

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)

The screenshot shows the AWS Management Console interface for an ECS cluster. On the left, there is a navigation menu with 'Amazon ECS' selected, and sub-items for 'Clusters', 'Task Definitions', 'Amazon ECR', and 'Repositories'. The main content area shows the cluster 'tbcluster1' with a status of 'ACTIVE'. Below the status, there are several metrics: '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), and 'Draining service count' (0 Fargate, 0 EC2). There are tabs for 'Services', 'Tasks', 'ECS Instances', 'Metrics', and 'Scheduled Tasks'. The 'Services' tab is active, showing a table with columns for 'Service Name', 'Status', 'Service...', 'Task De...', 'Desired...', 'Runnin...', 'Launch...', and 'Platfor...'. The table is currently empty, displaying 'No results'. There are buttons for 'Create', 'Update', and 'Delete' at the top of the table, and a 'Delete Cluster' button in the top right corner.

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 ▾

Builder

JSON

Task Definition Name

Task Role [ecsTaskExecutionRole](#)

Network Mode ⓘ

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 "dspace-db"

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



Task CPU maximum allocation for containers



Task Placement

Constraint No constraints

Container Definitions

Container Name	Image	CPU Units	Hard/Soft memory limits (MiB)	Essential
■ dspace	dspace/dspace.dspa...	0	--	true
■ dspace-db	dspace/dspace-post...	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
dspace	dspace/dspace:dspace-6_x-jdk8-test	0	--/--	true

Details			Mount Points														
<table border="1"> <thead> <tr> <th>Host Port</th> <th>Container Port</th> <th>Protocol</th> </tr> </thead> <tbody> <tr> <td>8080</td> <td>8080</td> <td>tcp</td> </tr> </tbody> </table>			Host Port	Container Port	Protocol	8080	8080	tcp	<table border="1"> <thead> <tr> <th>Container Path</th> <th>Source Volume</th> <th>Read only</th> </tr> </thead> <tbody> <tr> <td colspan="3">No Mount Points</td> </tr> </tbody> </table>			Container Path	Source Volume	Read only	No Mount Points		
Host Port	Container Port	Protocol															
8080	8080	tcp															
Container Path	Source Volume	Read only															
No Mount Points																	
<table border="1"> <thead> <tr> <th>Key</th> <th>Value/ValueFrom</th> </tr> </thead> <tbody> <tr> <td>db.url</td> <td>jdbc:postgresql://localhost:5432/dspace</td> </tr> </tbody> </table>			Key	Value/ValueFrom	db.url	jdbc:postgresql://localhost:5432/dspace	<table border="1"> <thead> <tr> <th>Source Container</th> <th>Read only</th> </tr> </thead> <tbody> <tr> <td colspan="2">No volumes from</td> </tr> </tbody> </table>			Source Container	Read only	No volumes from					
Key	Value/ValueFrom																
db.url	jdbc:postgresql://localhost:5432/dspace																
Source Container	Read only																
No volumes from																	
<table border="1"> <thead> <tr> <th>Key</th> <th>Value</th> </tr> </thead> <tbody> <tr> <td colspan="2">No docker labels</td> </tr> </tbody> </table>			Key	Value	No docker labels		<table border="1"> <thead> <tr> <th>Name</th> <th>Soft limit</th> <th>Hard limit</th> </tr> </thead> <tbody> <tr> <td colspan="3">No ulimit</td> </tr> </tbody> </table>			Name	Soft limit	Hard limit	No ulimit				
Key	Value																
No docker labels																	
Name	Soft limit	Hard limit															
No ulimit																	
<table border="1"> <thead> <tr> <th>Hostname</th> <th>IP address</th> </tr> </thead> <tbody> <tr> <td colspan="2">No host entries</td> </tr> </tbody> </table>			Hostname	IP address	No host entries		<table border="1"> <thead> <tr> <th>Key</th> <th>Value</th> </tr> </thead> <tbody> <tr> <td>awslogs-group</td> <td>/ecs/dspace</td> </tr> <tr> <td>awslogs-region</td> <td>us-west-2</td> </tr> <tr> <td>awslogs-stream-prefix</td> <td>ecs</td> </tr> </tbody> </table>			Key	Value	awslogs-group	/ecs/dspace	awslogs-region	us-west-2	awslogs-stream-prefix	ecs
Hostname	IP address																
No host entries																	
Key	Value																
awslogs-group	/ecs/dspace																
awslogs-region	us-west-2																
awslogs-stream-prefix	ecs																

dspacedb container

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

dspace/db dspace/dspace-postgres-pgcrypto 0 --/-- true

Details

Port Mappings

Host Port	Container Port	Protocol
5432	5432	tcp

Environment Variables

Key	Value/ValueFrom
<i>No Environment Variables</i>	

Docker labels

Key	Value
<i>No docker labels</i>	

Extra hosts

Hostname	IP address
<i>No host entries</i>	

Mount Points

Container Path	Source Volume	Read only
<i>No Mount Points</i>		

Volumes from

Source Container	Read only
<i>No volumes from</i>	

Ulimits

Name	Soft limit	Hard limit
<i>No ulimit</i>		

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

The screenshot shows the AWS ECS console interface. At the top, there are tabs for Services, Tasks, ECS Instances, Metrics, and Scheduled Tasks. Below the tabs are buttons for 'Run new Task', 'Stop', and 'Stop All'. A status bar indicates 'Desired task status: Running Stopped' and 'Last updated on November 14, 2018 4:11:36 PM (0m ago)'. A filter input and a 'Launch type' dropdown are visible. Below this is a table with columns: Task, Task definition, Container insta..., Last status, Desired status, Started By, Group, Launch type, and Platform version. The first row is highlighted in yellow.

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

Get the public IP address

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

The screenshot shows the 'Details' view of a task in the AWS ECS console. It includes sections for 'Details', 'Network', and 'Containers'. The 'Details' section lists various attributes like Cluster, Launch type, Platform version, Task definition, Group, Task role, Last status, Desired status, and Created at. The 'Network' section lists Network mode, ENI Id, Subnet Id, Private IP, Public IP (highlighted in yellow), and Mac address. The 'Containers' section shows a table of running containers with columns for Name, Container Id, Status, Image, and CPU.

Details

- 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

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

Open ip-address:8080/xmlui

