

Using Published DSpace Images on AWS Fargate

This approach is not recommended

It is fairly easy to start DSpace containers in AWS Fargate. AWS Fargate does not permit command line execution within a container, so regular DSpace command line tasks cannot be run.

See <https://stackoverflow.com/questions/52310447/is-it-possible-to-ssh-into-fargate-manged-container-instances/52310862#52310862>

I verified this limitation in a conversation with an AWS Developer Advocate.

- [Create an ECS Cluster \(specify Fargate when creating the cluster\)](#)
- [Create a Task Definition named "dspace" using "awsvpc" mode](#)
- [dspace container](#)
- [dspacedb container](#)
- [Select "Run Task" to start your task within the cluster you created.](#)
- [Customize the Security Group then Click "Run Task".](#)
- [Select the running task in order to get its IP address](#)
- [Get the public IP address](#)
- [Open ip-address:8080/xmlui](#)
- [Outstanding Issues](#)

This page describes how to run DSpace docker images on AWS Elastic Container Service (ECS) using [AWS Fargate](#).

Create an ECS Cluster (specify Fargate when creating the cluster)

Amazon ECS

Clusters

Task Definitions

Amazon ECR

Repositories

Clusters > tbcluster1

Cluster : tbcluster1

Delete Cluster

Get a detailed view of the resources on your cluster.

Status **ACTIVE**

Registered container instances 0

Pending tasks count 0 Fargate, 0 EC2

Running tasks count 0 Fargate, 0 EC2

Active service count 0 Fargate, 0 EC2

Draining service count 0 Fargate, 0 EC2

Services

Tasks

ECS Instances

Metrics

Scheduled Tasks

Create

Update

Delete

Last updated on November 14, 2018 3:45:04 PM (4m ago)



Filter in this page

Launch type

ALL

Service type

ALL



Service Name

Status

Service...

Task De...

Desired...

Runnin...

Launch...

Platfor...

No results

Create a Task Definition named "dspace" using "awsvpc" mode

Task Definition: dspace:11

View detailed information for your task definition. To modify the task definition, you need to create a new revision and then make the required changes to the task definition

Create new revision

Actions

BuilderJSON

Task Definition Name

dspace

Task Role

ecsTaskExecutionRole

Network Mode

awsvpc

If you choose <default>, ECS will start your container using Docker's default networking mode, which is Bridge on Linux and NAT on Windows. <default> is the only supported mode on Windows.

Compatibilities

EC2, FARGATE

Requires compatibilities

FARGATE

The Task Definition will contain 2 containers "dspace" and "dspacecdb"

Task execution IAM role

This role is required by tasks to pull container images and publish container logs to Amazon CloudWatch on your behalf. If you do not have the ecsTaskExecutionRole already, we can create one for you.

Task execution role

ecsTaskExecutionRole

Task size

The task size allows you to specify a fixed size for your task. Task size is required for tasks using the Fargate launch type and is optional for the EC2 launch type. Container level memory settings are optional when task size is set. Task size is not supported for Windows containers.

Task memory (MiB)

6144

Task CPU (unit)

2048

Task memory maximum allocation for container memory reservation

06144 shared of 6144 MiB

Task CPU maximum allocation for containers

02048 shared of 2048 CPU units

Task Placement

Constraint

No constraints

Container Definitions

	Container Name	Image	CPU Units	Hard/Soft memory limits (MiB)	Essential
▶	dspace	dspace/dspace:dspace	0	--/--	true
▶	dspacecdb	dspace/dspace-postgres	0	--/--	true

dspace container

This container will be built from dspace/dspace.

The awsvpc networking does not allow containers to be addressed by hostname, so a localhost reference must be used for the database. Fortunately, DSpace 6+ allows environment variables to override dspace.cfg and local.cfg settings. Note the override for db.url. Port 8080 is exposed.

Container Na...	Image	CPU Units	Hard/Soft memory limits (MiB)	Essential
▼ <div></div> dspace	dspace/dspace:dspace-6_x-jdk8-test	0	--/--	true

Details

Port Mappings

Environment Variables

Docker labels

Extra hosts

Host Port	Container Port	Protocol
8080	8080	tcp

Key	Value/ValueFrom
db.url	jdbc:postgresql://localhost:5432/dspace

Key	Value
No docker labels	

Hostname	IP address
No host entries	

Mount Points

Volumes from

Ulimits

Log Configuration

Container Path	Source Volume	Read only
No Mount Points		

Source Container	Read only
No volumes from	

Name	Soft limit	Hard limit
No ulimit		

Key	Value
awslogs-group	/ecs/dspace
awslogs-region	us-west-2
awslogs-stream-prefix	ecs

dspace db container

This container will be built from dspace/dspace-postgres-pgcrypto

dspacedb

dspace/dspace-postgres-pgcrypto

0

--/--

true

Details

Port Mappings

Host Port	Container Port	Protocol
5432	5432	tcp

Environment Variables

Key	Value/ValueFrom
No Environment Variables	

Docker labels

Key	Value
No docker labels	

Extra hosts

Hostname	IP address
No host entries	

Mount Points

Container Path	Source Volume	Read only
No Mount Points		

Volumes from

Source Container	Read only
No volumes from	

Ulimits

Name	Soft limit	Hard limit
No ulimit		

Log Configuration

Log driver: awslogs

Key	Value
awslogs-group	/ecs/dspace
awslogs-region	us-west-2
awslogs-stream-prefix	ecs

Select "Run Task" to start your task within the cluster you created.

Specify the Fargate launch type.

Select the VPC and Subnet associated with your cluster.

Click to Edit the Security Group so that port 8080 will be accessible.

Run Task

Select the cluster to run your task definition on and the number of copies of that task to run. To apply container overrides or target particular container instances, click Advanced Options.

Launch type ☒ FARGATE ☐ EC2 ⓘ

Task Definition dspace:10 ⓘ
This is your task definition. Note that it is versioned each time you update it.

Platform version LATEST ⓘ

Cluster tbcluster1 ⓘ
Select the cluster you created

Number of tasks 1

Task Group ⓘ

VPC and security groups

VPC and security groups are configurable when your task definition uses the awsvpc network mode.

Cluster VPC* vpc-0a032cc06f07a6968 (10.0.0.0/16) ⓘ

Subnets* subnet-076e093a16046ccf3 (10.0.0.0/24) | tbcluster1/Public - us-west-2 a assign ipv6 on creation: Disabled ⓘ
Select the VPC and subnet created for your cluster

Security groups* dspace-6567 ⓘ **Edit** ⓘ
Click Edit to enable access to port 8080 for the Cluster

Auto-assign public IP ENABLED ⓘ

▶ Advanced Options

Customize the Security Group then Click "Run Task".

Assigned security groups ☒ Create new security group
☐ Select existing security group

Security group name* dspace-6567 ⓘ

Description Wed Nov 14 2018 16:01:49 GMT-0800 (Pacif ⓘ

Inbound rules for security group

Type	Protocol	Port range	Source
HTTP	TCP	80	Anywhere
Custom TCP	TCP	8080	Anywhere

+ Add rule

Select the running task in order to get its IP address

Services

Tasks

ECS Instances

Metrics

Scheduled Tasks

Run new Task

Stop

Stop All

Last updated on November 14, 2018 4:11:36 PM (0m ago)

?

Desired task status:

Running

 Stopped

Filter in this page

Launch type ALL

< 1-1 > Page size 50

<input type="checkbox"/>	Task	Task definition	Container insta...	Last status	Desired status	Started By	Group	Launch type	Platform version
<input type="checkbox"/>	03db2c99-2faa-4...	dspace:10	--	PROVISIONING	RUNNING		family:dspace	FARGATE	1.2.0

Get the public IP address

Task : 03db2c99-2faa-4f6a-827f-586263a6db3f

Details

Logs

Cluster

tbcluster1

Launch type

FARGATE

Platform version

1.2.0

Task definition

dspace:10

Group

family:dspace

Task role

ecsTaskExecutionRole

Last status

RUNNING

Desired status

RUNNING

Created at

2018-11-14 16:11:35 -0800

Network

Network mode

awsvpc

ENI Id

eni-86593e6b

Subnet Id

subnet-076e093a16046ccf3

Private IP

10.0.0.149

Public IP

34.221.164.160

Mac address

02:bb:77:7d:d6:f6

Containers

Last updated on 11/14/2018 12:00:00 PM

	Name	Container Id	Status	Image	CPU
▶	dspace	b9e92d9e-492d-4e2d-951a-0337c69b263d	RUNNING	dspace/dspace:10	0
▶	dspace-db	ce10fa7d-9a85-40f5-a0af-c50b1efb12ce	RUNNING	dspace/dspace-db:10	0

Open ip-address:8080/xmlui

Not secure | 34.221.164.160:8080/xmlui/

MyZoom LIT Tech DG DataDog/NewRel LibService ASpace APT REST D7REST IIIF Alma LITStats Stats Talks

DSpace Repository

DSpace Home

DSpace Repository

DSpace is a digital service that collects, preserves, and distributes digital material. Repositories are important tools for preserving an organization's legacy; they facilitate digital preservation and scholarly communication.

Communities in DSpace

Select a community to browse its collections.

Search DSpace

[Advanced Search](#)




Browse

All of DSpace
[Communities & Collections](#)
[By Issue Date](#)
[Authors](#)
[Titles](#)
[Subjects](#)

My Account

[Login](#)
[Register](#)

RSS Feeds

 [RSS 1.0](#)
 [RSS 2.0](#)
 [Atom](#)

Outstanding Issues

Currently, we use docker exec to configure administrative permissions and to ingest content within a DSpace instance.

With Fargate deployment, it is unclear how to connect to an ECS container and issue commands within the container.

- <https://github.com/aws/amazon-ecs-cli/issues/143>