is running by testing the same from within the container. Then ssh into the host and then connect to your docker container using `docker exec` command (see above). From inside the container CURL the application URL using the IP 127.0.0.1 and port where the application is running. Confirm that this works. Then CURL the same URL using the IP address of the container instead of 127.0.0.1. The IP address can be obtained by running `ifconfig` command in the container.

If the connection from within the container is fine, then exit from the container into the host. Now CURL the same endpoint from the host i.e., using container IP and port. If this works then under the ELB UI in DuploCloud note down the host port that DuploCloud created for the given container endpoint. This will be in the range 10xxx or the same as container port. Now try connecting to the "HostIP" and `DuploMappedHostPort` just obtained. If this works as well but the service URL is still failing, contact your enterprise admin or duplolive-support@duploccloud.net.

- **What keys should I use in my application to connect to the AWS resources I have created in DuploCloud like S3, Dynamo, SQS, etc.?**

You don't need any. Use the AWS constructor that takes only the region as the argument (us-west-2). DuploCloud setup links your instance profile and the resources. The host in DuploCloud already has access to the resources within the same tenant\project.

DuploCloud AWS resources are reachable only from DuploCloud Hosts on the same account. YOU CANNOT CONNECT TO ANY DUPLOCLOUD AWS RESOURCE FROM YOUR LAPTOP.

- **What is rolling upgrade and how do I enable it?**

If you have multiple replicas in your service i.e., multiple containers then when you update your service like change an image or env variable then DuploCloud will make this change one container at a time i.e., it will bring down the first container and bring up the new one on that host with the updated config. If the new container fails to start or the health check URL set by the user is not returning 200 OK, then DuploCloud will pause the upgrade of remaining containers. A user must intervene by fixing the issue typically by updating the service with a newer image that has a fix. If no health check URL is specified, then DuploCloud only checks for the new container to be running. To specify health check, go the ELB menu and you will find the health check URL suffix.

- **I want to have multiple replicas of my mvc service, how do I make sure that only one of them runs migration?**

Enable health check for your service and make sure that the API does not return 200 OK till migration is done. Since DuploCloud waits for health check to be complete before upgrading the next service it is guaranteed that only one instance runs migration.

- **One or more of my containers are showing in a pending state, how can I debug?**

If it is Pending when the desired state is Running, then the image is being downloaded so wait a few minutes. If it's been more than five minutes check the faults from the button below the table. Check if your image name is correct and does not have spaces. Image names are case sensitive so it should be all lower case i.e., the image name in DockerHub should also be lower case.

If the current state is Pending when desired state is Delete then it means this container is the old version of the service. It is still running as the system is in rolling upgrade and the previous replica is not successfully upgraded yet. Check the faults in other containers of this service for which the Current State is Pending with Desired State as Running...

- **Some of my container status says "pending delete" what does it mean?**

This means DuploCloud is wanting to remove these containers. The most common reason for this is that another replica of the same service was upgraded but is now not operational. Hence DuploCloud has blocked the upgrade. You might see the other replicas in even "Running" state, but it is possible that health check is failing and that's why the rolling upgrade is blocked. To unblock, fix the service configuration (image, env, etc.) to an error free state.

- **How to create host with public IP?**

While creating host, click on **show advanced** and select the public subnet in the list of availability zone.

AWS Use Cases

Docker Webapp

In this demo, we will deploy a simple Hello World NodeJS web app. DuploCloud pulls Docker images from Docker Hub. You can choose a public image or provide credentials to access your private repository. For the sake of this demo, we will use a ready-made image available on Duplo's repository on Docker Hub.

- Login to your DuploCloud console.
- Select **DevOps → Hosts**, a Host is the instance in which your Docker container will run. You should choose a host with appropriate processing capacity for your application.
- Click on the **+Add** button to choose your host. Fill out the advanced options form if required and click submit.

The screenshot shows the 'Add Host' form in the DuploCloud console. The breadcrumb navigation is 'EC2 > DevOps > Hosts > EC2 > Add Host'. The form includes the following fields and values:

- Friendly Name:** demohost
- Instance Type:** 1 CPU 2 GB - t2.small
- Allocation Tags:** Allocation Tags
- Image Id:** Duplo-Docker-ul8
- Agent Platform:** Linux Docker/Native
- Availability Zone:** Zone A
- Disk Size:** 0
- Enable Block EBS Optimization:** No
- Enable Hibernation:** No
- Base64 Data:** 1
- Tags:** 1
- Volumes:** 1
- Network interfaces:** 1

At the bottom of the form, there are 'Cancel' and 'Add' buttons.

- You should now see your Host present in the table. Please give it a moment to instantiate.

- Next, we can create a Service. A Service is nothing but a container with user specified image and environment variables. Let's go ahead and click **+Add** button to create a new service.
- Name the service "Test-service". For this demo we will use the latest, nodejs-hello image from Duplo's public Docker hub repository. Fill in "duplocloud/nodejs-hello:latest" in the Docker Image field.
- Enter the desired number of replicas you want in the swarm. Please note that each replica runs in an individual Host, so the number of replicas must equal the number of Hosts. For the sake of this demo, we will choose 1.
- Fill in the desired environment variables, this is ideal for credentials or application specific configurations.
- Volume mapping is super easy, simply give the host path and container path as shown. Please note that we highly recommend keeping the Hosts stateless and using S3 for static assets. We will keep this field empty for this demo.

Services [Home](#) > [DevOps](#) > [Containers](#) > [EKS / Native](#) > [Services](#) > Add Service

Add Service

1 Basic Options
Minimal inputs to start service

2 Advanced Options
More options to configure service

Service Name ✓

Cloud x v

Platform x v

Docker Image ✓

Allocation Tag ✓

Replicas ✓

Docker Networks x v

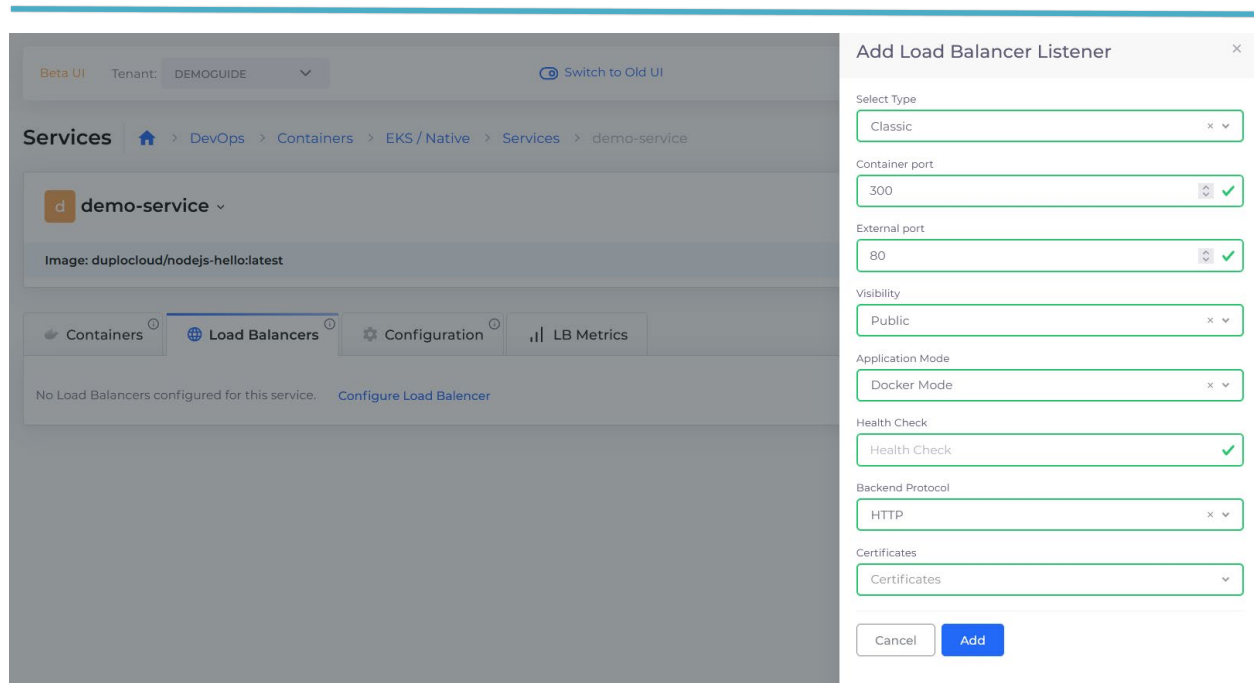
Daemonset v

Environment Variables

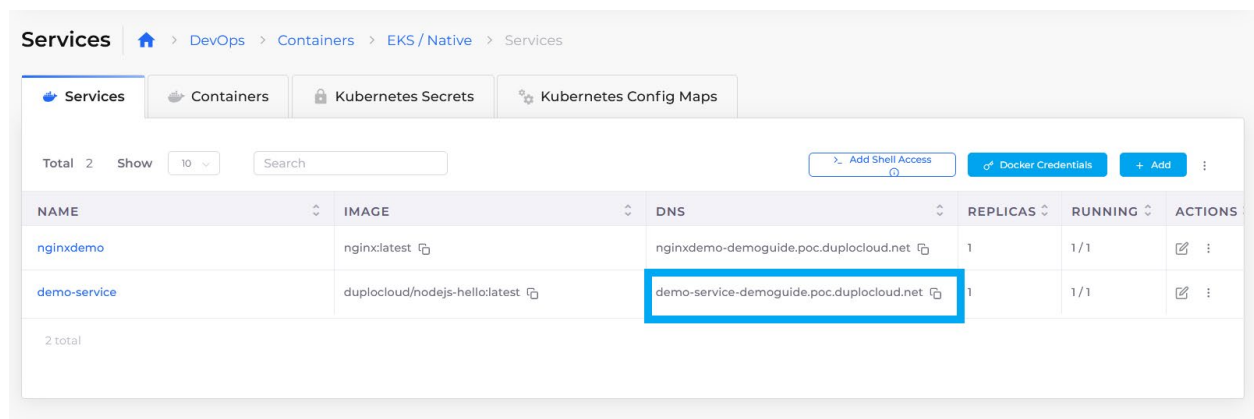
Previous

Next

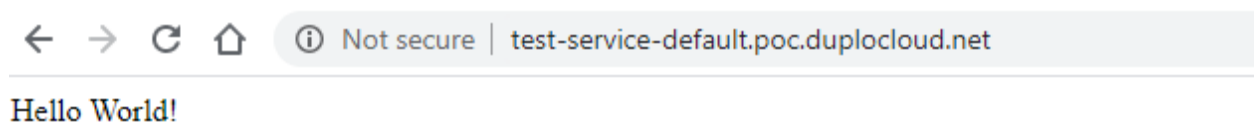
- Hit Create! Please wait a moment for the service to initialize.
- Almost there. Since the hello-nodejs image serves on port 3000 we need to create a load balancer (LB) configuration to map external port (LB) to internal port (container).
- Select the Test-service and click on **configure load balancer** on the load balancer tab. Fill the menu as shown below and click submit. This will also create a DNS name that we can use.



- Please wait for ~5 minutes as it can take a while for the DNS Route table changes to be reflected.



- Expand Test-service and copy the URL displayed under DNS. Hit the URL in the browser and You should see the Hello world serving our welcome message.



- Congratulations! You have just launched your first web service on Duplo!

S3 Backed Webapp

In the last tutorial we deployed a NodeJS based webserver and accessed it using the DNS name that DuploCloud created for us. In this tutorial we will take it a step further. We will create a text file with a message and upload it to S3. Next, we will modify our NodeJS application to access this file and display the message to every visitor. This purpose of the tutorial is to familiarize yourself with how easy it is to manage resources on AWS using DuploCloud.

- The code in the Docker container that we used in Part 1 can be found on this repo. It's a simple NodeJS web server. Clone this repository and make the following changes.
- Current app.js should look like this:

```
1 var express = require('express')
2 var app = express()
3
4 app.get('/', function (req, res) {
5   res.send('Hello World!')
6 })
7
8 app.listen(3000, function () {
9   console.log('Example app listening on port 3000!')
10 })
```


app.js hosted with ❤️ by GitHub

- Let's make the following changes

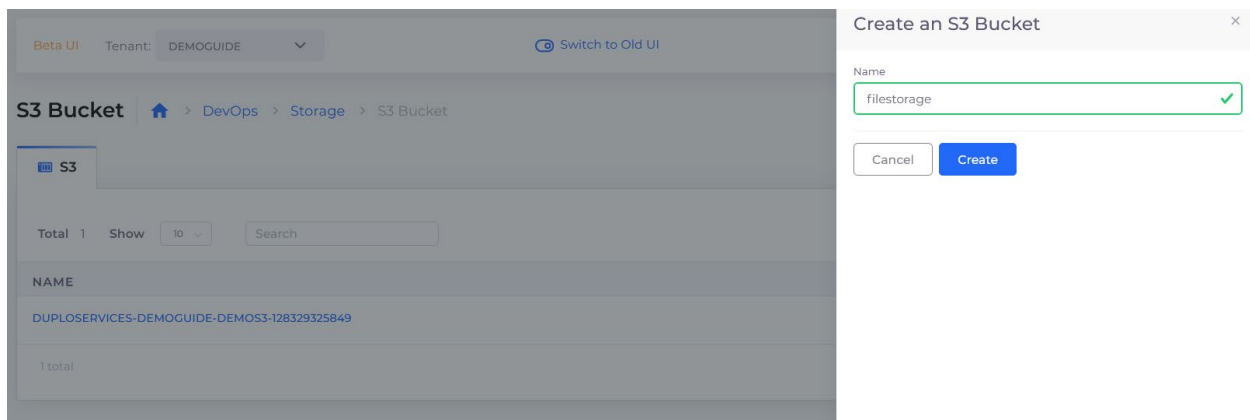
```
1 var express = require('express')
2 var app = express()
3 var request = require('request');
4 var message = '';
5 var link = '';
6
7 request(link, function (error, response, body) {
8   console.log('error:', error); // Print the error if one occurred
9   console.log('statusCode:', response && response.statusCode); // Print the response status code if a response was received
10  console.log('body:', body); // Print the HTML for the Google homepage.
11  message = body;
12 });
13
14 app.get('/', function (req, res) {
15   res.send('Hello ' + message)
16 })
17
18 app.listen(3000, function () {
19   console.log('Example app listening on port 3000!')
20 })
```

- In the above code we added a link variable that will point to the S3 bucket that we will create shortly, message variable is the message the text file in the S3 bucket will contain.
- We can now go ahead and create a text file with the message "My Name is Pranay",

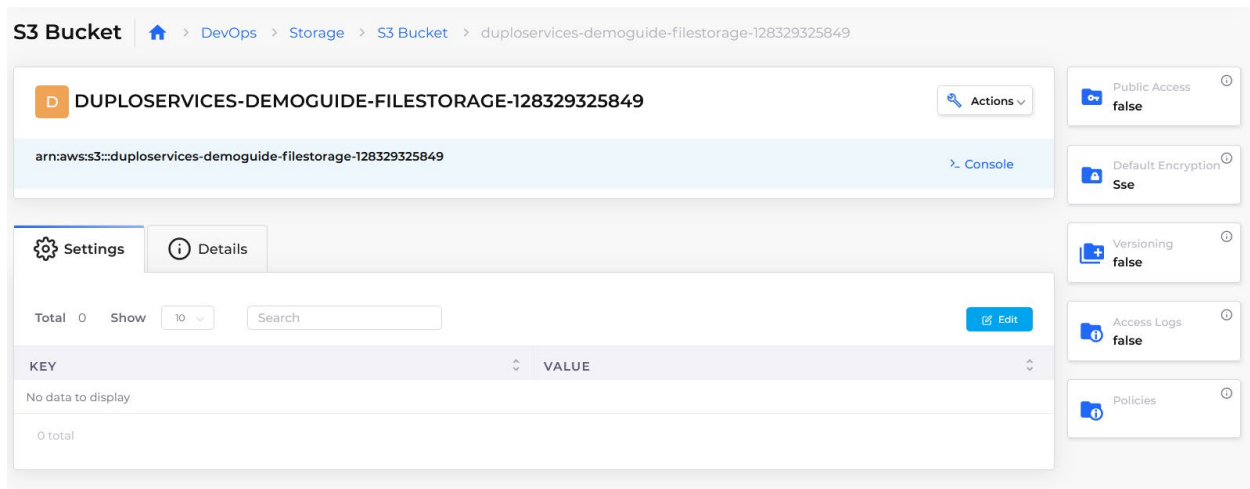
- Now that we have this file, we'll upload it to S3. Login to your DuploCloud Console and Navigate to **DevOps → Storage → S3**.



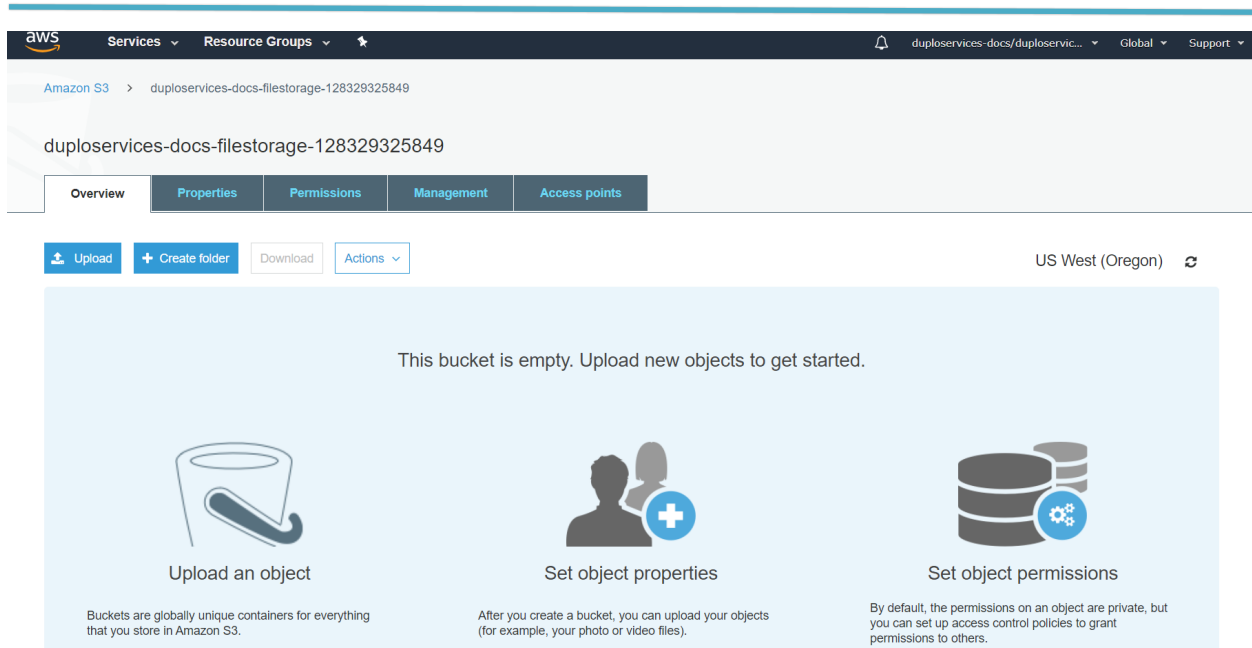
- Click **+Add** button to create a S3 bucket, provide the resource name and click on Create.



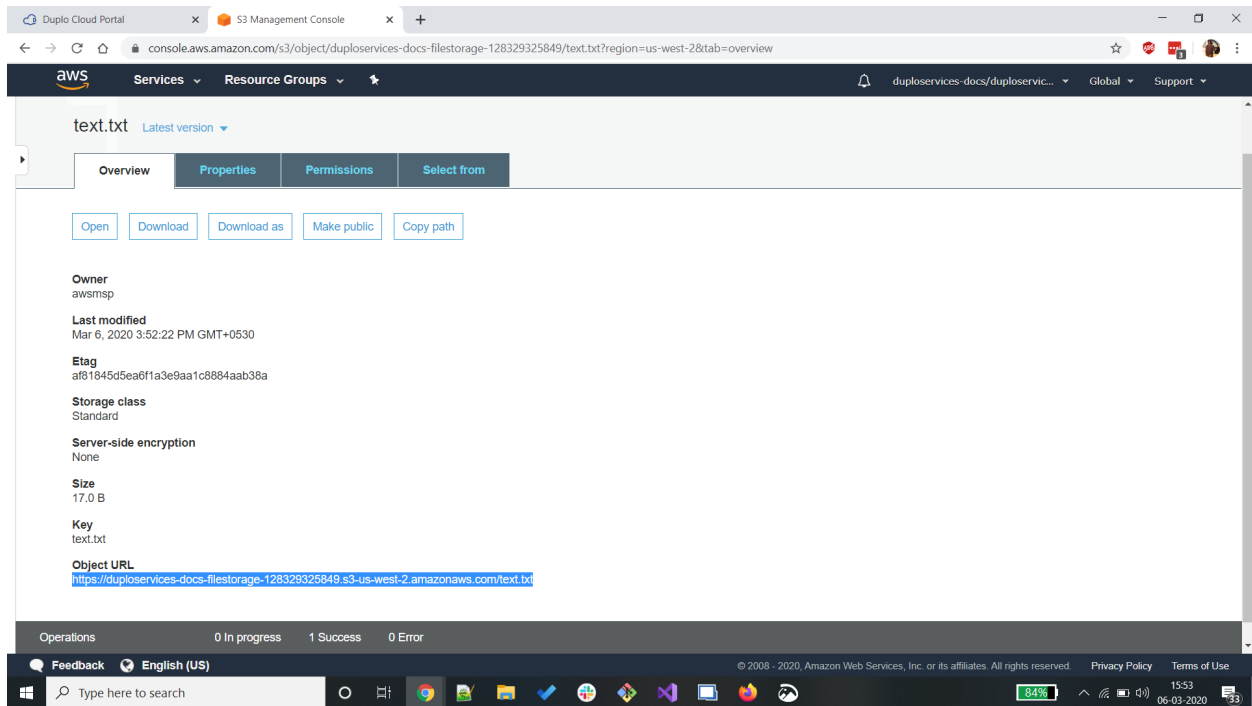
- This will create a S3 bucket



- Now click on the AWS Console button under Actions menu, this will take us directly to the S3 bucket in AWS console with limited access. Here, we can upload the file and modify permissions. Select Upload and choose the text file that we created earlier. Hit next till the end and the file should be uploaded.



- Click on the file name and copy the Link for the file in the bucket.



- Now let's assign this link to the link variable in app.js

```
1 var express = require('express')
2 var app = express()
3 var request = require('request');
4 var message = '';
5 var link = 'https://s3-us-west-2.amazonaws.com/duplocloudservices-670rp5-ac42540e-1401-49ee-939e-0c0e8c474263/myfile.txt';
6
7 request(link, function (error, response, body) {
8   console.log('error:', error); // Print the error if one occurred
9   console.log('statusCode:', response && response.statusCode); // Print the response status code if a response was received
10  console.log('body:', body); // Print the HTML for the Google homepage.
11  message = body;
12 });
13
14 app.get('/', function (req, res) {
15   res.send('Hello ' + message)
16 })
17
18 app.listen(3000, function () {
19   console.log('Example app listening on port 3000!')
20 })
```

- Almost there, all we need to do now is create a Docker file and push to Docker Hub so that DuploCloud can access it.
- I have created a simple dockerfile inspired from this Blog from NodeJS.

```
1 FROM node:boron
2
3 # Create app directory
4 RUN mkdir -p /usr/src/app
5 WORKDIR /usr/src/app
6
7 # Install app dependencies
8 COPY package.json /usr/src/app/
9 RUN npm install
10
11 # Bundle app source
12 COPY . /usr/src/app
13
14 EXPOSE 3000
15 CMD [ "nodejs", "app.js" ]
```

- We will now run the Docker build command to build out an image from this dockerfile.



- And now Docker push to push it to Docker Hub



- All set! We will create a new service as we did in the last tutorial to run this new Docker container.
- In the services tab click on the plus sign and fill the options as specified.

Services [Home](#) > [DevOps](#) > [Containers](#) > [EKS / Native](#) > [Services](#) > [Add Service](#)

Add Service

1 Basic Options
Minimal inputs to start service

2 Advanced Options
More options to configure service

Service Name

Cloud

Platform

Docker Image

Allocation Tag

Replicas

Docker Networks

Daemonset

Environment Variables

← Previous

Next →

- It's important to note here that this service along with the previous service will run in the single host that we created in the previous tutorial! All we need to do is create a new load balancer for this service so that we can access it.
- Click on configure Load balancer button under Load Balancer tab of the service and fill out the form.

Beta UI Tenant: DEMOGUIDE [Switch to Old UI](#)

Services [Home](#) > [DevOps](#) > [Containers](#) > [EKS / Native](#) > [Services](#) > test-service-s3

t test-service-s3

Image: duplocloud/nodejs-hello-world-s3:latest

Containers

Load Balancers

Configuration

LB Metrics

No Load Balancers configured for this service. [Configure Load Balancer](#)

Add Load Balancer Listener

Select Type

Container port

External port

Visibility

Application Mode

Health Check

Backend Protocol

Certificates

Cancel

Add

- As before, give it a minute and hit the freshly generated link in the browser. Voila!



- So now we have two services which we can access via two different load balancers running on a single host (EC2 instance)! Due to this tight coupling with S3 we can do all sorts of fun stuff like store static files as well as javascript in S3.

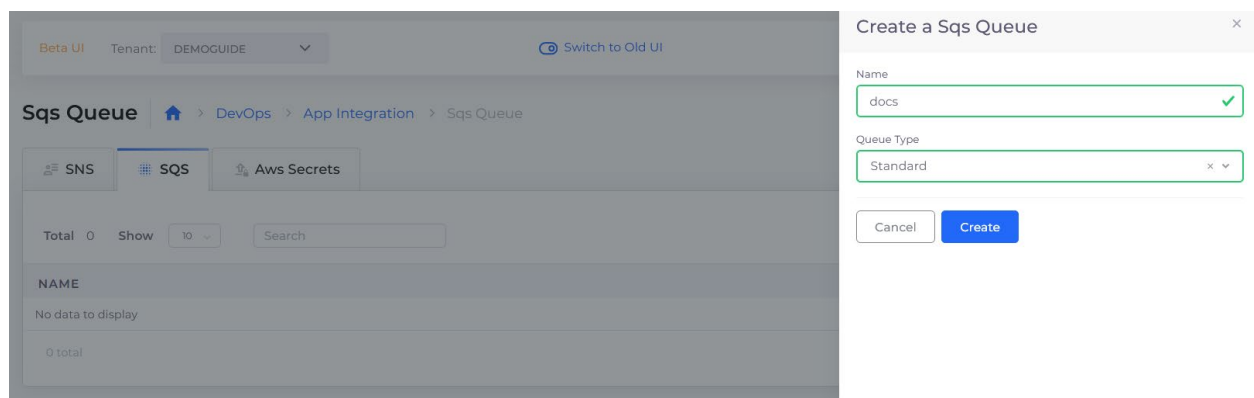
Serverless Webapp

DuploCloud makes deploying serverless applications a breeze. A step by step tutorial on how to deploy a Lambda application is at <https://www.linkedin.com/pulse/deploy-existing-web-applications-serverless-using-aws-thiruvengadam/>

Following is a video version of the tutorial: <https://youtu.be/MW62CrzVO2E>

Deployment is a three-step process.

- **Create a Zip file:** Generate a Zip package of your Lambda code. The Lambda function code should be at the root of the package. If you are using virtual env, all dependencies should be packaged. Refer to the AWS documentation at <http://docs.aws.amazon.com/lambda/latest/dg/lambda-python-how-to-create-deployment-package.html> for detailed instructions on how to generate a package. We personally prefer using tools like zappa and serverless.
- **Create an S3 bucket and upload the Zip file: DevOps → Storage → S3 → +Add** button above the table. Give a name for your bucket or leave it blank and a name will be auto generated. Then Click on AWS Console Button to get into AWS console for this S3 bucket and upload the zip file we just created.



- **Create Lambda Function: DevOps → Serverless → Lambda → +Add** button above the table. Give a name for the Lambda function and other values. This will create the lambda function. Click on AWS console to go to the AWS console for this function. Test the function. Refer to Lambda for more info. You can also look at the tutorial above for the same or look at AWS documentation.

Lambda Function [DevOps](#) > [Serverless](#) > [Lambda Function](#) > Create

Create a Lambda Function Cancel Submit

Name: ✓

Description:

Runtime: x v

Memory (MB): v

Timeout (Seconds): v

Function Handler:

Environment Variables (Optional)

```
1 {  
2   "Variables": {}  
3 }
```

S3 Bucket: v

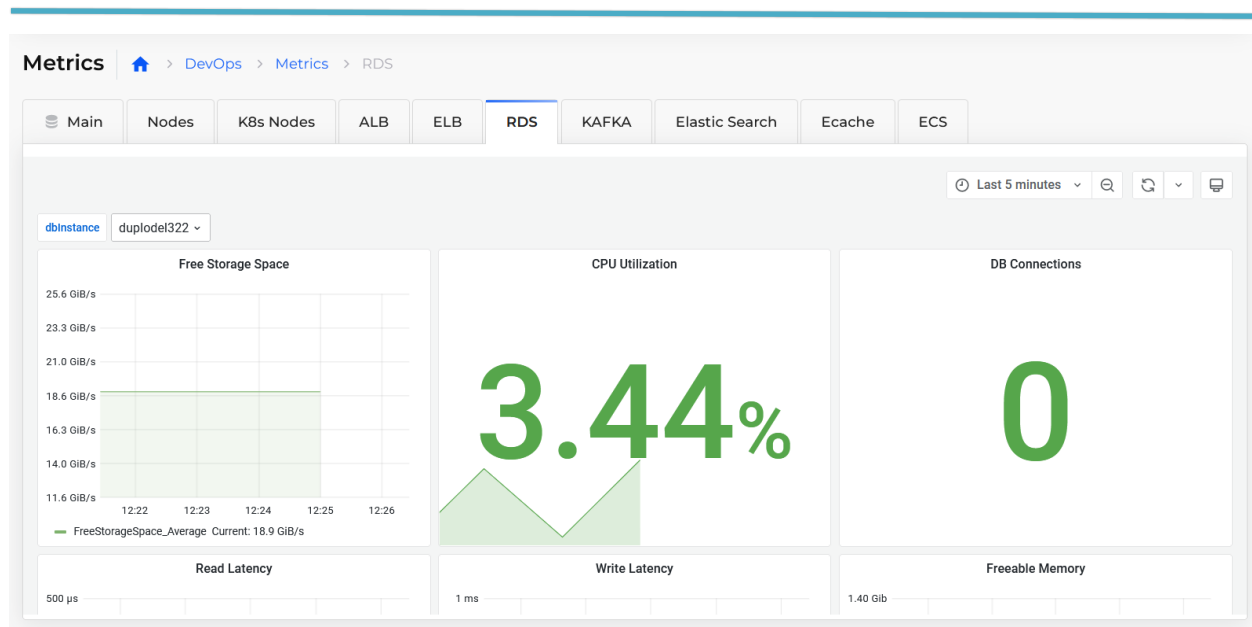
Function Package:

Passing Secrets

Passing secrets to the Docker based service/ Lambda function can be done using environmental variables. For example, you can create a SQL database from the **DevOps → Database → RDS** tab in DuploCloud, provide username and password, then in the Lambda menu give the same username and password. No secrets need to be stored anywhere outside like vault, git, repo, etc.

Monitoring

Metrics of the resources created/managed in DuploCloud can be tracked under **DevOps → Metrics**.



Faults/Alarms

Faults that happen in the system, be it Infrastructure creation, container deployments, Application health checks, or any Triggered Alarms can be tracked in the DuploCloud portal under Faults Menu. You can look at Tenant-specific faults under **DevOps → Faults** or all the faults in the system under **Administrator → Faults**. You can set the AWS Alarms for the individual metrics, click on the bell icon on any of the metrics. A form to create an alarm shows up. You can provide the necessary information and create the alarm.

Faults [Home](#) > [DevOps](#) > [Faults](#)

[Update Sentry Config](#)

	NAME	RESOURCE	DESCRIPTION
>	ContainerManagement_bebbbe50-fd92-46ce-b2c7-4b7bd312ef69	97513d1a-dbb1-446b-a77b-a6a6a1aa9bb4	Duplo Health Check for service nginxdemo container 507184eb7352cbf875809a71aa4609ec93885d818955c962b1924e3cf6e928b0 on host 10.188.20.242 failed
>	ContainerManagement_bebbbe50-fd92-46ce-b2c7-4b7bd312ef69	03fb0b23-79c0-40be-ac81-c4a6f868cec8	Create Container 03fb0b23-79c0-40be-ac81-c4a6f868cec8 Service test-service-s3 Error {"message": "No such image: duplocloud/nodejs-hello-world-s3:latest"}

2 total

Beta UI Tenant: DEFAULT [Switch to Old UI](#)

EC2 [Home](#) > [DevOps](#) > [Hosts](#) > [EC2](#) > i-0c89df4e529b46c44

DUPLOSERVICES-DEFAULT-MONITORING

Instance ID: i-0c89df4e529b46c44 Fleet Status: Connected

Containers

Metadata

Metrics

1 hour 6 hours 12 hours 1 day 7 days 30 days

Memory Utilization Average

CPU Utilization Average

Create Alarm

Statistic

Average

Metric

mem_used_percent

Operator

>=

Threshold

70

Period

5 Minutes

Cancel

Create

Alerting/Notification

In addition to notify you about the faults, DuploCloud integrates with Sentry, which will send an Email alert for the fault and acts as a single place to look at all the events. To configure Sentry, go to <https://sentry.io/> Under projects create a new project. Then go to **Settings → Projects → your project → Client keys (DSN)**. Click on show deprecated DSN, Get the deprecated DNS name and add it under **DevOps → Faults → Update Sentry Config**.

Settings > duplocloud > dev-test > Client Keys

Client Keys Generate New Key

To send data to Sentry you will need to configure an SDK with a client key (usually referred to as the `SENTRY_DSN` value). For more information on integrating Sentry with your application take a look at our [documentation](#).

DEFAULT Configure Disable Trash

DSN
The DSN tells the SDK where to send the events to. [Hide deprecated DSN](#)

`https://5ed27cd2e27ea2@o197006.ingest.sentry.io/5` Copy

Deprecated DSN includes a secret which is no longer required by newer SDK versions. If you are unsure which to use, follow installation instructions for your language.

`https://59f5ed27cd2e27ea2-69d700754c3c4ac0b78504afd1ae45d5@o197006.ingest.sentry.io/5` Copy

Security Header Endpoint
Use your security header endpoint for features like CSP and Expect-CT Expand

Beta UI Tenant: DEMOGUIDE Switch to Old UI

Faults DevOps > Faults

	NAME	RESOURCE	DESCRIPTION
>	ContainerManagement_bebbce50-fd92-46ce-b2c7-4b7bd312ef69	97513d1a-dbb1-446b-a77b-a6a6a1aa9bb4	Duplo Health Check for service nginxdemo cc507184eb7352cbf875809a71aa4609ec93885d8
>	ContainerManagement_bebbce50-fd92-46ce-b2c7-4b7bd312ef69	03fb0b23-79c0-40be-ac81-c4a6f868cec8	Create Container 03fb0b23-79c0-40be-ac81-c4a6f868cec8 image: duplocloud/nodejs-hello-workd-s3:late

2 total

Set Sentry Config ×

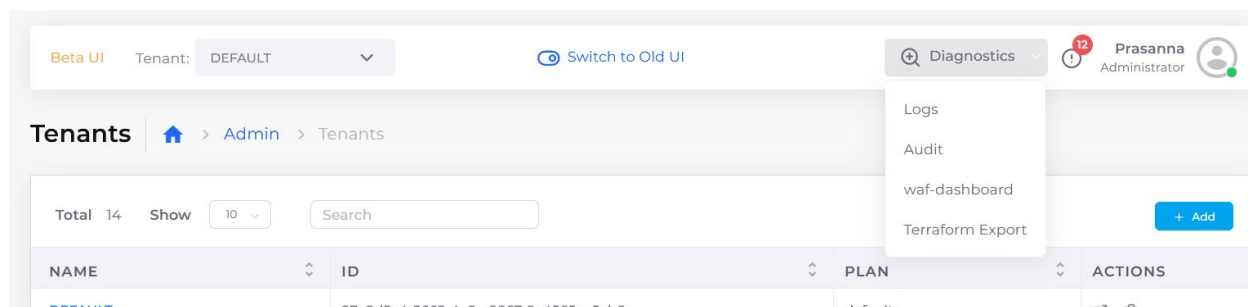
Sentry DSN ✓
`Sentry DSN`

Alerts Frequency (Seconds)
90

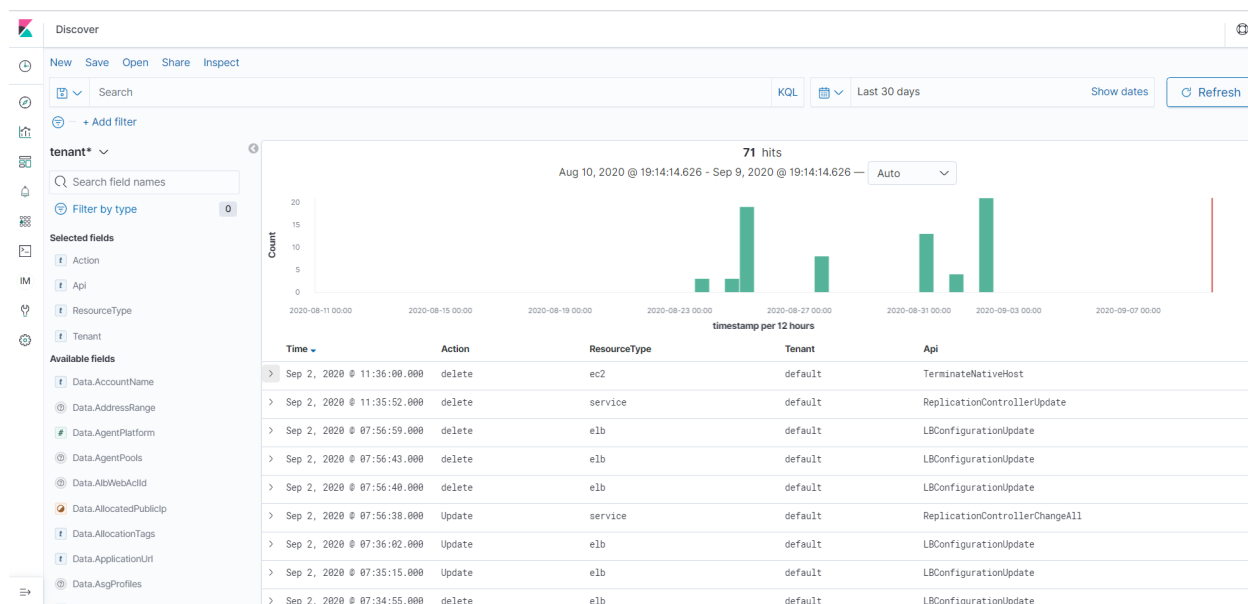
Cancel Update

Central Logging

All the activity in the DuploCloud is logged which can be used for auditing. All the logs are saved into Elasticsearch and can be visualized in Kibana. The URL for the Kibana is available under Diagnostics.



The ES & Kibana will be sitting inside the VPC and cannot be accessed from outside. Connect to the VPN and access these URL.



Backup & Restore

Backup of the resources, be it DB Snapshots, VM Snapshots/VM AMI can be done directly under the resource in the DuploCloud. Once the backups are done, they will be available to be restored on the next instance creation.

RDS > DevOps > Database > RDS

Rds | Ec2 | DynamoDb

Total 1 Show 10 Search + Add

NAME	ENDPOINT	SIZE	STATUS	ENGINE	ACTIONS
DUPLODEMODB	duplodemodb.cl13p7eyfk8n.us-west-2.rds.amazonaws.com:5432	db.t2.small	available	PostgreSQL	⋮

1 total

- View Instance Json
- Restart Instance
- Start Instance
- Stop Instance
- Create Snapshot**
- Update Delete Protection
- Update Password
- Delete Instance

RDS > DevOps > Database > RDS

Create a RDS Cancel Create

Rds Name

Rds Name

Create from Snapshot (Optional)

Choose Snapshot

duplodemodb-2021-08-19t20-41-37

rds:duplodemodb-2021-08-19-13

User password

User Password

Rds Engine

Choose Rds Engine

Rds Engine version (Optional)

version

Database Size

Choose Database Size

Storage size in GB

size

DB Parameter Group (Optional)

Choose DB Parameter Group

Encryption Key (Optional)

Choose Encryption Key

Disaster Recovery

Documentation TBD. Please contact DuploCloud team for assistance.

Cost Management

Usage cost for the resources can be viewed in the DuploCloud under **Administrator** → **Billing**. You can explore the cost for any period and for any set of resources.

All
Admin
Billing
All

Billing
Tenant

Billing Summary (14 Months)

<div>Aug 2021</div> <div>\$1531.67</div> <div>Top 5 Services</div> <ul style="list-style-type: none"> Elastic Compute Cloud - Compute\$373.83 Elastic Load Balancing\$267.70 EC2 - Other\$264.80 SageMaker\$103.14 Elastic Container Service for Kubernetes\$99.97 <div>More Details</div>	<div>Jul 2021</div> <div>\$2918.12</div> <div>Top 5 Services</div> <ul style="list-style-type: none"> Elastic Load Balancing\$482.83 Elastic Compute Cloud - Compute\$477.30 EC2 - Other\$472.06 Elastic Container Service for Kubernetes\$219.46 SageMaker\$180.10 <div>More Details</div>	<div>Jun 2021</div> <div>\$2413.33</div> <div>Top 5 Services</div> <ul style="list-style-type: none"> EC2 - Other\$479.43 Elastic Load Balancing\$438.56 Elastic Compute Cloud - Compute\$380.90 Elastic Container Service for Kubernetes\$211.23 SageMaker\$173.92 <div>More Details</div>
<div>May 2021</div> <div>\$4949.39</div> <div>Top 5 Services</div> <ul style="list-style-type: none"> EC2 - Other\$476.57 Elastic Load Balancing\$420.44 Elastic Compute Cloud - Compute\$344.38 Elastic Container Service for Kubernetes\$223.20 AWS Cost Explorer\$211.70 <div>More Details</div>	<div>Apr 2021</div> <div>\$1929.57</div> <div>Top 5 Services</div> <ul style="list-style-type: none"> Elastic Load Balancing\$361.96 Elastic Compute Cloud - Compute\$309.92 EC2 - Other\$297.99 Elastic Container Service for Kubernetes\$216.00 SageMaker\$191.20 <div>More Details</div>	<div>Mar 2021</div> <div>\$1907.48</div> <div>Top 5 Services</div> <ul style="list-style-type: none"> EC2 - Other\$408.51 Elastic Compute Cloud - Compute\$303.90 Elastic Load Balancing\$297.57 Elastic Container Service for Kubernetes\$228.79 SageMaker\$174.44 <div>More Details</div>

Bigdata/ETL

- **Use case:**
 - **Collection of data from using various methods/sources**
 - Web scraping: Selenium using headless chrome/firefox.
 - Web crawling: statis website sing crawling
 - API to Data collection: It could be REST or GraphQL API
 - Private internal customer data collected over various transactions
 - Private external customer data collected over secured SFTP
 - The data purchased from 3rd party
 - The data from various social networks
 - **Correlate data from various sources**
 - Clean up and Process data and apply various statistical methods, create
 - Correlate terabytes of data from various sources and make sense from the data.

-
- Detect anomalies, summarize, bucketize, and various aggregations
 - Attach meta-data to enrich data.
 - Create data for NLP and ML models for predictions of future events.
 - **AI/ML pipelines and life-cycle management**
 - Make data available to data science team
 - Train models and do continues improvement trials, reinforcement learning.
 - Create anomalies, bucketize data, summarize and do various aggragations.
 - Train NLP and ML models for predictions of future events based on history
 - Create history for models/hyper parameters and data at various stages.
 - **Deploying Apache Spark™ cluster**

In this tutorial we will create a Spark cluster with a Jupyter notebook. A typical use case is ETL jobs, for example reading parquet files from S3, processing and pushing reports to databases. The aim is to process GBs of data in faster and cost-effective way. The high-level steps are: (1) Create 3 VMs one for each Spark master, Spark worker and Jupyter notebook. (2) Deploy Docker images for each of these on these VMs.

- **Create a host for Spark master:** Click **DevOps → Hosts → EC2 → +Add** button to create hosts and click **show advanced**. Change the value of instance type to 'm4.xlarge' and add an allocation tag 'sparkmaster'.

Add Host

Friendly Name: spark-master ✓

Instance Type: Other x

Other Instance Type: m4.xlarge ✓

☒ Advanced Options

Allocation Tags: sparkmaster ✓

Image Id: Duplo-Docker-ubuntu x

Agent Platform: Linux Docker/Native x

Availability Zone: Zone B x

Disk Size: 0

Enable Block EBS Optimization: No

Enable Hibernation: No

Base64 Data: 1

Tags: 1

Volumes: 1

Network interfaces: 1

Cancel Add

- **Create host for spark worker:** Create another host for the worker. Change the value of instance type to 'm4.4xlarge' and add an allocation tag 'sparkworker'. Click on submit. The number of workers depends on how much load you want to process. You should add one host for each worker. They should all have the same allocation tag 'sparkworker'. You can add/remove workers and scale up or down the worker Spark service as many times as you want. We will see in the next steps:

Add Host

Friendly Name: spark-worker-01 ✓

Instance Type: Other x

Other Instance Type: m4.4xlarge ✓

☒ Advanced Options

Allocation Tags: sparkworker ✓

Image Id: Duplo-Docker-ubuntu x

Agent Platform: Linux Docker/Native x

Availability Zone: Zone B x

Disk Size: 0

Enable Block EBS Optimization: No

Enable Hibernation: No

Base64 Data: 1

Tags: 1

Volumes: 1

Network interfaces: 1

Cancel Add

- **Create host for jupyter:** Create one more host for Jupyter notebook. Choose the value of instance type to 'm4.4xlarge' and add the allocation tag as 'jupyter'.

- **Create spark master Docker service:** Go to **Containers → EKS/ Native → Services** and create new service. Under name choose 'sparkmaster', image 'duplocloud/anyervice:spark_v6', add the allocation tag 'sparkmaster'. In the Docker host config select Host network. By setting this in Docker Host config you are making the container networking the same as the VM i.e., container IP is same as VM.

- **Create Jupyter service:** First we need the IP address of Spark master. Click on Spark master service and on the right expand the container details and copy the host IP. Create another service, under name choose 'jupyter', image 'duplocloud/anyervice:spark_notebook_pyspark_scala_v4', add the allocation jupyter and select Host network for Docker Host Config, Add

volume mapping `"/home/ubuntu/jupyter:/home/jovyan/work"`, Also provide the environment variables `{"SPARK_OPTS": " --master spark://<>:7077 --driver-memory 20g --executor-memory 15g --executor-cores 4 "}`, replace `<>` with the IP you just got above.

The screenshot shows the 'Services' page in the AWS Management Console for a service named 'sparkmaster'. The service is running with 1 replica. The container 'sparkmaster' is shown with its image 'duplocloud/anyervice:spark_v6' and host IP '10.188.56.39' highlighted in a blue box. The 'HOST IP' column is highlighted in a blue box.

CONTAINER ID	NAME	IMAGE	HOST	HOST IP	DESIRED	CURRENT	ACTIONS
4cfb27e36320d29d072b	sparkmaster	duplocloud/anyervice:spark_v6	i-034598d2d78aedd88	10.188.56.39	Running	Stopped	

The screenshot shows the 'Add Service' form in the AWS Management Console. The form is in the 'Basic Options' tab. The service name is 'jupyter', cloud is 'AWS', platform is 'Linux Docker/Native', Docker image is 'duplocloud/anyervice:spark_notebook_pyspark_scala_v4', allocation tag is 'jupyter', and replicas is '1'. The environment variables section shows 'SPARK_OPTS' set to ' --master spark://10.188.56.39:7077'.

- **Create Spark workers:** Create another service name 'sparkworker', image 'duplocloud/anyervice:spark_v7', add the allocation sparkworker and select Host network for Docker Host Config. Also provide the environment variables `{"node": "worker", "masterip": "<>"}`, replace `<>` with the IP you just got above. Depending on how many worker hosts you have created, use the same number under replicas and that is the way you can scale up and down. At any time, you can add new hosts, set the allocation tag sparkworker and then under services, edit the sparkworker service and update the replicas.

Add Service

1 Basic Options Minimal inputs to start service > **2 Advanced Options** More options to configure service

Service Name ✓ Cloud x Platform x

Docker Image ✓ Allocation Tag ✓ Replicas

Docker Networks x Daemonset v

Environment Variables

```

1 {
2   "node": "worker",
3   "masterip": "10.188.56.39"
4 }

```

← Previous Next →

- **Create Docker services shell** : Add/update shell access by clicking on >_ icon. This gives you easy access into the container shell. You will need to wait for 5 minutes for the shell to be ready. Make sure you are connected to VPN if you choose to launch the shell as internal only
- **Open Jupyter docker shell**: Select Jupyter service and expand the container. Copy the hostip and then click on >_ icon.

jupyter Actions v

Image: duplocloud/anyervice:spark_notebook_pyspark_scala_v4' Replicas:1

Containers Load Balancers Configuration LB Metrics

Total 1 Show 10 Search

CONTAINER ID	NAME	IMAGE	HOST	HOST IP	DESIRED	CURRENT	ACTIONS
	jupyter	duplocloud/anyervice:spark_notebook_pyspark_scala_v4'	i-0962345199957df3e	10.188.57.203	Running	Running	<ul style="list-style-type: none"> Logs State >_ Shell

1 total

Running 0/1

LB Status NA

DNS

LB Visibility External

Operating System Linux

Container Platform Linux Docker/Native

- **Get Jupyter URL**: Once you are inside the shell. Enter 'jupyter notebook list' to get the URL along with auth token. Replace the Ip with Jupyter ip you copied previously.

```
Container Shell
jovyan@ip-10-165-21-204:~$ jupyter notebook list
Currently running servers:
http://0.0.0.0:8888/?token=11f85f9e3e95d4c59e37a4c1003575247d6cf9c49050447f :: /home/jovyan/work
jovyan@ip-10-165-21-204:~$
```

- **Open Jupyter:** In your browser, invoke the Jupyter URL and you should be able to see the UI.



Now you can use Jupyter to connect to data sources and destinations and do ETL jobs. Sources and destinations can include various SQL & NoSQL databases, S3 and various reporting tools including big data and GPU-based Deep learning.

- **Jupyter with an ETL Use-case**

In this tutorial we will create a Jupyter notebook and show some basic web scraping, using Spark for preprocessing, exporting into schema, do ETLs, join multiple dataframes (parquets), and export reports into MySQL.

- **Scrape data from Internet:** Connect to a website and parse html (using jsoup)


```

In [ ]: import org.jsoup.nodes._
import org.jsoup.Jsoup
val doc = Jsoup.connect("http://download.cms.gov/provider_data.html").get()

In [ ]: Files.write(Paths.get(provider_data_html),
doc.outerHtml().getBytes(StandardCharsets.UTF_8))

In [31]: val all_links = doc.select("li").asScala
val zip_file_links = all_links.filter(_.text().indexOf("ZIP format") > 0)
val zipFiles = zip_file_links.map(el => ZipFile(el.select("a").first().attr("href").replace("./", "") , el.text()))
val df_zipFiles = spark.createDataFrame( zipFiles)

zipFiles.foreach( (file:ZipFile) => println(file.file_name + " \n : " + file.description))

```

- **Download data:** Extract the downloaded zip. This particular file is 8 GB in size and has 9 million records in csv

```

In [ ]: import java.net.URL
import java.io.File
import org.apache.commons.io.FileUtils
val tmpFile = new File(folder_full+"/full.zip") |
FileUtils.copyURLToFile(new URL("http://download.cms.gov/providers/"+monthly_zip_file), tmpFile)

In [ ]: val result_ls = "ls -al " + tmpFile !
val result_rm = "rm -rf " + folder_full !
val result_unzip = "unzip " + tmpFile + " -d " + folder_full !

```

- **Save to S3** Upload the data to Aws S3

```

In [19]: import java.net.URI
import com.amazonaws.services.s3.transfer.{TransferManager, TransferManagerBuilder}
import com.amazonaws.services.s3.{AmazonS3, AmazonS3ClientBuilder}
import com.amazonaws.services.s3.model.{InitiateMultipartUploadRequest, InitiateMultipartUploadResult}
val s3Client = new AmazonS3Client(new DefaultAWSCredentialsProviderChain())
val transferManager: TransferManager = TransferManagerBuilder.standard()
.withS3Client(s3Client)
.build()
def pushFile(storageUrl:URI, localPath:File):Unit = {
if( storageUrl.getScheme != "s3" ){
println("S3InvalidProtocol " + storageUrl.getScheme)
return
}
if( !localPath.exists() ){
println("S3InvalidLocalPath " + localPath)
return
}
val s3Bucket = storageUrl.getHost
val s3Key = storageUrl.getPath.replaceFirst("^/", "")

if (localPath.isDirectory) {
transferManager.uploadDirectory(s3Bucket, s3Key, localPath, true).waitForCompletion()
} else {
transferManager.upload(s3Bucket, s3Key, localPath).waitForCompletion()
}
}

s3Client = com.amazonaws.services.s3.AmazonS3Client@24b8b208

```

- **Connect to spark cluster** Also Configure session with required settings to read & write from Aws S3

```
In [10]: val spark = SparkSession.builder()
        .appName("SparkNLPJob")
        .master("spark://10.165.25.255:7077") // get from env
        .config("spark.hadoop.fs.s3a.impl", "org.apache.hadoop.fs.s3a.S3AFileSystem")
        .config("spark.hadoop.fs.s3.impl", "org.apache.hadoop.fs.s3native.NativeS3FileSystem")
        .getOrCreate()
        spark.sparkContext.hadoopConfiguration.set("fs.s3a.impl", "org.apache.hadoop.fs.s3a.S3AFileSystem")
        spark.sparkContext.hadoopConfiguration.set("fs.s3.impl", "org.apache.hadoop.fs.s3native.NativeS3FileSystem")

        val spark_sc = spark.sparkContext

        spark = org.apache.spark.sql.SparkSession@77fa5058
        spark_sc = org.apache.spark.SparkContext@643a61d7

Out[10]: org.apache.spark.SparkContext@643a61d7
```

○ Load data in spark cluster

```
In [ ]: val txmy_df = spark.read.parquet(txmy_path)
        val txmy2_df = spark.read.option("header", "true")
        .csv(txmy_2_path)
```

○ Define schema

```
In [ ]: import org.apache.spark.sql.types.StructType
        val schemaUntyped = new StructType()
        .add("icd_name", "string")
        .add("icd_code", "string")
        .add("cpt_name", "string")
        .add("cpt_code", "string")
        .add("incidencerate", "double")
        .add("cptfrequency", "double")
        .add("visit25", "integer")
        .add("visit50", "integer")
        .add("visit75", "integer")
        .add("visitmean", "double")
        .add("costmean", "double")
        .add("odgcaa", "integer")
        .add("paymentflag", "string")
```

○ Do data processing

```
In [39]: import org.apache.spark.sql.functions.{when,col}
        val df2 = df_colssel.na.fill("NULL").withColumn("newbie", when($"Entity Type Code" === 1, "Individual").otherwise("Organ")
        .drop($"Entity Type Code").withColumnRenamed("newbie", "Entity Type Code")
        .withColumn("newbie", when($"Provider Gender Code" === "M", "Male").when($"Provider Gender Code" ===
        .drop($"Provider Gender Code").withColumnRenamed("newbie", "Provider Gender Code")

        df2 = [NPI: string, Replacement NPI: string ... 37 more fields]

Out[39]: [NPI: string, Replacement NPI: string ... 37 more fields]

In [40]: import org.apache.spark.sql.functions.{when,col}
        val df2 = df_colssel.na.fill("NULL").withColumn("newbie", when($"Entity Type Code" === 1, "Individual").otherwise("Organ")
        .drop($"Entity Type Code").withColumnRenamed("newbie", "Entity Type Code")
        .withColumn("newbie", when($"Provider Gender Code" === "M", "Male").when($"Provider Gender Code" ===
        .drop($"Provider Gender Code").withColumnRenamed("newbie", "Provider Gender Code")

        df2 = [NPI: string, Replacement NPI: string ... 37 more fields]

Out[40]: [NPI: string, Replacement NPI: string ... 37 more fields]

In [41]: val df3 = df2.withColumn("NPI Activation Status", when(($"NPI Reactivation Date" === "Null")&&($"NPI Deactivation Date
        df3 = [NPI: string, Replacement NPI: string ... 38 more fields]
```

○ Spark SQL

```
In [ ]: val spark2 = spark
import spark2.implicitly._
import org.apache.spark.sql.functions._
```

```
In [ ]: df.createOrReplaceTempView("provData_Full")
```

```
In [ ]: val df_colssel = spark.sql("""select `PID`,
`Entity Type Code`,
`Replacement NPI`,
```

- **Spark SQL joins** 20 GB of data from multiple sources

```
In [ ]: val df_join = df3.join(taxy_trans, col("Healthc_1") === col("Tax_Code"), "left_outer")
```

```
In [ ]: df_join.show
```

- **Export reports to RDS for UI consumption** Generate various charts and graphs

```
In [54]: val mysql_server="duplonpidata.aaaaa.us-west-2.rds.amazonaws.com"
val username="dupploUser"
val password="dupploUser123"
val myssql_db_name="duplonpidata"

val connectionProperties = new java.util.Properties
connectionProperties.setProperty("driver", "com.mysql.jdbc.Driver")
connectionProperties.setProperty("user", username)
connectionProperties.setProperty("password", password)
connectionProperties.setProperty("useSSL", "false")
val jdbcUrl = "jdbc:mysql://" + mysql_server + ":3306/" + myssql_db_name // "jdbc:mysql://localhost:3306/data"
```

Security & Compliance

Refer to [Compliance](#)

Cloud Migration (SMS)

Documentation TBD. Please contact DuploCloud team for assistance.

Remote Working

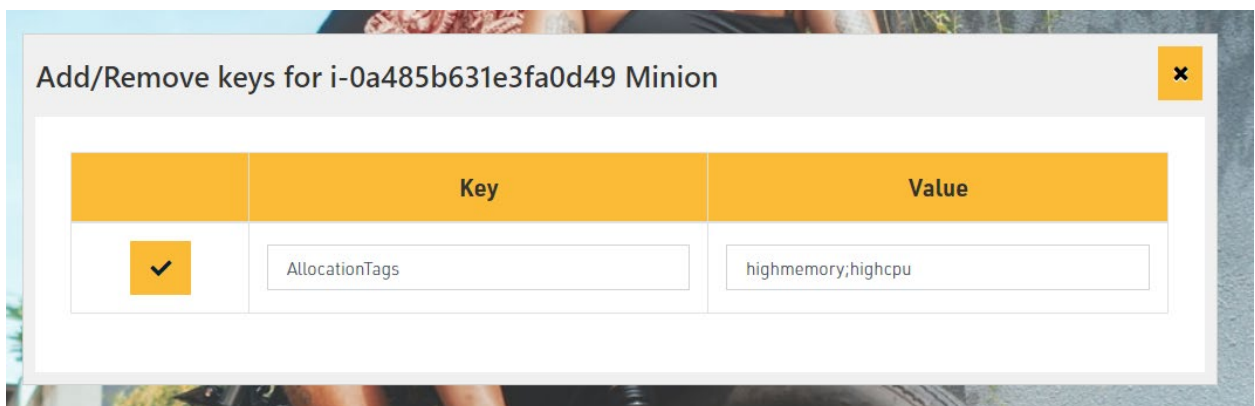

Refer to [Workspaces](#)

Additional Deployment functionalities

Advanced Functions

- **Allocation tags:** By default, DuploCloud spreads the container replicas across the hosts. By virtue of allocation tag the user can pin a container to a set of hosts that have a specific tag. Click on the edit icon next to allocation tag under host.

"AllocationTags" as key is already pre-populated, configure the value of your choice for example "highmemory;highcpu" or "serviceA" and click on save on the left. Then, when creating a service, set the allocation tag value to be a substring of the above tag, for example highmemory or serviceA. The allocation algorithm tries to pick a host where AllocationTag specified in the service is a substring of the AllocationTags of the host.



BYOH

If you have a host already running somewhere in the cloud or on-premise, you can bring that to DuploCloud using BYOH functionality and let DuploCloud manage the host, let it be running the containers or installing the compliance agents. To configure BYOH go to **Deployments → Hosts → BYO-HOSTS**. Click on **add** and provide the name, IP Address and in the Fleet type select

Native App (if you don't, DuploCloud will not manage the containers but manage the compliance agents on the host), provide the username, password/Private key file. Make sure the SSH access to the host is opened to access from the DuploCloud. After the host is added in about 5mins you can go to **Security → Agents → Select Agent** and see that the agent on host is in Active state.

Add a New BYO Host

×

Friendly Name

test-byoh

Direct Address

57.45.23.127

Fleet Type

Native App

Username

ec2-user

Password

Password

PrivateKey

CPcZw+nQ83QYUUh8glc30L5pLqFeQV9xzMnjdFvrHBQzYlft+mZHCYrvq5qe8/a0IdeG+36a
-----END RSA PRIVATE KEY-----

SUBMIT

CANCEL

Tenants

AGENT CONFIG

Search

DEFAULT

STG1

COMPLIANCE

OSSEC

Agent details

TRAFFIC MIRRORING

Agent details

CLOUDWATCHAGENT_0

Agent details

CLAMAV_V0

Agent details

CLAMAV_SCANNER_V1

Agent details

Servers: 2

Search

duploservices-default-test-i-0b75d04c217fd5061

Monitored: Yes

Agent Status: Active

IP: 10.137.27.77

OS Platform: Linux

Last Active: 2020-06-02 07:51:05

duploservices-default-test-bboh

Monitored: Yes

Agent Status: Active

IP: 10.137.6.45

OS Platform: Linux

Last Active: 2020-06-02 07:51:09

Administrator's Guide

Configuration Scopes

- **Base Configuration:** This must be set up by the user outside DuploCloud and provided as an input to DuploCloud. This is a once in a lifetime configuration. This configuration includes one VPC and one subnet with appropriate routes. It is recommended that you setup at least 2 Availability Zones, one public and one private subnet in each availability zone. All resources (EC2 instances, RDS, Elastic-cache, etc.) are placed in the private

subnet by default. ELB is placed in public or private subnet based on user (tenant) choice. DuploCloud provides a default cloud formation template that can be used if desired.

- **Plan Configuration:** Every tenant is part of one and only one plan. Configuration applied at the plan level is applied to all tenants in a plan. A plan can be used to denote say a dev environment, a class of tenants (private\public facing) etc. Following are the plan configurations. Except of VPC and subnets rest of the parameters can be changed at will. Plan parameters include:

```
{
    "Name": "devplan", /* Name of the plan */
    "Images": [ /* AMIs that are made available for the
tenants */
        {
            "Name": "Duplo-Host-Docker-v1.17", /* User
friendly Name of the AMI displayed to the user */
            "ImageId": "ami-0861ea68", /* AWS AMI ID */
            "OS": "Linux"
        },
        {
            "Name": "ubuntu-dev",
            "ImageId": "ami-d83af7b8",
            "OS": "Linux"
        },
        {
            "Name": "Windows-Docker-Fleet",
            "ImageId": "ami-1977d979",
            "OS": "Windows"
        },
        {
            "Name": "Unmanaged Ubuntu-16.04",
            "ImageId": "ami-5e63d13e",
            "OS": "Linux"
        }
    ],
    "Quotas": [ /* Limit of on the number of resources
that be used by each tenant. If none is set for a given resource type,
then there is no limit. */
        {
            "ResourceType": "ec2",
            "CumulativeCount": 2,
            "InstanceQuotas": [
                {
                    "InstanceType": "t2.medium",
                    "MetaData": "(2CPU, 4GB)",
                    "Count": 2
                },
                {
                    "InstanceType": "t2.small",
                    "MetaData": "(1CPU, 2GB)",
                    "Count": 2
                }
            ]
        }
    ]
}
```

```
]
},
{
  "ResourceType": "rds",
  "CumulativeCount": 1,
  "InstanceQuotas": [
    {
      "InstanceType": "db.t2.small",
      "MetaData": "(1CPU, 1.7GB)",
      "Count": 1
    }
  ]
},
{
  "ResourceType": "ecache",
  "CumulativeCount": 1,
  "InstanceQuotas": [
    {
      "InstanceType": "cache.t2.small",
      "MetaData": "(1CPU, 1.7GB)",
      "Count": 1
    }
  ]
},
{
  "ResourceType": "s3",
  "CumulativeCount": 1,
  "InstanceQuotas": [
    {
      "InstanceType": "bucket",
      "MetaData": "bucket",
      "Count": 1
    }
  ]
},
{
  "ResourceType": "sqs",
  "CumulativeCount": 1,
  "InstanceQuotas": [
    {
      "InstanceType": "queue",
      "MetaData": "queue",
      "Count": 1
    }
  ]
},
{
  "ResourceType": "dynamodb",
  "CumulativeCount": 1,
  "InstanceQuotas": [
    {
      "InstanceType": "table",
      "MetaData": "table",
      "Count": 1
    }
  ]
},
},
```

```

        {
            "ResourceType": "elb",
            "CumulativeCount": 2,
            "InstanceQuotas": [
                {
                    "InstanceType": "elb",
                    "MetaData": "elb",
                    "Count": 1
                }
            ]
        },
        {
            "ResourceType": "sns",
            "CumulativeCount": 1,
            "InstanceQuotas": [
                {
                    "InstanceType": "topic",
                    "MetaData": "topic",
                    "Count": 1
                }
            ]
        }
    ],
    "AwsConfig": {
        "VpcId": "vpc-1eafce79", /* VPC for this tenant */
        "AwsHostSg": "sg-c2d899ba", /* list of security
groups separated by ;. All hosts will be placed in this security
groups. */
        "AwsElbSg": "sg-b9d899c1", /* Security group in
which the ELB will be placed. This applies to traffic into the external
(VIP) of the elb. Internal Sg between hosts and ELB will be auto setup.
*/
        "AwsPublicSubnet": "subnet-0066b449;subnet-
24e25943", /* Subnets in which hpublic facing ELBs must be placed. Each
subnet corresponds to an AZ. */
        "AwsInternalElbSubnet": "subnet-0e66b447;subnet-
23e25944", /* Subnets in which internal resources like Ec2 hosts, rds
ecaache etc will be placed. If you like hosts to be public facing then
put the same public subnets here. */
        "AwsElastiCacheSubnetGrp": "duplo-cache-
aww6gk5hsy4n",
        "AwsRdsSubnetGrp": "duplo-vpc-21-resources-
dbsubnetduplo-z9fh3g59362z",
        "CommonPolicyARNs": "", /*Set of IAM policy ARNs
that will be applied to all tenants. */
        "Domain name": "" /* Name of route 53 domain that
will be used for tenant custom dns names*/
        "CertMgrResourceArns": "" /* List of certificate
arns that will be available to the tenant for SSL termination on the
ELB. */
    },
    "UnrestrictedExtLB": false,
    "Capabilities": {
        "DisableNativeApps": false,
        "DisablePublicEps": false,
        "DisablePrivateElb": false,
        "AssignInstanceElasticIp": false,

```

```
"BlockEbsOptimization": false,
"DisableSumoIntegration": false,
"DisableSignalFxIntegration": false,
"EnableTenantExpiry": false
}
}
```

- **Tenant Configuration:** These are the configuration that we covered in the deployment guide.

Access Control

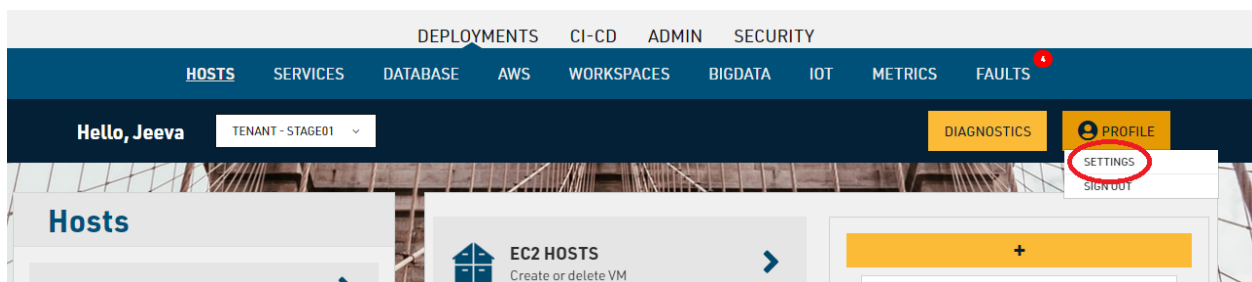
There are 2 types of roles, user and administrator. Each user can have access to one or more tenants. Each tenant can be accessed by one or more users. Administrator has access to all tenants plus the administrative functions like plan configuration, system dashboard, system faults, etc.

Accessing Cloud Server

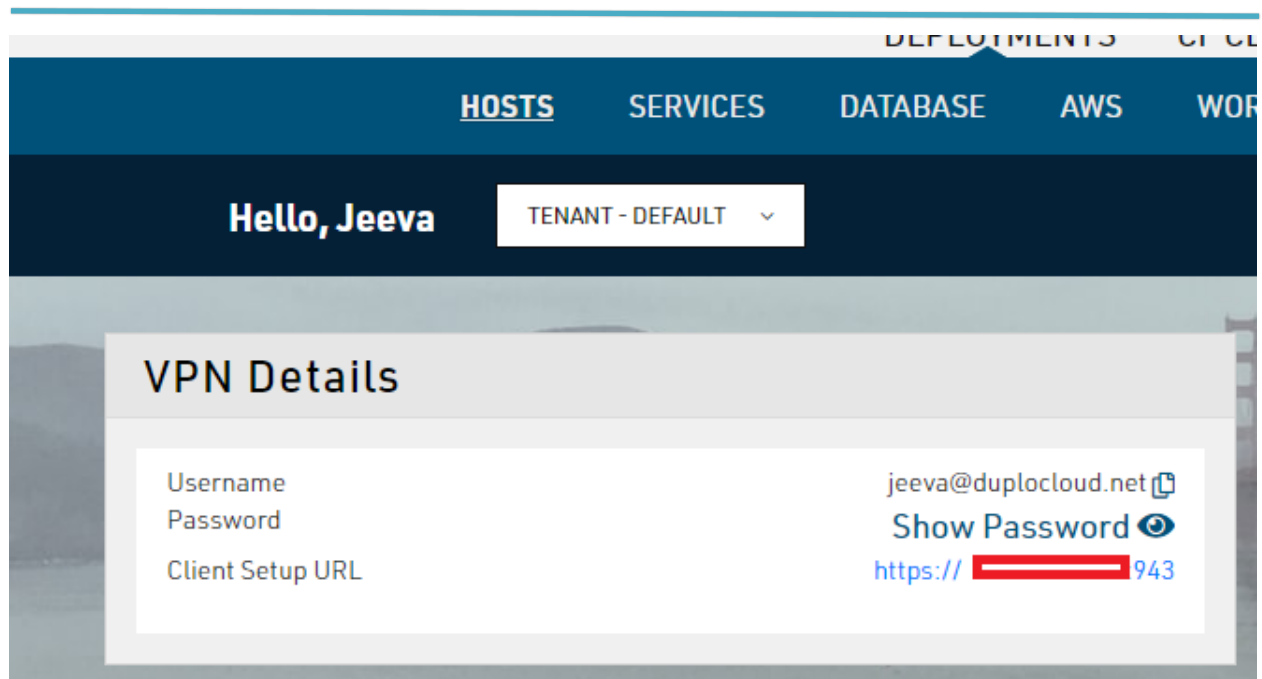
VPN Access

DuploCloud provisions the OpenVPN server for users to connect with the cloud resources like EC2, RDS instances, Elastic Search, Elastic Cache etc. Below steps will guide you to connect with VPN server using OpenVPN client.

- **Setup OpenVPN client**
Login to OpenVPN web portal to download OpenVPN config and client.
- Click on **Profile → Settings** as shown in image below



- Click the openvpn link which will open a new tab.



- Login using Username and Password provided.

A screenshot of the OpenVPN Access Server User Login form. The form has a header with the OpenVPN Access Server logo. Below the logo, the text "User Login" is displayed. There are two input fields: "Username" with an email icon and "Password" with a key icon. A "Sign In" button is located at the bottom of the form.

- Download the openvpn config and client.



- Now you are ready to connect with VPN using OpenVPN client.

Sharing Encrypted DB

Sharing Unencrypted DB to other accounts is very simple and straight forward. But sharing encrypted DB is slightly difficult. Here we will go through the steps that needs to be followed to share the encrypted DB:

- Create a new customer managed key in AWS KMS, in the Define key usage permissions provide the account id of the other account.

Step 1

Configure key

Step 2

Add labels

Step 3

Define key administrative permissions

Step 4

Define key usage permissions

Step 5

Review and edit key policy

Define key usage permissions

This account
Select the IAM users and roles that can use the CMK in cryptographic operations. [Learn more](#)

< 1 2 3 4 5 >

<input type="checkbox"/>	Name	Path	Type
<input checked="" type="checkbox"/>	jeeva@duplocloud.net	/	User
<input type="checkbox"/>	pravin@duplocloud.net	/	User
<input checked="" type="checkbox"/>	srikar@duplocloud.net	/	User
<input type="checkbox"/>	AWSConfigServiceRolePolicy-custom	/	Role
<input type="checkbox"/>	AWSServiceRoleForAccessAnalyzer	/aws-service-role/access-analyzer.amazonaws.com/	Role
<input type="checkbox"/>	AWSServiceRoleForAmazonElasticFileSystem	/aws-service-role/elasticfilesystem.amazonaws.com/	Role
<input type="checkbox"/>	AWSServiceRoleForAmazonElasticsearchService	/aws-service-role/es.amazonaws.com/	Role
<input type="checkbox"/>	AWSServiceRoleForAmazonGuardDuty	/aws-service-role/guardduty.amazonaws.com/	Role
<input type="checkbox"/>	AWSServiceRoleForAmazonInspector	/aws-service-role/inspector.amazonaws.com/	Role
<input type="checkbox"/>	AWSServiceRoleForAutoScaling	/aws-service-role/autoscaling.amazonaws.com/	Role

Other AWS accounts
Specify the AWS accounts that can use this key. Administrators of the accounts you specify are responsible for managing the permissions that allow their IAM users and roles to use this key. [Learn more](#)

arn:aws:iam::

227120241369

:root

Remove

Add another AWS account

Cancel

Previous

Next

- Once the key is created, go to **RDS → Snapshots**, select the snapshot and click Copy Snapshot. In the encryption change the master key to the key we created before.

Encryption

Encryption

Info

☒ Enable encryption [Learn more](#) [🔗](#)
Select to encrypt the given instance. Master key ids and aliases appear in the list after they have been created using the Key Management Service(KMS) console.

☐ Disable encryption

Master key

Info

rds-key

▼

Account

128329325849

KMS key ID

d9483dc3-9ae4-4597-bffe-b353042e201d

Cancel

Copy Snapshot

- Once the copied snapshot is ready, as usual share the snapshot to another account by clicking share snapshot and providing the other account id.
- 4. Now go to the other **AWS account** → **RDS** → **Shared with me**. Select the shared snapshot and click copy-snapshot again and change the encryption key to the encryption key in the account.

Encryption

Encryption

Info

☒ Enable encryption

[Learn more](#)

Select to encrypt the given instance. Master key ids and aliases appear in the list after they have been created using the Key Management Service(KMS) console.

☐ Disable encryption

Master key

Info

rds-custom-key

▼

Account

227120241369

KMS key ID

eb60cbf2-a05b-48e6-a0ee-c5b80625a65f

Cancel

Copy Snapshot

- In the copied snapshot add a tag with Key as "Name" and Value as "duploservices-{tenantname}" where tenantname is the tenant where u want to launch an RDS with this snapshot.

Add tags

×

Add tags to your RDS resources to organize and track your Amazon RDS costs. [Learn more](#)

Tag key	Value
Name	duploservices-awsdemo01

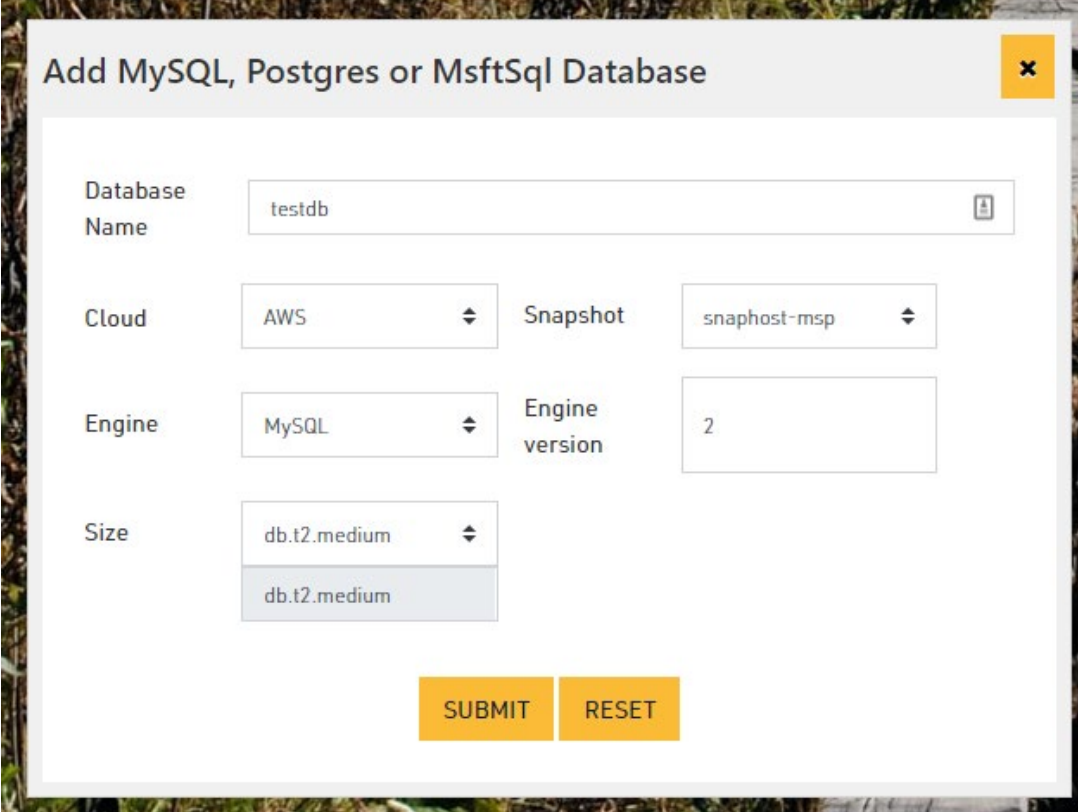
Add another Tag

Cancel

Add

- Go to DuploCloud portal select the tenant. **Open RDS → Add new DB (+ icon) → Give** name for the new DB. In the snapshot select the new snapshot. Enter instance

type and hit submit. In few mins, the DB will be created with the data from the snapshot. You must use the existing username and password to access the DB



The screenshot shows a web form titled "Add MySQL, Postgres or MsftSql Database". The form has the following fields and values:

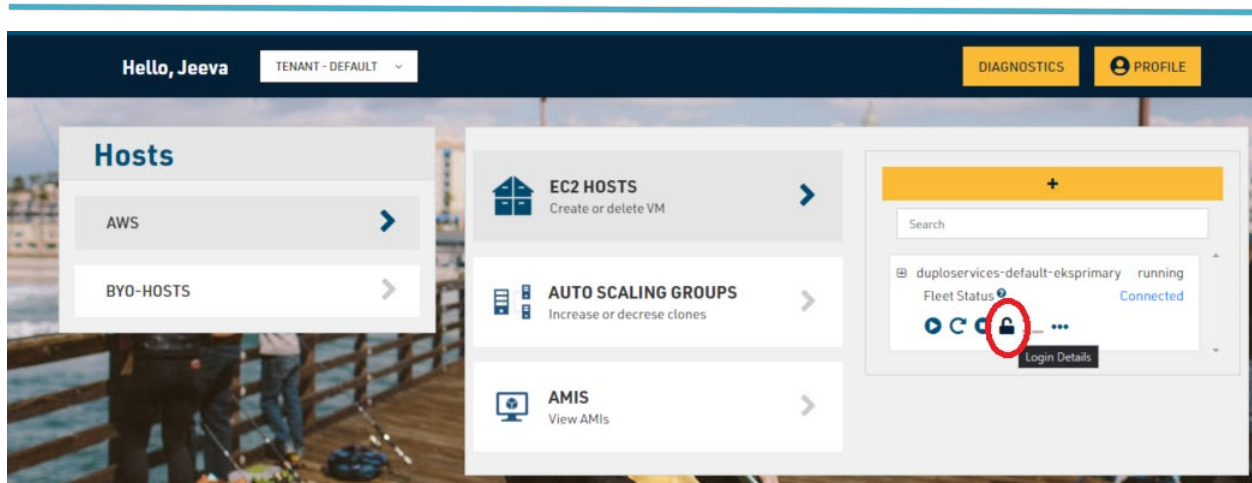
- Database Name: testdb
- Cloud: AWS
- Snapshot: snaphost-msp
- Engine: MySQL
- Engine version: 2
- Size: db.t2.medium

At the bottom of the form are two buttons: "SUBMIT" and "RESET".


SSH EC2 Instance

Once you are connected to VPN. You can use the privatekey to SSH into EC2 instance.

- **Download Private key**
Select Tenant and navigate to Deployments --> Hosts.
- Click on the unlock icon as shown below.



- Click on the key icon to download the private key. Change the permission to key file as shown and ssh into the ec2 instance.



- Now you should be connected to the server.

CI/CD Guide – Katkit: DuploCloud’s CI/CD Component

DuploCloud provides a CI/CD framework that allows you to build, test and deploy your application from Git HUB commits and PRs. We call it Katkit. Katkit is a arbitrary code execution engine which allows the user to run arbitrary code before and after deployment. Katkit follows the same notion of a "Tenant" or environment. Thus, tying together CI and CD. In other words, the tests are run against the application in same underlying AWS topology where one's code is running as against running them in a separate fleet of servers which does not capture the interactions of the application with the AWS infrastructure like IAM, Security groups ELB etc.

At a high level, Katkit functions as follows:

- A repository is linked to a Tenant.
- User chooses a GIT commit to run test and deploy

-
- Katkit deploys a service in the same tenant with the docker image provided by DuploCloud, which is essentially like a jenkins worker and had the Katkit agent in it.
 - Katkit agent in the CI container checks out the code at that commit inside the container. It then executes ci.sh from the checked-out code. Essentially each build is a short-lived service that is removed once the ci.sh execution is over.
 - User can put any arbitrary code in ci.sh
 - Katkit allows, for a given run of a commit, the user to execute code in "phases" where in each phase Katkit repeats the above steps with a difference in the ENV variables that are set in each phase. The code inside ci.sh is to read the env variables and perform different actions corresponding to each phase
 - Katkit has a special phase called "deployment" where it does not run ci.sh but it looks for the servicedescription.js file (details below), replaces the docker image tag and replaces it with the git commit sha. It is assumed that the user, before invoking the deployment phase, has gone through a prior phase where he build a docker image which was tagged with the git commit sha. The sha is available as an ENV variable in every phase.

First Deployment

Before using CI/CD, the first deployment of the application needs to be done via DuploCloud menus described above. Make sure that the application works as expected. Katkit is used only for upgrades of container images and run tests that can be written to run before and after.

Environments

In DuploCloud a standard practice is to have a separate tenant for a given logical application for each deployment environment. For example, say an application called taskrunner would be created as three tenants called d-taskrunner, b-taskrunner and p-taskrunner to represent dev, beta and prod environment. In each tenant one can specify an arbitrary name for the env say "DEV" in the menu Dashboard-->ENV. This string will be set by Katkit as an ENV variable when it runs the tests during CI/CD and thus your test code can use this value to determine what tests should be run in each env or for that matter take any other action.

Export Service Description

Service Description represents the topology of a service. It is a JSON file that is used by Katkit to upgrade the running service. Go to **Deployment → Services → Export**. This will give a json file. Save this as servicedescription.js under the servicedescription folder that must exist at the root

of your repository. In this file search for "DockerImage": and here change the image tag to the word <hubtag> for example change "DockerImage": "nginx:latest" to "DockerImage": "nginx:<hubtag>". Remove the ExtraConfig and Replicas field from the file. These have env variables and replicas which would vary from one environment to other. Hence during deployment Katkit will retain what is already present in the current running service.

Link Repository

Once the above steps have been performed, you can link your GitHub Repository to your tenant. In addition to the repository name, you also need to specify the "Home Branch" which is the branch for which the PRs will be monitored by Katkit for the user to run deployments. Same repository and branch combination can be linked in several tenants. If your repository has several services for different tenants, then each service can be represented by a separate folder at the root. This is Folder Path field. Katkit looks for service description file under /servicedescription/servicedescription.js Same repository but different folders can also be used in different tenant. Same tenant can also have different repositories.

CREATE REPOSITORY LINK

Repository Name: duploclouddemo^demoservice

Folder Path: / Home Branch: master

Repo Params: Key Value

SUBMIT RESET

Phases

Each CI/CD run comprises of one or more phases. There are two types of phases - execution and deployment. In execution phase Katkit will invoke ci.sh file from the repository. The difference between two execution phases is in ENV variables based on which user code in ci.sh can operate differently. There can be only one deployment phase in each run. Katkit does not run ci.sh in deployment phase but it looks for the servicedescription.js file (details below), replaces the docker image tag <hubtag> and replaces it with the git commit sha. It is assumed that the user, before invoking the deployment phase, has gone through a prior phase where they build a docker

image which was tagged with the git commit sha. The sha is available as an ENV variable in every phase.

DEPLOY & TEST - DUPLODEMO/DEMOSERVICE

HIDE ADVANCED

✕

Sha: a201c7991f8b1e1e4c0266725b72c6eb5e844a28

Path: / Branch: thvenket-patch-10

Fleet Local Fleet

	Name	Image	Params	Order
<input checked="" type="checkbox"/>	PRE_DEPLOY_BUILD	duplocloud/zbuilder:v7	PHASE=PRE_DEPLOY_BUILD, FOO=BAR	0
<input checked="" type="checkbox"/>	DEPLOY	duplocloud/zbuilder:v7	PHASE=DEPLOY	1
<input checked="" type="checkbox"/>	POST_DEPLOY_VERIFICATION	duplocloud/zbuilder:v7	PHASE=POST_DEPLOY_VERIFICATION, FOO=BAR	2

SUBMITRESET

Katkit Config

The above configuration customizations like Phases, ENV, etc. can be saved in the repository in a config file called katkitconfig.js Following is an example of one such file

```
[
  {
    "EnvName": "default",
    "LocalFleet": "true",
    "Workflow" : [
      {
        "Name": "PRE_DEPLOY_BUILD",
        "PhaseType": 4,
        "BuildParams": "PHASE=PRE_DEPLOY_BUILD, FOO=BAR",
        "Order": 0,
        "Parallelism": 1,
        "ContainerImage": "duplocloud/zbuilder:v7"
      },
      {
        "Name": "DEPLOY",
        "PhaseType": 1,
        "BuildParams": "PHASE=DEPLOY",
        "Order": 1,
        "Parallelism": 1,
        "ContainerImage": null
      }
    ]
  }
]
```

Advanced Functions

- **Bring-your-own-image:** By default, all tenant CI/CD runs are executed in a docker image specified by the administrator. This image would typically have the most common packages for your organization. But a user can bring his own builder image and specify the same. The image should have the Katkit agent that can be copied from the default builder image.
- **Bring-your-own-fleet or Local Fleet:** By default, Katkit will run the builder containers in a separate set of hosts, but the user can also choose to run the build container in the same tenant hosts which is being tested.