



User Guide

AWS App2Container



AWS App2Container: User Guide

Copyright © 2025 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

Amazon's trademarks and trade dress may not be used in connection with any product or service that is not Amazon's, in any manner that is likely to cause confusion among customers, or in any manner that disparages or discredits Amazon. All other trademarks not owned by Amazon are the property of their respective owners, who may or may not be affiliated with, connected to, or sponsored by Amazon.

Table of Contents

What is AWS App2Container?	1
How App2Container works	1
Accessing AWS through App2Container	2
Pricing	2
Compatibility guide	3
Operating system compatibility	3
Containerization features	5
Deployment features	6
Pipeline support	8
Supported applications	9
Complex Windows .NET apps	13
Step 1: Setup and initialization	13
Step 2: Analysis phase	13
Step 3: Containerization	16
Step 4: Deployment	17
Getting started	28
Understand Docker containers	29
Decide where containerization will run	29
Prerequisites: Set up your servers	30
Sign up for AWS	31
Grant permissions to run AWS App2Container commands	31
Enable remote access for a worker machine (optional)	32
Configure your AWS profile	34
Install the Docker engine	35
Step 1: Install App2Container	37
Step 2: Initialize App2Container	40
Step 3: Analyze your application	43
Step 4: Transform your application	45
Step 5: Deploy your application	47
Step 6: Clean up	48
App2Container Automation runbook	50
Prerequisites	50
Create policies and roles for the automation	51
Attaching the IAM role	60

Run the automation	61
Runbook parameters	61
Running the automation	64
Reviewing output from the automation	64
Complete the modernization process	65
Configuring your application	66
Manage secrets	66
Create remote access secrets	67
Create secrets for Jenkins pipelines	69
Create secrets for Microsoft Azure DevOps pipelines	72
Configure containers	74
Configure deployment	94
deployment.json file	94
Configure pipelines	113
pipeline.json file	113
Product and service integrations	123
Automatic storage and registration using Amazon Elastic Container Registry	123
Deploy to Amazon ECS	124
Prerequisites	125
Amazon ECS integration for App2Container workflow	126
Deploy to Amazon EKS	129
Prerequisites	129
Amazon EKS integration for App2Container workflow	129
Deploy to App Runner	132
Prerequisites	133
App Runner integration for App2Container workflow	133
Set up CodePipeline pipelines	135
Validation	136
Output	137
Set up Jenkins pipelines	138
Prerequisites	138
Jenkins integration for App2Container workflow	139
Set up Azure DevOps pipelines	141
Prerequisites	142
Azure DevOps integration for App2Container workflow	144
Route logs using FireLens	149

FireLens log routing for Linux	149
Security	162
Data protection	162
Data encryption	163
Internet network traffic privacy	164
Identity and access management	164
Create IAM resources for general use	166
Create IAM resources for deployment	181
Update management	182
Command reference	183
Containerization phases	183
Initialize	183
Analyze	184
Transform	185
Deploy	186
Utility commands	187
analyze	187
Syntax	188
Parameters and options	188
Output	188
Examples	189
containerize	190
Syntax	190
Parameters and options	190
Output	191
Examples	192
extract	194
Syntax	194
Parameters and options	195
Output	195
Examples	195
generate app-deployment	196
Syntax	198
Parameters and options	198
Output	199
Examples	201

generate pipeline	205
Syntax	207
Parameters and options	207
Output	208
Examples	210
help	216
Syntax	217
Parameters and options	217
Output	217
Examples	217
init	218
Syntax	218
Parameters and options	218
Output	219
Examples	220
inventory	222
Syntax	222
Parameters and options	222
Output	223
Examples	223
remote analyze	225
Syntax	225
Parameters and options	226
Output	226
Examples	227
remote configure	228
Syntax	228
Parameters and options	228
Input	228
Output	230
Examples	230
remote extract	231
Syntax	232
Parameters and options	232
Output	232
Examples	233

remote inventory	233
Syntax	234
Parameters and options	234
Output	230
Examples	230
upgrade	237
Syntax	237
Options	238
Output	238
Examples	238
upload-support-bundle	238
Syntax	239
Options	239
Output	239
Examples	239
Troubleshooting	241
Access App2Container logs on your server	241
Access application logs inside of a running container	242
AWS resource creation fails for the generate command	242
Description	242
Cause	243
Solution	243
Troubleshoot Java applications on Linux	243
Troubleshoot .NET applications on Windows	245
Troubleshoot generate pipeline build for Jenkins	246
Release notes	248
Document history	277

What is AWS App2Container?

AWS App2Container (A2C) is a command line tool to help you lift and shift applications that run in your on-premises data centers or on virtual machines, so that they run in containers that are managed by Amazon ECS, Amazon EKS, or AWS App Runner. For a console-based experience, you can use the *Replatform applications to Amazon ECS* template in the [AWS Migration Hub Orchestrator console](#). For more information, see [Replatform applications](#) to Amazon ECS in the *AWS Migration Hub Orchestrator User Guide*.

Moving legacy applications to containers is often the starting point toward application modernization. There are many benefits to containerization:

- Reduces operational overhead and infrastructure costs
- Increases development and deployment agility
- Standardizes build and deployment processes across an organization

Contents

- [How App2Container works](#)
- [Accessing AWS through App2Container](#)
- [Pricing](#)

How App2Container works

You can use App2Container to generate container images for one or more applications running on Windows or Linux servers that are compatible with the Open Containers Initiative (OCI). This includes commercial off-the-shelf applications (COTs). App2Container does not need source code for the application to containerize it.

You can use App2Container directly on the application servers that are running your applications, or perform the containerization and deployment steps on a worker machine.

App2Container performs the following tasks:

- Creates an inventory list for the application server that identifies all running ASP.NET (Windows) and Java applications (Linux) that are candidates to containerize.

- Analyzes the runtime dependencies of supported applications that are running, including cooperating processes and network port dependencies.
- Extracts application artifacts for containerization and generates a Dockerfile.
- Initiates builds for the application container.
- Generates AWS artifacts and optionally deploys the containers on Amazon ECS, Amazon EKS, or AWS App Runner. For example:
 - a CloudFormation template to configure required compute, network, and security infrastructure to deploy containers using Amazon ECS, Amazon EKS, or AWS App Runner.
 - An Amazon ECR container image, Amazon ECS task definitions, or AWS CloudFormation templates for Amazon EKS or AWS App Runner that incorporate best practices for security and scalability of the application by integrating with various AWS services.
 - When deploying directly, App2Container can upload AWS CloudFormation resources to an Amazon S3 bucket, and create a CloudFormation stack.
- Optionally creates a CI/CD pipeline with AWS CodePipeline and associated services, to automate building and deploying your application containers.

Accessing AWS through App2Container

When you initialize App2Container, you provide it with your AWS credentials. This allows App2Container to do the following:

- Store artifacts in Amazon S3, if you configured it to do so.
- Create and deploy application containers using AWS services such as Amazon ECS, Amazon EKS, and AWS App Runner.
- Create CI/CD pipelines using AWS CodePipeline.

Pricing

App2Container is offered at no additional charge. You are charged only when you use other AWS services to run your containerized application, such as Amazon ECR, Amazon ECS, Amazon EKS, and AWS App Runner. For more information, see [AWS Pricing](#).

App2Container compatibility

The following documentation provides information for the operating systems, software, and tooling that you can use with App2Container.

Contents

- [Operating system compatibility](#)
- [Containerization features](#)
- [Deployment features](#)
- [Pipeline support](#)

Operating system compatibility

The following table contains information about the applications that App2Container supports for each operating system.

Compatibility item	Linux	Windows
Supported application server operating systems ¹	• Ubuntu (version 18.04 and later) • CentOS (version 8 and later) • RHEL (version 7 and later) • Amazon Linux 2 (AL2) • Amazon Linux 2023 (AL2023)	• Windows Server 2016 and later ²
Container hosts	The container host can be any supported application server operating system. The major kernel version of the container host must match with the container image.	The container host operating system must be either Windows Server 2016, 2019, or 2022. The Windows Server operating system version of the container host must match the container image.

Compatibility item	Linux	Windows
		App2Container automatically deploys the container host using the same operating system used for the container ization process.
Application types	• Java applications • .NET applications	• IIS .NET applications
Supported frameworks	• Java (JDK 1.8 and later) • Tomcat • TomEE • JBoss (standalone mode) • .NET applications • .NET Core 3.1 • .NET 5 • .NET 6 • .NET 7 • .NET 8 • .NET 9	.NET Framework version 3.5 and 4.x
Unsupported application features	High Availability (HA) clusters	• IIS applications that use files and registries outside of IIS web application directories

Compatibility item	Linux	Windows
Additional system requirements	- Docker version 17.07 and later ³ - kubectl versions up to v1.30 for Amazon EKS deployments.	- Docker version 17.07 and later ³ - kubectl versions up to v1.30 for Amazon EKS deployments. - Windows IIS (7.5 and later) - Windows PowerShell version 5.1 or PowerShell version 6 and later

¹ We have only tested the operating systems and configurations listed. Other operating systems could be compatible, but have not been tested.

² App2Container v1.47 is the last version that supports Windows 2008 and 2012. These operating systems also require a worker machine. We recommend using a later version of the Windows Server operating system to be able to use the latest version of App2Container. For more information, see [Applications you can containerize using AWS App2Container](#).

³ Docker must be installed to use App2Container. For more information, see [Prerequisites: Set up your servers](#).

Note

Windows client operating systems such as Windows 7 and Windows 10 aren't supported.

Containerization features

App2Container supports the following containerization features.

Containerization feature	Linux	Windows
gMSA for connection with Active Directory	Not supported	Supported

Containerization feature	Linux	Windows
Containerization of multiple applications in the same container	Not supported	Supported *
Containerization of applications that use multiple ports	Not supported	Supported

* Containerizing multiple applications in the same container for Windows requires that the applications are nested under a main IIS site.

For more information about configuring Windows containers with additional ports and multiple applications, see [Configuring application containers](#).

For more information about group managed service accounts (gMSAs), see [Configuring container deployment](#).

Deployment features

The following table lists the deployment services that App2Container supports.

Deployment feature	Linux			Windows		
	Amazon ECS (AWS Fargate only)	Amazon EKS (Amazon EC2 only)	AWS App Runner	Amazon ECS (AWS Fargate)	Amazon ECS (Amazon EC2)	Amazon EKS (Amazon EC2 only)
Modify memory usage	Supported	Supported	Not supported	Supported ¹	Supported	Supported
Modify CPU usage	Supported	Supported	Not supported	Supported	Supported	Supported

Deployment feature	Linux			Windows		
	Application Load Balancer	Application Load Balancer, Network Load Balancer with Nginx	N/A	Application Load Balancer	Application Load Balancer	Application Load Balancer, Network Load Balancer with Nginx
Load balancer types						
Reuse VPC ²	Supported	Supported	Not supported	Supported	Supported	Supported
Reuse cluster previously deployed with App2Container ²	Supported	Supported	N/A	Supported	Supported	Supported
FireLens logging	Supported	Not supported	Not supported	Not supported	Not supported	Not supported
gMSA for connection with Active Directory	Not supported	Not supported	Not supported	Not supported	Supported	Supported
Deploy complex .NET applications ³	N/A	N/A	N/A	Supported	Supported	Supported

¹ AWS Fargate only supports certain Windows Server operating systems for running Windows containers. Select a Windows Server operating system that both Fargate and App2Container

support. For more information, see [Windows platform versions](#) in the *Amazon ECS User Guide for AWS Fargate*.

² You can reuse certain components that App2Container created for a prior deployment. For more information about the `reuseResources` object, see [Configuring container deployment](#).

³ A complex .NET application has multiple Windows .NET application components running in a single container. For more information, see [Containerizing complex Windows .NET applications with App2Container](#).

For more information about FireLens for Amazon ECS, see [Custom log routing](#) in the *Amazon Elastic Container Service Developer Guide*.

For more information about deployment settings for group managed service accounts (gMSAs), see [Configuring container deployment](#).

Pipeline support

App2Container supports AWS CodePipeline, Jenkins, and Azure DevOps Services pipeline types for both Windows and Linux. For more information about configuring pipelines, see [Configuring container pipelines](#) and [Examples](#).

Applications you can containerize using AWS App2Container

App2Container supports the following application types:

- Java applications (Linux)
- ASP.NET applications (Windows, Linux)

For supported application frameworks, App2Container targets only the application files and dependencies that are needed for containerization, thereby minimizing the size of the resulting container image. This is known as *application mode*.

If App2Container does not find a supported framework running on your application server, or if you have other dependent processes running on your server, App2Container takes a conservative approach to identifying dependencies. This is known as *process mode*. For process mode, all non-system files on the application server are included in the container image.

For more details on application and framework support, expand the section that matches the platform that your application runs on.

Important

App2Container does not containerize database layer components. If your application requires access to a database, you must configure your application container to have access to the database server.

Supported applications for Linux

App2Container supports identification and containerization of Java and ASP.NET applications running on Linux.

Supported Linux distributions:

- Ubuntu
- CentOS

- RHEL
- Amazon Linux

For supported frameworks, and other language-specific details, choose the tab that matches the language your application is written in.

Java

For Java applications, App2Container identifies Java processes, and can generate container images that replicate the running state of each process. App2Container determines which files to include in the application container image, based on the Java application framework.

Application mode is supported for the following Java application frameworks:

Supported frameworks

- Tomcat
- TomEE
- JBoss (standalone mode)

Note

Containerization is not supported for Java applications running on frameworks that are using Cluster/HA mode.

ASP.NET

For ASP.NET applications running on Linux, App2Container detects the .NET runtime version and containerizes the application using the corresponding runtime base images.

Supported .NET Core runtime versions

- **.NET Core 3.1** – uses [SDK version 3.1](#) as the base image for generic .NET Core applications (or the highest version if multiple versions are used).
- **.NET 5** – uses [SDK version 5.0](#) as the base image for generic .NET Core applications (or the highest version if multiple versions are used).

- **.NET 6** – uses [SDK version 6.0](#) as the base image for generic .NET Core applications (or the highest version if multiple versions are used).
- **.NET 7** – uses [SDK version 7.0](#) as the base image for generic .NET Core applications (or the highest version if multiple versions are used).
- **.NET 8** – uses [SDK version 8.0](#) as the base image for generic .NET Core applications (or the highest version if multiple versions are used).

Important

- Process mode is not supported for ASP.NET applications running on .NET Core.
- If you are using .NET Core 3.1 or .NET 5, you must update the `analysis.json` file's `containerBaseImage` parameter to `mcr.microsoft.com/dotnet/sdk:3.1` or `mcr.microsoft.com/dotnet/sdk:5.0`, respectively. For more information, see [Configuring application containers](#).

Supported applications for Windows

App2Container supports containerization of ASP.NET applications deployed on IIS, including IIS-hosted WCF applications, running on Windows Server 2016, 2019, and 2022. It uses Windows Server Core as a base image for its container artifacts, matching the Windows Server Core version to the operating system (OS) version of the server where you run containerization commands.

If you use a worker machine to containerize your application, the version matches your worker machine OS. If you are running containerization directly on application servers, the version matches your application server OS.

If your applications are running on Windows Server 2008 or 2012 R2, you might still be able to use App2Container by setting up a worker machine for containerization and deployment steps. App2Container does not support applications running on Windows client operating systems, such as Windows 7 or Windows 10.

Application framework and system requirements

- Containerization commands must run on Windows OS versions that support containers—Windows Server 2016, 2019, or 2022. This can be the worker machine, if you configure one, or the application server.

- If you use a worker machine to run containerization commands, App2Container supports Windows Server 2008 and up for the application server.

Note

Windows Server 2008 and 2012 only support App2Container up to version 1.47. For more information, see [Release notes for AWS App2Container](#).

- IIS 7.5 or later.
- .NET framework version 3.5 or later.
- Docker version 17.07 or later (to install).

Note

App2Container does not support applications running on Windows client operating systems, such as Windows 7 or Windows 10.

Supported applications

- Simple ASP.NET applications running in a single container
- A Windows service running in a single container
- Complex ASP.NET applications that depend on WCF, running in a single container or multiple containers
- Complex ASP.NET applications that depend on Windows services or processes outside of IIS, running in a single container or multiple containers
- Complex, multi-node IIS or Windows service applications, running in a single container or multiple containers

Unsupported applications

- ASP.NET applications that use files and registries outside of IIS web application directories
- ASP.NET applications that depend on features of a Windows operating system version prior to Windows Server Core 2016

Containerizing complex Windows .NET applications with App2Container

Containerization for complex multi-tier Windows .NET applications requires careful planning. When functionality is shared between the root application and one or more lower-level or system applications, you need to make decisions about packaging, deployment, and orchestration for all of the components.

To summarize how AWS App2Container works to containerize a complex Windows .NET application, we'll visit each step in the App2Container workflow, and call out the highlights and things to consider.

Step 1: Setup and initialization

Setup and initialization are the same for complex Windows .NET applications as for other types of applications. Setup tasks include installing software, configuring your AWS profile and IAM permissions, and deciding which servers the App2Container commands should run on. To learn more about setting up your environment before running App2Container for the first time, see [Prerequisites: Set up your servers](#).

After you have completed the setup tasks, but before you use App2Container for the first time, you must initialize the servers where you plan to run App2Container commands. To learn more about initialization and worker machine configuration, see the [Initialize](#) section in the [App2Container command reference](#).

Step 2: Analysis phase

After you have completed setup and initialization tasks on your servers, App2Container helps you to take an inventory of your running applications, and perform analysis to determine what should be included in your application containers.

Inventory

The first step in the analysis phase is to take an inventory of your applications. When you run the **app2container inventory** command (or the **app2container remote inventory** command, if you have configured a worker machine), App2Container detects the applications that are running in IIS. It also detects the Windows services that could be configured as dependent application components.

App2Container identifies each IIS application or Windows service as a separate application, with its own application ID in the `inventory.json` file. App2Container makes an effort to exclude basic operating system services that you would not want to add to your containers. However, even when these services are excluded, the inventory list can still be quite long.

To narrow the results of the **app2container inventory** or **app2container remote inventory** commands, you can specify what type of application you are looking for with the `--type` option:

- To run an inventory of your IIS applications, you can set the `--type` option to "iis".
- To run an inventory of your Windows services, you can set the `--type` option to "service".

If you don't want App2Container to filter inventory results at all, you can use the `--nofilter` option. This option prevents App2Container from filtering out default system services when building the inventory list. For more information and command syntax, see the [inventory](#) or [remote inventory](#) command in the command reference section.

Analysis

When you run the **app2container analyze** or **app2container remote analyze** commands, App2Container analyzes the application component that you specify with the `--application-id` parameter.

App2Container creates the folder structure for the application component, inside of the App2Container directory on your application server or worker machine. It produces the `analysis.json` file, and saves it to the new folder structure, along with other artifacts that are required for containerization. The `analysis.json` file is where you begin to define your container structure.

Tip

Run the **app2container analyze** or **app2container remote analyze** command for every component in your multi-tier application before you configure your container structure.

You can implement the following container structures for a multi-tier Windows .NET application:

- **Multiple application components running in separate containers (recommended)**

In this scenario, each application component in your multi-tier Windows .NET application runs in a separate container. Relationships between the root application and up to two dependent applications are configured in the `deployment.json` file for the root application. This file is produced during the containerization phase.

When your application components are running in separate containers, leave the `additionalApps` array in the `analysis.json` file empty for all components.

- **Multiple application components running in a single container**

In this scenario, the application components in a multi-tier Windows .NET application run together in one container. We recommend that packaging multiple application components in a single container is only done when there are cross-dependencies between the components.

To specify multiple application components running in a single container, you can include up to five dependent component application IDs in the `additionalApps` array in the `analysis.json` file for the root application.

 Note

This configuration has the following limitations:

- Only the port that is defined for the root application is exposed to outside traffic through your load balancer. Ports that are defined for other application components are exposed only from the container, and are not accessible through the load balancer.
- If you are using remote commands on a worker machine, all of the application components in a multi-tier application must be running on the same application server if you want them to run in a single container.

To learn more about configuring containers, see [Configuring application containers](#). To compare configuration examples for a simple .NET application, and for complex multi-tier .NET applications, expand the **Containers running on Windows** section, and explore the example tabs.

For more information and command syntax, see the [analyze](#) or [remote analyze](#) command in the command reference section.

Step 3: Containerization

This phase creates containers for your application, based on the output of the analysis phase and on your configuration in the `analysis.json` file.

Extract

If you are using a worker machine to run App2Container commands, or if you want to store an application archive for reference, this phase starts with an **app2container extract** or **app2container remote extract** command. Because this has no effect on the configuration for multi-tier application containers, we will not cover that here.

Containerize

The **app2container containerize** command performs the following tasks for the application that's specified in the `--application id` parameter:

- Extracts application artifacts from the server it runs on, or reads from an extract archive. For complex multi-tier applications, the extract includes all artifacts that are needed for all of the components running in the container.
- Generates a Dockerfile and a container image, based on the application artifacts and the application settings in the `analysis.json` file.
- Creates the `deployment.json` file that defines initial settings for container deployment during the deployment phase.

You must run the **app2container containerize** command for the root application container, and for each additional application component that runs in a separate container. Do not run the command for any components that are included in the root application container. The command displays real-time task completion messages, followed by instructions for next steps. This includes the AWS commands that you run if you are deploying manually.

To configure the `deployment.json` file for a complex multi-tier application, refer to the following scenario that describes your implementation:

- **Multiple application components running in separate containers**

In this scenario, each application component is running in a separate container, and each has its own deployment file. Before running the **generate app-deployment** command, configure the

deployment.json file for the root application to include all dependent applications or services in the dependentApps array, including the application ID, private root domain, and DNS record name for each one.

- **Multiple application components running in a single container**

If you are running multiple application components in a single container, the process for configuring the deployment.json file is the same as for any other containerized application. Leave the dependentApps array empty.

Note

If you are deploying to a specific VPC, make sure that all components point to that VPC in the vpcId parameter in the reuseResources array in the deployment.json file.


To learn more about configuring your deployment.json file, see [Configuring container deployment](#). For more information and command syntax for creating your application container, see the [containerize](#) command in the command reference section.

Step 4: Deployment


Deployment steps for complex Windows .NET applications with multiple application components running in a single container are handled the same as any other application deployment. For more information and command syntax for deploying your application container, see the [generate app-deployment](#) command in the command reference section.

The remainder of the content in this section applies to complex Windows .NET applications that have multiple application components running in separate containers, similar to the application example shown in the following diagrams:

Amazon ECS deployment



Amazon EKS deployment



Normally, you run the **generate app-deployment** command for each application container that you create. However, with complex Windows .NET applications that have dependent applications running in separate containers, App2Container takes care of some of that for you. When you run the **generate app-deployment** command for the root application, App2Container completes the following tasks for the root application *and each of its dependent application components*:

- Checks for AWS and Docker prerequisites.
- Creates an Amazon ECR repository.
- Pushes the container image to the Amazon ECR repository.
- Generates the following artifacts, depending on your target container management service:

Amazon ECS

- An Amazon ECS task definition.
- The `ecs-master.yml` file that you can use for Amazon ECS deployment.

Amazon EKS

- The Kubernetes `eks-master.yml` file that you can use for Amazon EKS deployment.
- The `eks_deployment.yaml` and `eks_service.yaml` files that you can use with the **kubectl** command.
- Generates a `pipeline.json` file.

Additionally, if you use the `--deploy` option, App2Container takes care of all of those deployments in the order in which they need to run, and configures shared infrastructure settings. When App2Container handles the deployment for you, it follows these conventions:

- The root application and all dependent application components are deployed to the same cluster.
- All dependent application components are configured with an internal load balancer only.
- Each application component has its own Amazon ECS or Amazon EKS service running in a shared cluster.

If you want to customize the deployment artifacts, you can deploy manually, using the AWS Management Console or AWS CLI when you are ready.

For deployment steps, choose the tab that matches your deployment scenario.

Automated (A2C)

Follow these steps if you are using the App2Container automated deployment.

1. Verify that the values are set correctly in the `deployment.json` files for all of your application components, before running the **generate app-deployment** command for your root application, as follows:
 - None of the application components in the multi-tier application should specify `reuseCfnStack`.
 - Dependent application components should not specify any of the following parameters: `vpcId`, `gMSAParameters`.
 - The following parameters can be specified in the root application, and App2Container applies the same values for all dependent application components: `vpcId`, `resourceTags`, and `gMSAParameters`.
2. The following example shows the **generate app-deployment** command for the root application in our sample multi-tier application, using the `--deploy` option, with the `--application-id` parameter set to the application ID for the root application. This example handles the full deployment for all application components.

```
PS> app2container generate app-deployment --deploy --application-id iis-  
colormvciis-b69c09ab --profile admin-profile  
✓ AWS prerequisite check succeeded  
✓ Docker prerequisite check succeeded  
... [more notifications as deployment steps are completed for each dependent  
application component, followed by the root application and shared  
configurations]  
Deployment successful for application iis-colormvciis-b69c09ab
```

The URL to your Load Balancer Endpoint is:
`a2c-i-Publi-1A2BCD3EFGRW-4567890123.us-west-2.elb.amazonaws.com`

Successfully created Amazon ECS stack `a2c-iis-colormvciis-b69c09ab-ECS`. Check the AWS CloudFormation Console for additional details.

3. Set up a pipeline for your application stack using app2container:

```
app2container generate pipeline --application-id iis-colormvciis-  
b69c09ab
```

The first deployment for a dependent application component creates shared AWS resources, such as the VPC and Amazon ECS or Amazon EKS cluster. After the first dependent application component is successfully deployed, App2Container updates deployment artifacts for all of the other application components to reference the shared AWS resources prior to completing the remaining deployments.

Manual (AWS CLI)

Follow these steps to customize your deployment files and use the AWS CLI to deploy manually. We do not include AWS Management Console instructions here. However, you can follow the same general order of operations in the console.

1. Verify that the values are set correctly in the deployment .json files for all of your application components, before running the **generate app-deployment** command for your root application, as follows:
 - None of the application components in the multi-tier application should specify `reuseCfnStack`.
 - Dependent application components should not specify any of the following parameters: `vpcId`, `gMSAParameters`.
 - The following parameters can be specified in the root application, and App2Container applies the same values for all dependent application components: `vpcId`, `resourceTags`, and `gMSAParameters`.
2. The following example shows the **generate app-deployment** command for the root application in our sample multi-tier application, with the `--application-id` parameter set to the application ID for the root application. The `--deploy` option is not used in this case, as we plan to customize deployment files and then deploy using AWS CLI commands to control deployment for each application component.

Note

App2Container creates deployment artifacts for all application components in the complex Windows .NET application when you run the **generate app-deployment** command for the root application.

Use the **generate app-deployment** command, specifying the application ID for your root application, as follows:

```
PS> app2container generate app-deployment --application-id iis-colormvciis-b69c09ab --profile admin-profile
✓ AWS prerequisite check succeeded
✓ Docker prerequisite check succeeded
... [more notifications as deployment steps are completed for each dependent component, followed by the root application and shared configurations]
CloudFormation templates and additional deployment artifacts generated successfully for application iis-colormvciis-b69c09ab
```

You're all set to use AWS CloudFormation to manage your application stack.

Next Steps:

1. Create application stacks for first dependent application using the AWS CLI or the AWS Console. AWS CLI commands:

```
aws cloudformation deploy --template-file C:\Users\Administrator\AppData\Local\app2container\iis-dependentappb-12345bcd\EcsDeployment\ecs-master.yml --capabilities CAPABILITY_NAMED_IAM CAPABILITY_AUTO_EXPAND --stack-name a2c-iis-dependentappb-12345bcd-ECS
```

2. Required! Reuse the VpcId, ClusterId and PublicSubnets from above CloudFormation console outputs and assign them in master templates of **service-colorwindowsservice-69f90194**, **iis-colormvciis-b69c09ab**

If your other dependent application(s) that share the same root domain, also assign HostedZoneId to their master template(s).

Create application stacks for remaining applications using the AWS CLI or the AWS Console. AWS CLI commands:

```
aws cloudformation deploy --template-file C:\Users\Administrator\AppData\Local\app2container\service-colorwindowsservice-69f90194\EcsDeployment\ecs-master.yml --capabilities CAPABILITY_NAMED_IAM CAPABILITY_AUTO_EXPAND --stack-name a2c-service-colorwindowsservice-69f90194-ECS
```

```
aws cloudformation deploy --template-file C:\Users\Administrator\AppData\Local\app2container\iis-colormvciis-b69c09ab\EcsDeployment\ecs-master.yml --capabilities CAPABILITY_NAMED_IAM CAPABILITY_AUTO_EXPAND --stack-name a2c-iis-colormvciis-b69c09ab-ECS
```

3. Set up a pipeline for your application stack using app2container:

```
app2container generate pipeline --application-id iis-colormvciiis-b69c09ab
```

3. Review the deployment artifacts that were generated in the prior step, and customize the YAML deployment templates and other deployment artifacts as needed.

Manual deployment follows this step, beginning with one of the dependent applications. The first deployment creates any shared infrastructure that is required.

 Note

If you are using an existing VPC, the `vpcId` that you specified in the deployment `.json` file for the root application should be reflected in the YAML deployment templates for all of the dependent applications.

4. To deploy your first dependent application and create shared infrastructure, run the following command in the AWS CLI, using your dependent application's details.

```
PS> aws cloudformation deploy --template-file C:\Users\Administrator\AppData\Local\app2container\iis-dependentappb-12345bcd\EcsDeployment\ecs-master.yml --capabilities CAPABILITY_NAMED_IAM CAPABILITY_AUTO_EXPAND --stack-name a2c-iis-dependentappb-12345bcd-ECS
```

5. After your first stack is ready (stack status is `CREATE_COMPLETE`), update the YAML deployment templates for all remaining application components in your application to reference the following shared infrastructure in the parameters for existing resources:

- `VpcId`
- `PublicSubnets`
- `ClusterId`

Additionally, for any remaining dependent applications, update the following references:

- `DomainName`
- `RecordName`
- `ExistingHostedZoneId` – update this if dependent applications share the root domain, or if they are using an existing domain.

- RecordExist – set this to "true" (string) if the record already exists in the hosted zone. If you are creating a new domain, set this to "false". The default value is "true".
6. Deploy any remaining dependent applications, using your application component information and the updated YAML deployment templates, with the **cloudformation deploy** command. The following command example deploys the service component in our sample multi-tier application.

```
PS> aws cloudformation deploy --template-file C:\Users\Administrator\AppData
\Local\app2container\service-colorwindowsservice-69f90194\EcsDeployment\ecs-
master.yml --capabilities CAPABILITY_NAMED_IAM CAPABILITY_AUTO_EXPAND --stack-
name a2c-service-colorwindowsservice-69f90194-ECS
```

7. After you've created all of your dependent component stacks, deploy your root application with the **cloudformation deploy** command. The following command example deploys the root application in our sample multi-tier application.

```
PS> aws cloudformation deploy --template-file C:\Users\Administrator\AppData
\Local\app2container\iis-colormvc-iis-b69c09ab\EcsDeployment\ecs-master.yml --
capabilities CAPABILITY_NAMED_IAM CAPABILITY_AUTO_EXPAND --stack-name a2c-iis-
colormvc-iis-b69c09ab-ECS
```

Tip

It can take a few minutes to spin up a CloudFormation stack, along with the other infrastructure that is created for your deployment. You can use one of the following methods to check the stack status for your deployment:

- Sign in to the AWS Management Console and open the AWS CloudFormation console at <https://console.aws.amazon.com/cloudformation>.

In the console, you can see stacks that are being created, as well as existing stacks. For more information, see [Viewing AWS CloudFormation stack data and resources on the AWS Management Console](#) in the *AWS CloudFormation User Guide*.

- Use one of these AWS CloudFormation commands in the AWS CLI: **list-stacks** or **describe-stacks**. For more information, see **Available Commands** in the [AWS CLI Command Reference](#).

- Use one of these AWS CloudFormation API commands: **ListStacks** or **DescribeStacks**. For more information, see [Actions](#) in the *AWS CloudFormation API Reference*.

Getting started with AWS App2Container

AWS App2Container is a tool that helps you break down the work of moving your applications into containers, and configuring them to be hosted in AWS using the Amazon ECS, Amazon EKS, or App Runner container management services.

The following sections demonstrate the initial setup of your containerization environment, starting with prerequisites and initial workflow decisions. Then we take you step by step through containerizing a basic application using App2Container. We generate the artifacts that you can use to deploy it on Amazon ECS, Amazon EKS, or AWS App Runner, and then we clean up.

Note

To avoid creating billable AWS resources, we stop before the final deployment. You can review the deployment artifacts that are created by the **generate app-deployment** command to see what we would create.

For an overview of the command phases that includes summary information and command reference links for all of the App2Container commands, see the [App2Container command reference](#).

Contents

- [Understand Docker containers](#)
- [Decide where containerization will run](#)
- [Prerequisites: Set up your servers](#)
- [Step 1: Install App2Container](#)
- [Step 2: Initialize App2Container](#)
- [Step 3: Analyze your application](#)
- [Step 4: Transform your application](#)
- [Step 5: Deploy your application](#)
- [Step 6: Clean up](#)

Understand Docker containers

The following resources can help you get the most out of your application containers by understanding what goes into them.

- To learn more about Docker containers on AWS, see [What is Docker?](#).
- Use the Docker command line reference to look up Docker commands. See [Use the Docker command line](#).

Decide where containerization will run

To use App2Container on the server where the applications are running, you must set up an AWS profile, install App2Container, and install the Docker engine. If your server does not meet the requirements to containerize your application and deploy it to AWS, or if you do not want to install the Docker engine on the application server, you can set up and use a worker machine. On the worker machine, you can run the steps to containerize your application and deploy it to AWS, or you can set up connectivity between the worker machine and the application servers to run remote commands from the worker machine, targeting the application servers.

The following are example situations where you might decide to set up a worker machine:

- Your application servers are running in an on-premises data center and they do not have internet access.
- Your application server is running on a Windows operating system that does not support containers. For more information, see [Supported applications](#).
- You prefer to use a dedicated server to run the containerization and deployment steps.
- You want to consolidate your work by using a worker machine to run commands for all of your application servers.

When you set up a worker machine to handle the steps to containerize and deploy your applications, it must have the same operating system platform as your application server (Linux or Windows), and the operating system must support containers. We recommend that you launch an Amazon EC2 instance as the worker machine, using an Amazon Machine Image (AMI) that is optimized for Amazon ECS.

Prerequisites: Set up your servers

Before you use App2Container for the first time, make sure that your application environment meets all of the requirements that are listed for your operating system (OS) platform in the [Supported applications](#) section of this guide.

Choose the tab that matches your operating system (OS) platform to continue:

Linux

To run on a Linux platform, App2Container has the following additional requirements for the servers where you run App2Container commands. This includes application servers and the worker machine, if you have one configured.

- There are one or more Java applications running on each application server whose inventory is the subject of the **analyze** or **remote analyze** command.
- You have root access on the servers.
- The servers have **tar** and at least 20 GB of free space.

Windows

To run on a Windows platform, App2Container has the following additional requirements for the servers where you run App2Container commands. This includes application servers and the worker machine, if you have one configured.

- There are one or more applications running in IIS on each application server whose inventory is the subject of the **analyze** or **remote analyze** command.
- You are logged in as the Administrator user on the servers.
- The servers have Windows PowerShell version 5.1 or PowerShell version 6 or later and at least 20-30 GB of free space.

Complete the following tasks before you use App2Container for the first time.

- [Sign up for AWS](#)
- [Grant permissions to run AWS App2Container commands](#)
- [Enable remote access for a worker machine \(optional\)](#)
- [Configure your AWS profile](#)

- [Install the Docker engine](#)

Sign up for AWS

When you sign up for Amazon Web Services (AWS), your AWS account is automatically signed up for all services in AWS. You are charged only for the services that you use.

If you do not have an AWS account already, use the following procedure to create one.

To create an AWS account

1. Open <https://portal.aws.amazon.com/billing/signup>.
2. Follow the online instructions.

Part of the sign-up procedure involves receiving a phone call or text message and entering a verification code on the phone keypad.

When you sign up for an AWS account, an *AWS account root user* is created. The root user has access to all AWS services and resources in the account. As a security best practice, assign administrative access to a user, and use only the root user to perform [tasks that require root user access](#).

Grant permissions to run AWS App2Container commands

App2Container needs access to AWS services in order to run most of its commands. There are two very different sets of permissions needed to run **app2container** commands.

- The general purpose IAM user, group, or role can run all of the commands *except* commands that are run with the `--deploy` option.
- For deployment, App2Container must be able to create or update AWS objects for container management services (Amazon ECR with Amazon ECS, Amazon EKS, or AWS App Runner), and to create CI/CD pipelines with AWS CodePipeline. This requires elevated permissions that should only be used for deployment.

We recommend that you create general purpose IAM resources, and if you plan to use App2Container to deploy your containers or create pipelines, that you create separate IAM resources for deployment.

For instructions on how to set up your IAM resources for App2Container, and policy examples that include resources and actions that App2Container needs access to, see [Identity and access management in App2Container](#).

Note

You can use an instance profile to pass an IAM role to an Amazon EC2 instance. App2Container detects if there is an instance profile associated with the application server or worker machine when you run the **init** command. If it detects an instance profile, the **init** command prompts if you want to use it. To find out more about using instance profiles, see [Using instance profiles](#) in the *IAM User Guide*.

Enable remote access for a worker machine (optional)

To enable your worker machine to run remote commands for your application servers, you must ensure that the worker machine can connect.

For the required setup to enable remote access, choose the operating system tab that matches your application server.

Linux

For Linux application servers, you can use SSH key-based or SSH certificate-based connections. You must ensure that there is network connectivity between the worker machine and the application server, and verify that your worker machine can connect.

Certificate-based connections

By default, App2Container trusts the certificate, and does not verify its validity before connecting to your application server. To change this behavior, set the `acceptCerts` attribute to `false` in the `init.json` file.

Windows

To connect to a Windows application server from a Windows Server 2016, 2019, or 2022 worker machine, use the WinRM protocol. Your application server must meet the requirements that are listed for Windows in the [Supported applications](#) section of this user guide.

Note

App2Container does not support applications running on Windows client operating systems, such as Windows 7 or Windows 10.

1. Worker machine

To ensure that you can run PowerShell scripts on the worker machine, set the PowerShell Execution Policy to one of the following values:

RemoteSigned

Example:

```
PS> Set-ExecutionPolicy RemoteSigned
```

Unrestricted

Example:

```
PS> Set-ExecutionPolicy Unrestricted
```

2. Application servers

Complete the following steps on each application server to enable remote access from the worker machine.

1. Ensure network connectivity to the application server over WinRM port 5986.
2. Download the WinRMSetup.ps1 PowerShell script to your application server from the following location: [WinRMSetup.ps1](#).

Note

Checksum files for this script can be downloaded using the following links:

- [WinRMSetup.ps1.sha256](#)
- [WinRMSetup.ps1.md5](#)

3. Download the [New-SelfsignedCertificateEx.ps1](#) PowerShell script from the Microsoft Technet gallery. The WinRMSetup.ps1 PowerShell script from step 2 uses it to generate a self-signed certificate.

 Note

This script must run from the same directory where the WinRMSetup.ps1 PowerShell script from step 2 is located.

4. Run the WinRMSetup.ps1 PowerShell script on the application server. The script ensures that WinRM is enabled, and generates self-signed certificates that are used to secure the connection from the worker machine.

Configure your AWS profile

AWS App2Container requires command line access to AWS resources for containerization and deployment commands. It uses information from your AWS profile to configure access to AWS resources for your account. To run App2Container commands, you must install and configure a command line tool on the application servers and worker machines where you run the commands.

 Note

- AWS Tools for Windows PowerShell is required for running App2Container commands in PowerShell on a Windows server.
- Tools for Windows PowerShell comes pre-installed on Windows-based Amazon Machine Images (AMIs). If your application server or worker machine is an Amazon EC2 instance that was launched from one of these AMIs, you can skip to configuring your AWS profile. See [Shared credentials](#) in the *AWS Tools for PowerShell User Guide* for more details.

To install the AWS Command Line Interface (AWS CLI) or AWS Tools for Windows PowerShell command line tools, and to configure your AWS profile, follow the instructions on the tab that matches your command line tool.

AWS CLI

To install the AWS CLI and set up your AWS profile, follow these steps:

1. Install the AWS CLI according to the instructions in the *AWS Command Line Interface User Guide*. For more information, see [Installing the AWS CLI](#).
2. To configure your AWS default profile, use the **aws configure** command. For more information, see [Configuration basics](#) in the *AWS Command Line Interface User Guide*.

Tools for Windows PowerShell

To install Tools for Windows PowerShell and set up your AWS profile, follow these steps:

1. Install the Tools for Windows PowerShell according to the instructions in the *AWS Tools for PowerShell User Guide*. For more information see [Installing the AWS Tools for Windows PowerShell](#).
2. To set up your AWS default profile, use the [Initialize-AWSDefaultConfiguration](#) cmdlet. For more information about shared credentials in Tools for Windows PowerShell, see [Shared credentials](#) in the *AWS Tools for PowerShell User Guide*.

After you containerize your applications, you can also use the AWS CLI or Tools for Windows PowerShell to deploy them on AWS, though we recommend using the `--deploy` option with the **generate app-deployment** and **generate pipeline** commands to do your deployment.

Install the Docker engine

App2Container uses the Docker engine (Docker CE) to create container images and generate Dockerfiles that run the containers hosted on Amazon ECS, Amazon EKS, or AWS App Runner. You must install the Docker engine on the application server or worker machine that you'll use to containerize the application using the **containerize** command.

Linux

Use the following procedure to install Docker on Linux.

To install the Docker engine

1. **Install Docker version 17.07 or later**

Choose your Linux distribution from the following options, and follow instructions to download and install the Docker engine, using the links provided.

Amazon Linux

To download and install the Docker engine on Amazon Linux instances, see [Docker basics for Amazon ECS](#) in the *Amazon Elastic Container Service Developer Guide*. This works with any Amazon Linux instance.

RHEL

Recent versions of RHEL do not natively support the Docker engine. However, you can still download and install the Docker engine on RHEL to create containers that will be hosted and run on Amazon ECS, Amazon EKS, or AWS App Runner. To do this, follow the instructions given for CentOS on the Docker website: [Install Docker engine](#).

All other supported distributions (CentOS, Ubuntu)

To download and install the Docker engine for other supported Linux distributions, follow the instructions for your Linux distribution on the Docker website: [Install Docker engine](#).

2. Verify the Docker installation

To verify that your Docker installation was successful, run the following command.

```
$ docker run -it hello-world
```

When the command runs, it pulls the latest hello-world application from the Docker repository, if applicable. When the application has finished downloading, it displays a "Hello" message followed by information on how this command verified your installation of Docker.

Windows

Use the following procedure to install Docker on Windows.

To install the Docker engine

1. Install Docker version 17.07 or later

To download and install the Docker engine on Windows, see [Get started: Prep Windows for containers \(Install Docker section\)](#).

2. Verify the Docker installation

To verify that your Docker installation was successful, run the following command.

```
PS> docker run -it hello-world
```

When the command runs, it pulls the latest hello-world application from the Docker repository, if applicable. When the application has finished downloading, it displays a "Hello" message followed by information on how this command verified your installation of Docker.

Step 1: Install App2Container

To get started with App2Container, the first step is to download and install the application. To help ensure a successful installation, you can verify the integrity and authenticity of the binary file before installing it.

Tip

For Amazon EC2 instances, you can perform Step 1, Step 2, Step 3, and Step 4 by using an AWS Systems Manager Automation runbook. For more information, see [App2Container Automation runbook](#).

If you prefer, you can replatform your applications running on Amazon EC2 to containers and deploy them to Amazon ECS on AWS Fargate with a console-based experience by using the *Replatform applications to Amazon ECS* template in the [Migration Hub Orchestrator console](#). For more information, see the [AWS Migration Hub Orchestrator User Guide](#).

Choose the tab that matches your operating system (OS) platform to continue:

Linux

App2Container for Linux is packaged as a tar.gz archive. The archive contains an interactive shell script that installs App2Container on your server. If you use an application server and a worker machine, you must install App2Container on both.

To download and install App2Container for Linux

1. Download the installation file in one of the following ways:

- Use the **curl** command to download the App2Container installation package from Amazon S3.

```
$ curl -o AWSApp2Container-installer-linux.tar.gz https://app2container-release-us-east-1.s3.us-east-1.amazonaws.com/latest/linux/AWSApp2Container-installer-linux.tar.gz
```

- Use your browser to download the installer from the following URL: <https://app2container-release-us-east-1.s3.us-east-1.amazonaws.com/latest/linux/AWSApp2Container-installer-linux.tar.gz>.

2. Extract the package to a local folder on the server.

```
$ sudo tar xvf AWSApp2Container-installer-linux.tar.gz
```

3. Run the install script that you extracted from the package and follow the prompts.

```
$ sudo ./install.sh
```

To check the downloaded tar.gz installer archive for integrity, you can validate the SHA256 hash of the local file against the published hash file.

Verify the integrity of the download

1. **Generate hashes to verify**

From the directory where you downloaded your tar.gz installer, run the following command to generate the hash of the downloaded tar.gz file.

```
$ sha256sum AWSApp2Container-installer-linux.tar.gz
9482952019adb6df96c7be773aa20ecb8de559083b99c270c67c34da56dd8dee
AWSApp2Container-installer-linux.tar.gz
```

2. **Verify hashes against the public file**

Download the App2Container hash file from Amazon S3 with the following link, and compare the contents to the hash that you generated in step 1:

- Download the App2Container hash file from Amazon S3: [AWSApp2Container-installer-linux.tar.gz.sha256](https://app2container-release-us-east-1.s3.us-east-1.amazonaws.com/latest/linux/AWSApp2Container-installer-linux.tar.gz.sha256).

To verify the authenticity of the download, run the following commands to download the certificate and signature files, and verify the signature.

Verify the authenticity of the download

1. Download the App2Container certificate:

```
curl -o app2container.cert https://app2container-keys.s3.us-east-1.amazonaws.com/latest/app2container.cert
```

2. Download the App2Container signature file:

```
curl -o app2container.sig https://app2container-release-us-east-1.s3.us-east-1.amazonaws.com/latest/linux/app2container.sig
```

3. Verify the signature:

```
openssl dgst -sha256 -verify app2container.cert -signature app2container.sig /usr/bin/app2container
```

Windows

App2Container for Windows is packaged as a zip archive. The package contains a PowerShell script that installs App2Container. If you use an application server and a worker machine, you must install App2Container on both.

To download and install App2Container for Windows

1. Download the App2Container installation package, [AWSApp2Container-installer-windows.zip](#).
2. Extract the package to a local folder on the server and navigate to that folder.

Note

App2Container automatically enables NTFS long paths for all supported Windows versions so that you can use file paths longer than 260 characters. For more information about this setting, see [How to enable NTFS Long Paths in Windows 10 / Windows Server 2016 / 2019 or newer](#) (PDF).

3. Run the install script from the folder where you extracted it, and follow the prompts.

```
PS> .\install.ps1
```

4. (Optional) To verify the authenticity of the download, use the Get-AuthenticodeSignature PowerShell command as follows to get the Authenticode Signature of the App2Container executable.

```
PS> Get-AuthenticodeSignature C:\Users\Administrator\app2container  
AWSApp2Container\bin\app2container.exe
```

To check the downloaded zip archive for integrity, you can validate the SHA256 hash of the local file against the published hash file.

To verify the integrity of the download

1. **Generate hashes to verify**

From the directory where you downloaded your zip archive, run the following command to generate the hash of the downloaded archive file.

```
PS> Get-FileHash C:\Users\Administrator\Downloads\AWSApp2Container-installer-  
windows.zip -Algorithm SHA256
```

2. **Verify hashes against the public file**

Download the App2Container hash file from Amazon S3 with the following link, and compare the contents to the hash that you generated in step 1:

- Download the App2Container hash file from Amazon S3: [AWSApp2Container-installer-windows.zip.sha256](#).

Step 2: Initialize App2Container

The containerization process consists of several distinct phases. This step focuses on the initialization phase, during which you initialize App2Container's global settings, and configure remote command settings if you are using a worker machine.

The **init** command performs one-time initialization tasks for App2Container. This interactive command prompts for the information required to set up the local App2Container environment. Run this command before you run any other App2Container commands. For more information, see the [init](#) command reference page.

If you are using a worker machine to run commands remotely on application servers, you must also run the **remote configure** command on the worker machine. For more information, see the [remote configure](#) command reference page.

Choose the tab that matches your operating system (OS) platform to continue:

Linux

On each server where you installed App2Container, run the [init](#) command as follows.

```
$ sudo app2container init
```

You are prompted to provide the following information. Choose *<enter>* to accept the default value.

- Workspace directory path – A local directory where App2Container can store artifacts during the containerization process. The default is `/root/app2container`.
- AWS profile – Contains information needed to run App2Container, such as your AWS access keys. For more information about AWS profiles, see [Configure your AWS profile](#).

Note

If App2Container detects an instance profile for your server, the **init** command prompts if you want to use it. If you don't specify any value, App2Container uses your AWS default profile.

- Amazon S3 bucket – You can optionally provide the name of an Amazon S3 bucket where you can extract artifacts using the **extract** command. The **containerize** command uses the extracted components to create the application container if the Amazon S3 bucket is configured. The default is no bucket.
- You can optionally upload logs and command-generated artifacts automatically to App2Container support when an `app2container` command crashes or encounters internal errors.

- Permission to collect usage metrics – You can optionally allow App2Container to collect information about the host operating system, application type, and the **app2container** commands that you run. The default is to allow the collection of metrics.
- Whether to enforce signed images – You can optionally require that images are signed using Docker Content Trust (DCT). The default is no.

Windows

On each server where you installed App2Container, run the [init](#) command as follows.

```
PS> app2container init
```

You are prompted to provide the following information. Choose *<enter>* to accept the default value.

- Workspace directory path – A local directory where App2Container can store artifacts during the containerization process. The default is C:\Users\Administrator\AppData\Local\app2container.
- AWS profile – Contains information needed to run App2Container, such as your AWS access keys. For more information about AWS profiles, see [Configure your AWS profile](#).

Note

If App2Container detects an instance profile for your server, the **init** command prompts if you want to use it. If you don't specify any value, App2Container uses your AWS default profile.

- Amazon S3 bucket – You can optionally provide the name of an Amazon S3 bucket where you can extract artifacts using the **extract** command. The **containerize** command uses the extracted components to create the application container if the Amazon S3 bucket is configured. The default is no bucket.
- You can optionally upload logs and command-generated artifacts automatically to App2Container support when an app2container command crashes or encounters internal errors.
- Permission to collect usage metrics – You can optionally allow App2Container to collect information about the host operating system, application type, and the **app2container** commands that you run. The default is to allow the collection of metrics.

- Whether to enforce signed images – You can optionally require that images are signed using Docker Content Trust (DCT). The default is no.

Step 3: Analyze your application

After you have completed setup and initialization tasks on your servers, you can begin to analyze your applications. During the analysis phase, you take inventory of the applications running on your application servers, and analyze specific applications within your inventory.

Choose the tab that matches your operating system (OS) platform to continue:

Linux

On the application server, follow these steps to prepare to containerize the applications.

Prepare for containerization

1. Run the [inventory](#) command as follows to list the Java applications that are running on your server.

```
$ sudo app2container inventory
```

The output includes a JSON object collection with one entry for each application. Each application object will include key/value pairs as shown in the following example.

```
"java-app-id": {  
  "processId": pid,  
  "cmdline": "/user/bin/java ...",  
  "applicationType": "java-apptype"  
}
```

2. Locate the application ID for the application to containerize in the JSON output of the **inventory** command, and then run the [analyze](#) command as follows, replacing *java-app-id* with the application ID that you located.

```
$ sudo app2container analyze --application-id java-app-id
```

The output is a JSON file, `analysis.json`, stored in the workspace directory that you specified when you ran the **init** command.

3. (Optional) You can edit the information in the `containerParameters` section of `analysis.json` as needed before continuing to the next step.

Windows

On the application server, follow these steps to prepare to containerize your applications.

Prepare for containerization

1. Run the [inventory](#) command as follows to list the ASP.NET applications that are running on your server.

```
PS> app2container inventory
```

The output includes a JSON object collection with one entry for each application. Each application object will include key/value pairs as shown in the following example.

```
"iis-app-id": {  
  "siteName": My site name,  
  "bindings": "http/*:80:",  
  "applicationType": "iis",  
  "discoveredWebApps": [  
    "app1",  
    "app2"  
  ]  
}
```

2. Locate the application ID for the application to containerize in the JSON output of the **inventory** command, and then run the [analyze](#) command as follows, replacing *iis-app-id* with the application ID that you located.

```
PS> app2container analyze --application-id iis-app-id
```

The output is a JSON file, `analysis.json`, stored in the workspace directory that you specified when you ran the **init** command.

3. (Optional) You can edit the information in the `containerParameters` section of `analysis.json` as needed before continuing to the next step.

Step 4: Transform your application

Now that your application has gone through the analysis phase, it's ready for containerization. The transform phase creates the containers that your application runs in after you deploy it to Amazon ECS, Amazon EKS, or App Runner, if eligible. For more information about how App2Container integrates with container management services and other products, see [Product and service integrations for AWS App2Container](#).

Choose the tab that matches your operating system (OS) platform to continue:

Linux

The transform phase depends on whether you are running all steps on the application server, or are using the application server for the analysis and a worker machine for containerization and deployment.

To containerize the application on the application server

If you are using an application server for all steps, run the [containerize](#) command as follows.

```
$ sudo app2container containerize --application-id java-app-id
```

The output is a set of deployment files that are stored in the workspace directory that you specified when you ran the **init** command.

To containerize the application on a worker machine

If you are using a worker machine for containerization and deployment, use the following procedure to transform the application.

1. On the application server, run the [extract](#) command as follows.

```
$ sudo app2container extract --application-id java-app-id
```

2. If you specified an Amazon S3 bucket when you ran the **init** command, the archive is extracted to that location. Otherwise, you can manually copy the resulting archive file to the worker machine.
3. On the worker machine, run the [containerize](#) command as follows.

```
$ sudo app2container containerize --input-archive /path/extraction-file.tar
```

The output is a set of deployment artifacts that are stored in the workspace directory that you specified when you ran the **init** command.

Windows

The transform phase depends on whether you are running all steps on the application server or using the application server for the analysis and a worker machine for containerization and deployment.

To containerize the application on the application server

If you are using an application server for all steps, run the [containerize](#) command as follows.

```
PS> app2container containerize --application-id iis-app-id
```

The output is a set of deployment files stored in the workspace directory that you specified when you ran the **init** command.

To containerize the application on a worker machine

If you are using a worker machine for containerization and deployment, use the following procedure to transform the application.

1. On the application server, run the [extract](#) command as follows.

```
PS> app2container extract --application-id iis-app-id
```

2. If you specified an Amazon S3 bucket when you ran the **init** command, the archive is extracted to that location. Otherwise, you can manually copy the resulting archive file to the worker machine.
3. On the worker machine, run the [containerize](#) command as follows.

```
PS> app2container containerize --input-archive drive:\path\extraction-file.zip
```

The output is a set of deployment artifacts that are stored in the workspace directory that you specified when you ran the **init** command.

Step 5: Deploy your application

After your application has gone through containerization, it's ready to deploy to Amazon ECS, Amazon EKS, or App Runner, if eligible. When you run the **generate app-deployment** command, App2Container creates an Amazon ECR repository where it stores your application container artifacts for deployment. It also creates deployment configuration files that you can deploy as follows:

- You can customize the deployment files, and have complete control over the deployment by running the AWS commands for your destination container management environment. When you run the **generate app-deployment** command without the `--deploy` option, App2Container returns instructions that you can use to deploy manually.
- If you're sure that you won't need to customize your deployment files, App2Container can optionally deploy your application containers directly to the container management environment that you have configured. To choose this option, run the **generate app-deployment** command with the `--deploy` option. You can verify the settings that App2Container used for the deployment by reviewing the deployment configuration files.

The deployment phase includes the option to create a deployment pipeline using the **generate pipeline** command. That step is not covered here, in order to prevent any unexpected charges for AWS resources. For more information, see [app2container generate pipeline command](#) in the command reference section.

Choose the tab that matches your operating system (OS) platform to continue:

Linux

Run the [generate app-deployment](#) command as follows to deploy the application on AWS.

```
$ sudo app2container generate app-deployment --application-id java-app-id
```

You have now created deployment artifacts for your application! You can find the deployment artifacts that the **generate app-deployment** command created for you in the local directory for your application.

Windows

Run the [generate app-deployment](#) command as follows to deploy the application on AWS.

```
PS> app2container generate app-deployment --application-id iis-smarts-51d2dbf8
```

You have now created deployment artifacts for your application! You can find the deployment artifacts that the **generate app-deployment** command created for you in the local directory for your application.

Applications using Windows authentication

For applications using Windows authentication, you can use the `gMSAParameters` inside of the `deployment.json` file to set the gMSA-related artifacts automatically during generation of your AWS CloudFormation template.

Perform the actions in the list below once per Active Directory domain before you update the gMSA parameters.

- Set up a secret in SecretsManager that stores the Domain credentials with the following key value pairs:

Key	Value
Username	<DomainNetBIOSName>\<Domain User>
Password	<DomainUserPassword>

- For the VPC with the Domain Controller, verify that the DHCP options are set to reach the Domain Controller. The options for `DomainName` and `DomainNameServers` must be set correctly. See [DHCP options sets](#) for more information about how to set DHCP options.

Step 6: Clean up

If you explored deployment options outside of the steps that we covered for this tutorial, make sure that you tear down any application stacks that might have been created, and verify that you have removed any artifacts that were created in the process.

Choose the tab that matches your operating system (OS) platform to continue:

Linux

To remove App2Container from your application server or worker machine, delete the `/usr/local/app2container` folder where it is installed, and then remove this folder from your path.

To clean up your AWS profile, use the **aws configure set** command. For more information, see [Set and view configuration settings](#) in the *AWS Command Line Interface User Guide*.

Windows

To remove App2Container from your application server or worker machine, delete the `C:\Users\Administrator\app2container` folder where it is installed, and then remove this folder from your path.

To clean up your AWS profile, see [Removing Credential Profiles](#) in the *AWS Tools for PowerShell User Guide*.

App2Container Automation runbook

AWS App2Container provides the `AWSApp2Container-ReplatformApplications` Automation runbook for use on Amazon EC2 instances. Automation is a capability of AWS Systems Manager. The runbook performs the installation of App2Container as well as the initialize, analyze, and transform phases for replatforming supported applications. If desired, the automation can also push the containerized application to Amazon Elastic Container Registry (Amazon ECR). For more information, see [App2Container compatibility](#) and [Applications you can containerize using AWS App2Container](#).

You must have access to Systems Manager to use the runbook. For more information about Systems Manager Automation, see [AWS Systems Manager Automation](#) in the *AWS Systems Manager User Guide*.

Tip

To containerize your applications with a console-based experience and deploy them on Amazon ECS on AWS Fargate, you can use the *Replatform applications to Amazon ECS* template on the [AWS Migration Hub Orchestrator console](#). For more information, see [Replatform applications to Amazon ECS](#) in the *AWS Migration Hub Orchestrator User Guide*.

Contents

- [Prerequisites](#)
 - [Create policies and roles for the automation](#)
 - [Attaching the IAM role](#)
- [Run the automation](#)
 - [Runbook parameters](#)
 - [Running the automation](#)
 - [Reviewing output from the automation](#)
- [Complete the modernization process](#)

Prerequisites

Before you run the automation, you must have:

- An S3 bucket to store your containerized application artifacts. This bucket must be in the same AWS account and Region as your Amazon EC2 instances being containerized. For more information, see [Creating a bucket](#) in the *Amazon Simple Storage Service User Guide*.
- An IAM service role with the permissions necessary for Automation, a capability of AWS Systems Manager, to run the automation on your behalf.
- An IAM role for your EC2 instances that permits the necessary actions to run the automation in your target instances.
- (Optional) A customer managed key in AWS KMS to use as your own server-side encryption key for Amazon S3. For more information, see [Customer managed keys](#) in the *Amazon Simple Storage Service User Guide*.
- If you are using AWS Application Migration Service and running this automation as a post-launch action, you must configure the EC2 launch template setting **Auto-assign public IP** to Enabled. For more information, see [Full launch template setting review](#) in the *AWS Application Migration Service User Guide*.

Topics

- [Create policies and roles for the automation](#)
- [Attaching the IAM role](#)

Create policies and roles for the automation

You must create the required policies and roles before running the automation. You can create the roles using AWS CloudFormation or manually.

Creating policies and roles with AWS CloudFormation

You can use the following AWS CloudFormation template to create a stack which will create the roles and policies required to run the automation. You can create a stack using the [AWS CloudFormation console](#) or the [AWS Command Line Interface \(AWS CLI\)](#).

```
AWSTemplateFormatVersion: "2010-09-09"
Parameters:
  A2CServiceRoleName:
    Type: String
    Description: Name of the A2C Service Role
    Default: "a2cServiceRole"
```

A2CInstanceRoleName:

Type: String

Description: Name of the A2C Instance Role

Default: "a2cinstancerole"

Resources:

A2CServiceRole:

Type: "AWS::IAM::Role"

Properties:

RoleName: !Ref A2CServiceRoleName

AssumeRolePolicyDocument:

Version: "2012-10-17"

Statement:

- Effect: "Allow"

Principal:

Service: ["ssm.amazonaws.com"]

Action: "sts:AssumeRole"

Policies:

- PolicyName: "a2cServicePolicy"

PolicyDocument:

Version: "2012-10-17"

Statement:

- Sid: "EC2DescribeAccess"

Effect: "Allow"

Action:

- "ec2:DescribeInstances"

Resource: "*"

- Sid: "IAMRoleAccess"

Effect: "Allow"

Action:

- "iam:AttachRolePolicy"

- "iam:GetInstanceProfile"

Resource: "*"

- Sid: "ApplicationTransformationAccess"

Effect: "Allow"

Action:

- "application-transformation:StartRuntimeAssessment"

- "application-transformation:GetRuntimeAssessment"

- "application-transformation:PutMetricData"

- "application-transformation:PutLogData"

Resource: "*"

- Sid: "SSMSendCommandAccess"

Effect: "Allow"

Action:

```

        - "ssm:SendCommand"
    Resource:
        - "arn:aws:ec2:*:*:instance/*"
        - "arn:aws:ssm:*:*:document/AWS-RunRemoteScript"
    - Sid: "SSMDescribeAccess"
      Effect: "Allow"
      Action:
        - "ssm:DescribeInstanceInformation"
        - "ssm:ListCommandInvocations"
        - "ssm:GetCommandInvocation"
        - "ssm:GetParameters"
      Resource: "arn:aws:ssm:*:*:*"
    - Sid: "S3ObjectAccess"
      Effect: "Allow"
      Action:
        - "s3:GetObject"
        - "s3:PutObject"
      Resource:
        - "arn:aws:s3::*/application-transformation*"
    - Sid: "S3ListAccess"
      Effect: "Allow"
      Action:
        - "s3:ListBucket"
        - "s3:GetBucketLocation"
      Resource: "arn:aws:s3:::*"
    - Sid: "KmsAccess"
      Effect: "Allow"
      Action:
        - "kms:GenerateDataKey"
        - "kms:Decrypt"
      Resource:
        - "arn:aws:kms:*:*:key/*"
      Condition:
        StringLike:
          kms:ViaService:
            - "s3.*.amazonaws.com"

```

A2CInstanceRole:

Type: "AWS::IAM::Role"

Properties:

RoleName: !Ref A2CInstanceRoleName

AssumeRolePolicyDocument:

Version: "2012-10-17"

Statement:

```

    - Effect: "Allow"
      Principal:
        Service: ["ec2.amazonaws.com"]
      Action: "sts:AssumeRole"
ManagedPolicyArns:
- "arn:aws:iam::aws:policy/AmazonSSMManagedInstanceCore"
Policies:
- PolicyName: "ApplicationTransformationAnalyzerPolicy"
  PolicyDocument:
    Version: "2012-10-17"
    Statement:
      - Sid: "S3BucketAccess"
        Effect: "Allow"
        Action:
          - "s3:GetBucketLocation"
        Resource:
          - "arn:aws:s3:::*"
      - Sid: "S3ObjectAccess"
        Effect: "Allow"
        Action:
          - "s3:PutObject"
          - "s3:GetObject"
        Resource:
          - "arn:aws:s3:::*/application-transformation*"
      - Sid: "KmsAccess"
        Effect: "Allow"
        Action:
          - "kms:GenerateDataKey"
          - "kms:Decrypt"
        Resource:
          - "arn:aws:kms:*:*:key/*"
        Condition:
          StringLike:
            kms:ViaService:
              - "s3.*.amazonaws.com"
      - Sid: "TelemetryAccess"
        Effect: "Allow"
        Action:
          - "application-transformation:PutMetricData"
          - "application-transformation:PutLogData"
        Resource:
          - "*"
a2cInstanceProfile:
  Type: AWS::IAM::InstanceProfile

```

Properties:

InstanceProfileName: !Ref A2CInstanceRoleName

Roles:

- !Ref A2CInstanceRole

Creating policies and roles manually

The following sections detail how you can manually create the roles and policies required to run the automation.

Creating policies to run the automation

To enhance the security posture of the App2Container automation execution, it is strongly recommended to scope down IAM S3 access permissions to allow access only to the bucket created for the App2Container automation execution. You can create least-privilege policies required to run the automation with the following procedures.

To create the service role policy for running the automation

1. Open the IAM console at <https://console.aws.amazon.com/iam/>.
2. In the navigation pane, choose **Policies** then choose **Create policy**.
3. Choose **JSON**, enter the following policy in the **Policy editor**, then choose **Next**:

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "EC2DescribeAccess",
      "Effect": "Allow",
      "Action": [
        "ec2:DescribeInstances"
      ],
      "Resource": "*"
    },
    {
      "Sid": "IAMRoleAccess",
      "Effect": "Allow",
      "Action": [
        "iam:AttachRolePolicy",
```

```

        "iam:GetInstanceProfile"
    ],
    "Resource": ["*"]
},
{
    "Sid": "ApplicationTransformationAccess",
    "Effect": "Allow",
    "Action": [
        "application-transformation:StartRuntimeAssessment",
        "application-transformation:GetRuntimeAssessment",
        "application-transformation:PutMetricData",
        "application-transformation:PutLogData"
    ],
    "Resource": "*"
},
{
    "Sid": "SSMSendCommandAccess",
    "Effect": "Allow",
    "Action": [
        "ssm:SendCommand"
    ],
    "Resource": [
        "arn:aws:ec2:*:*:instance/*",
        "arn:aws:ssm:*:*:document/AWS-RunRemoteScript"
    ]
},
{
    "Sid": "SSMDescribeAccess",
    "Effect": "Allow",
    "Action": [
        "ssm:DescribeInstanceInformation",
        "ssm:ListCommandInvocations",
        "ssm:GetCommandInvocation",
        "ssm:GetParameters"
    ],
    "Resource": "arn:aws:ssm:*:*:*"
},
{
    "Sid": "S3ObjectAccess",
    "Effect": "Allow",
    "Action": [
        "s3:GetObject",
        "s3:PutObject"
    ],

```

```

        "Resource": [
            "arn:aws:s3::*/application-transformation*"
        ]
    },
    {
        "Sid": "S3ListAccess",
        "Effect": "Allow",
        "Action": [
            "s3:ListBucket",
            "s3:GetBucketLocation"
        ],
        "Resource": "arn:aws:s3:::*"
    },
    {
        "Sid": "KmsAccess",
        "Effect": "Allow",
        "Action": [
            "kms:GenerateDataKey",
            "kms:Decrypt"
        ],
        "Resource": [
            "arn:aws:kms:*:key/*"
        ],
        "Condition": {
            "StringLike": {
                "kms:ViaService": [
                    "s3.*.amazonaws.com"
                ]
            }
        }
    }
}

```

4. Enter a value for the **Policy name**.
5. Choose **Create policy**.

To create the policy for the IAM role used by your instance profile

1. Open the IAM console at <https://console.aws.amazon.com/iam/>.
2. In the navigation pane, choose **Policies** then choose **Create policy**.
3. Choose **JSON**, enter the following policy in the **Policy editor**, then choose **Next**:

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "S3BucketAccess",
      "Effect": "Allow",
      "Action": [
        "s3:GetBucketLocation"
      ],
      "Resource": [
        "arn:aws:s3:::*"
      ]
    },
    {
      "Sid": "S3ObjectAccess",
      "Effect": "Allow",
      "Action": [
        "s3:PutObject",
        "s3:GetObject"
      ],
      "Resource": [
        "arn:aws:s3:::*/application-transformation*"
      ]
    },
    {
      "Sid": "KmsAccess",
      "Effect": "Allow",
      "Action": [
        "kms:GenerateDataKey",
        "kms:Decrypt"
      ],
      "Resource": [
        "arn:aws:kms:*:key/*"
      ],
      "Condition": {
        "StringLike": {
          "kms:ViaService": [
            "s3.*.amazonaws.com"
          ]
        }
      }
    }
  ]
}
```

```

        }
    },
    {
        "Sid": "TelemetryAccess",
        "Effect": "Allow",
        "Action": [
            "application-transformation:PutMetricData",
            "application-transformation:PutLogData"
        ],
        "Resource": [
            "*"
        ]
    }
]
}

```

4. Enter **ApplicationTransformationAnalyzerPolicy** for the **Policy name**.
5. Choose **Create policy**.

Creating the IAM service role for running the automation

You can use the following procedure to create an IAM service role.

To create an IAM role using the IAM console

1. Open the IAM console at <https://console.aws.amazon.com/iam/>.
2. In the navigation pane, choose **Roles** then choose **Create role**.
3. On the **Select trusted entity** page, choose **AWS service**, select the **Systems Manager** use case, and then choose **Next**.
4. On the **Add permissions** page, select the policy that you created for the IAM service role previously, and then choose **Next**.
5. On the **Name, review, and create** page, enter a name and description for the role and add tags if needed.
6. Choose **Create role**.

This role is used for the `AutomationAssumeRole` parameter in the [Run the automation](#) section.

Creating the instance profile role

You can use the following procedure to create an IAM role for your instance profile. The permissions provided by the instance profile role are used by your EC2 instances. For more information, see [Using an IAM role to grant permissions to applications running on Amazon EC2 instances](#) in the *AWS Identity and Access Management User Guide*.

Note

An instance profile can only contain one IAM role. If your target instances have an existing IAM role, the automation will add the `ApplicationTransformationAnalyzerPolicy` policy on execution to the instance profile role on your behalf. The existing role should provide the permissions required to make the instances managed nodes in AWS Systems Manager. For more information, see [Instance profiles](#) in the *Amazon Elastic Compute Cloud User Guide* and [Managed nodes](#) in the *AWS Systems Manager User Guide*.

To create an instance profile role using the IAM console

1. Open the IAM console at <https://console.aws.amazon.com/iam/>.
2. In the navigation pane, choose **Roles** then choose **Create role**.
3. On the **Select trusted entity** page, choose **AWS service**, select the **EC2** use case, and then choose **Next**.
4. On the **Add permissions** page, select both the `AmazonSSMManagedInstanceCore` policy and the policy you created for the instance profile role previously, and then choose **Next**.
5. On the **Name, review, and create** page, enter a name and description for the role and add tags if needed.
6. Choose **Create role**.

The instance profile role is used in the following section.

Attaching the IAM role

If your target instances don't have an existing IAM role, you can attach the previously created IAM role to them. The following steps assume you have already created the required policies and roles.

To attach an IAM role to an instance

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. In the navigation pane, choose **Instances**.
3. Select the instance, choose **Actions, Security, Modify IAM role**.
4. Select the IAM role to attach to your instance, and choose **Save**.

For more information, see [Attach an IAM role to an instance](#).

Run the automation

When you run the automation, the following processes occur:

- Discover – The instances you specified are scanned for supported applications to create an inventory of each server.
- Analyze – Once the discover phase has completed, the automation analyzes each application and creates an entry. The instances you specified are scanned for supported applications to create an inventory of each server. Once this discovery process has completed, the automation analyzes each application and creates an entry.

Note

Applications using Windows Server operating systems will use Windows Server Core as their base image. Applications using Linux operating systems will use a Linux based image.

Topics

- [Runbook parameters](#)
- [Running the automation](#)
- [Reviewing output from the automation](#)

Runbook parameters

You can specify the following parameters for the Automation runbook.

Parameter name	Type	Description	Default value	Required
AutomationAssumeRole	String	The ARN of the role that allows Automation to perform actions on your behalf.		TRUE
EnableContainerization	Boolean	Controls whether to containerize discovered applications. If enabled, the automation will use the artifacts uploaded to the S3 bucket to generate Open Containers Initiative (OCI) container images and push them to Amazon ECR.	FALSE	FALSE
OutputLocation	String	The S3 location in which to upload deployment artifacts. The bucket must be in the same account and		TRUE

Parameter name	Type	Description	Default value	Required
		Region of the EC2 instance. All artifacts will be created with a prefix of application-transformation .		
OutputEncryptionKey	String	The ARN of a customer managed KMS key to use for server-side encryption. For more information, see Protecting data with server-side encryption in the <i>Amazon Simple Storage Service User Guide</i> .		FALSE
InstanceId	String	An EC2 instance ID with applications to be assessed for replatforming. Only running applications are assessed.		TRUE

Running the automation

You can run the automation from the Systems Manager console.

To run the automation

1. Access the AWS Systems Manager Automation console at <https://console.aws.amazon.com/systems-manager/automation>.
2. Choose **Execute automation**.
3. Under **Automation runbook**, enter `AWSApp2Container-ReplatformApplications`, and search the repository.
4. Choose the `AWSApp2Container-ReplatformApplications` runbook, then choose **Next**.
5. Enter the required parameters, and any optional ones you require:
 - a. For `AutomationAssumeRole`, enter the ARN of the service role you created previously.
 - b. For `EnableContainerization`, specify `TRUE` if you want your containerized applications pushed to Amazon ECR.
 - c. For `OutputLocation`, specify the S3 path to upload artifacts to.
 - d. For `OutputEncryptionKey`, you can specify the ARN of a KMS key if you want to encrypt the uploaded objects with your customer managed key.
 - e. For `InstanceId`, specify the instance ID for the automation to take action on.
6. Choose **Execute**.

Reviewing output from the automation

Once the automation has completed, you can access the output in the S3 location that you provided.

To review output from the automation

1. Access the AWS Systems Manager Automation console at <https://console.aws.amazon.com/systems-manager/automation>.
2. Choose the **Execution ID** to review.
3. Select **Outputs** and review the **Finalize.report** output.
4. For more details, review the text file indicated in the **Finalize.reportS3Location** output.

Complete the modernization process

You can complete the modernization process using AWS Migration Hub Orchestrator to create a workflow based on the *Replatform applications to Amazon ECS* template to deploy your applications on Amazon ECS on AWS Fargate. This template can use the application artifacts App2Container uploaded to Amazon S3. For more information, see [Replatform applications to Amazon ECS](#) in the *AWS Migration Hub Orchestrator User Guide*.

To continue the containerization process without Migration Hub Orchestrator, you can use the App2Container CLI extraction and containerization process. For more information, see [Step 4: Transform your application](#).

After performing the containerization process with App2Container, continue with the deployment phase to complete the modernization process. You can use either App2Container or proprietary deployment tools. If you use the App2Container CLI, you can generate the required AWS CloudFormation templates. For more information about deploying your containerized application using App2Container, see [Step 5: Deploy your application](#).

Configuring your application

Containerizing your application and creating pipelines with App2Container requires configuration throughout the process. This section of the guide describes the configuration files that are created by **app2container** commands, the fields that they contain, and which fields are configurable. App2Container commands primarily generate JSON configuration files, using standard JSON notation. Field details for the files included here indicate where there are specific requirements for the values.

App2Container also generates YAML format CloudFormation templates when you run the **generate app-deployment** command. However, those are not covered in this section, as their content is dictated by the target container management environment, such as Amazon ECS, Amazon EKS, or AWS App Runner. For more information about how App2Container works with these services, see [Product and service integrations for AWS App2Container](#).

Creating IAM resources is also covered separately, under the Security section. For more information and instructions about how to set up IAM resources for App2Container, see [Identity and access management in App2Container](#).

You can consolidate your containerization workload by configuring connections to your application servers to run containerization workflows remotely, using App2Container remote commands from your worker machine. Prior to running remote commands, you must configure the connections that the worker machine uses for its target application servers. For more information on configuring connections, see the [remote configure](#) command reference page.

Contents

- [Manage secrets for AWS App2Container](#)
- [Configuring application containers](#)
- [Configuring container deployment](#)
- [Configuring container pipelines](#)

Manage secrets for AWS App2Container

App2Container uses AWS Secrets Manager to manage the credentials necessary to connect your worker machine to application servers and run remote commands. Secrets Manager encrypts your secrets for storage and provides an Amazon Resource Name (ARN) so that you can access

the secret. When you run the **remote configure** command, you provide the secret ARN that App2Container uses to connect to your target server when you run the remote command.

For more information about Secrets Manager, see [What Is AWS Secrets Manager?](#) For information specifically related to costs, see [Pricing for AWS Secrets Manager](#) in the *AWS Secrets Manager User Guide*.

Create remote access secrets

The secret that App2Container uses to connect to an application server varies with the application server's operating system (OS) platform. To create a remote access secret for your application server, choose the tab that matches your OS platform.

Linux

For Linux, you can store either the SSH private key or the Certificate and SSH private key in Secrets Manager. To create a secret in Secrets Manager so that you can access your application server remotely, follow the steps shown in the [Create a secret](#) page in the *AWS Secrets Manager User Guide*. Provide the information that App2Container needs to run remote commands as follows.

Step 1 Choose secret type

- **Secret type** – To store a key that App2Container uses programmatically, through API calls, choose the **Other type of secrets** option.
- Specify the following **Key/value pairs** to store in the secret. To add the next key/value pair, choose **+ Add row**.

Username key

- **Key name (box 1):** `username`
- **Key value (box 2):** Enter the plaintext username value to use with SSH.

SSH private key

- **Key name (box 1):** `key`
- **Key value (box 2):** Copy the base64-encoded string that represents your private key file into the second box.

Note

To base64-encode your key file, you can use the following command, where `.ssh/id_rsa` is the private key that encodes the file:

```
$ base64 .ssh/id_rsa
```

SSH Certificate key (optional)

- **Key name (box 1):** `cert`
- **Key value (box 2):** Copy the base64-encoded string that represents your signed certificate file into the second box.

Note

To base64-encode your signed certificate file, you can use the following command, where `.ssh/id_rsa-cert.pub` is the private key that encodes the file:

```
$ base64 .ssh/id_rsa-cert.pub
```

Step 2 Configure secret

- Enter a name for your secret in the **Secret name** box. You can also enter optional information to help identify your secret, such as **Description**, or you can enter tags in the **Tags** panel.

Windows

For Windows application servers, you can store the Username and Password for remote access. In most cases, the username and password translates to a set of credentials for a domain user with access to the application servers. [Create a secret](#) page in the *AWS Secrets Manager User Guide*

Step 1 Choose secret type

- **Secret type** – To store a key that App2Container uses programmatically, through API calls, choose the **Other type of secrets** option.

- Specify the following **Key/value pairs** to store in the secret. To add the next key/value pair, choose **+ Add row**.

Username key

- **Key name (box 1):** `username`
- **Key value (box 2):** In the second box, enter the plaintext username value to use with the connection credentials for your application server.

Password key

- **Key name (box 1):** `password`
- **Key value (box 2):** In the second box, enter the password value.

Step 2 Configure secret

- Enter a name for your secret in the **Secret name** box. You can also enter optional information to help identify your secret, such as **Description**, or you can enter tags in the **Tags** panel.

Create secrets for Jenkins pipelines

Integration with Jenkins requires secure authentication, both for the Git repository that Jenkins uses for automated container build pipelines, and for authentication to the Jenkins server itself. For secure authentication, App2Container uses Secrets Manager to store credentials, and provide access to the authentication secrets to Jenkins agent nodes.

Jenkins secrets

- [Authentication secret for Git](#)
- [Authentication secret for Jenkins server](#)

Authentication secret for Git

App2Container uses SSH to authenticate to the Git source repository that the Jenkins agent uses to update your pipeline. In the `pipeline.json` file, you provide the ARN from the authentication secret you create, in the `sshKeyArn` parameter value.

To create a secret in Secrets Manager so that App2Container can authenticate to the Git repository for the Jenkins agent, follow the steps shown in the [Create a secret](#) page in the *AWS Secrets*

Manager User Guide. Provide the information that App2Container needs to authenticate to the Git source repository as follows.

Step 1 Choose secret type

- **Secret type** – To store a key that App2Container uses programmatically, through API calls, choose the **Other type of secrets** option.
- Specify the following **Key/value pairs** to store in the secret. To add the next key/value pair, choose **+ Add row**.

Username key

- **Key name (box 1): username**
- **Key value (box 2):** In the second box, enter the plaintext username value that App2Container uses with SSH to authenticate to the Git source repository for Jenkins.

Username key

- **Key name (box 1): key**
- **Key value (box 2):** In the second box, copy the base64-encoded string that represents your private key file.

Note

To base64-encode your key file, you can use the following command, where `.ssh/id_rsa` is the private key that encodes the file:

```
$ base64 .ssh/id_rsa
```

Step 2 Configure secret

- Enter a name for your secret in the **Secret name** box. You can also enter optional information to help identify your secret, such as **Description**, or you can enter tags in the **Tags** panel.

Authentication secret for Jenkins server

Just as App2Container needs credentials to interact with AWS services on your behalf, so it also needs credentials to interact with the Jenkins server that runs your pipelines. In the

pipeline.json file, you provide the ARN from the authentication secret you create, in the apiTokenArn parameter value.

Generate a Jenkins authentication token

Before you store your Jenkins authentication secrets in Secrets Manager, generate an API token from your Jenkins server. To generate a Jenkins API authentication token, follow these steps:

1. Log in to your Jenkins server.
2. In the upper right corner of the interface, choose your name.
3. From the left side navigation menu, choose **Configure**.
4. In the **API Token** panel, choose **Add new Token**.
5. After Jenkins generates the token, give it a name. Keep track of the name. You will need it for the secret key you enter in Secrets Manager.
6. Choose the copy icon to copy the token value, or select and copy the value manually. You will need it for the secret value that you enter in Secrets Manager. You can't see the value again after you log out of Jenkins.

Note

Ensure that you revoke tokens that you no longer need.

Store your Jenkins authentication token in Secrets Manager

To create a secret in Secrets Manager for the Jenkins authentication token, follow the steps shown in the [Create a secret](#) page in the *AWS Secrets Manager User Guide*. Provide the information that App2Container needs to authenticate to the Jenkins server that runs your pipelines as follows.

Step 1 Choose secret type

- **Secret type** – To store a key that App2Container uses programmatically, through API calls, choose the **Other type of secrets** option.
- Specify the following **Key/value pairs** to store in the secret. To add the next key/value pair, choose **+ Add row**.

Username key

- **Key name (box 1):** username

- **Key value (box 2):** In the second box, enter the plaintext username value so that App2Container can log in to the Jenkins server.

Username key

- **Key name (box 1):** `apitoken`
- **Key value (box 2):** In the second box, copy the base64-encoded string that represents your Jenkins authentication token.

Note

To base64-encode a string, you can use the following command:

```
$ echo string-to-encode | base64
```

Step 2 Configure secret

- Enter a name for your secret in the **Secret name** box. You can also enter optional information to help identify your secret, such as **Description**, or you can enter tags in the **Tags** panel.

Create secrets for Microsoft Azure DevOps pipelines

To integrate with Azure Repos Git repositories and Azure DevOps pipelines, App2Container uses secure authentication. App2Container authenticates with a Microsoft Azure Personal Access Token (PAT) that you store as a secret in Secrets Manager.

In the `apiTokenArn` parameter value of the `pipeline.json` file, provide the ARN from the authentication secret that you create.

Generate a Microsoft Azure Personal Access Token (PAT)

Before you generate a Personal Access Token (PAT), you first must have an active Microsoft Azure account, with an organization and project already defined. For more information about how to set up Azure DevOps, see [Prerequisites](#).

To generate a PAT for your Microsoft Azure account, sign in to your Azure organization and create a new token with a **Custom defined** scope. For instructions, see [Create a PAT](#) in the *Azure DevOps*

Services documentation on the Microsoft documentation website. Choose the settings for your custom scope as follows.

- **Agent Pools:** Read and manage
- **Build:** Read and execute
- **Code:** Full
- **Extensions:** Read and manage
- **Release:** Read, write, execute, and manage
- **Service Connections:** Read and query

 Note

If you don't see all of the settings, choose **Show all scopes** to show the complete list.

Store your PAT in Secrets Manager

To create a secret in Secrets Manager for the PAT, follow the procedure on the [Create a secret](#) page in the *AWS Secrets Manager User Guide*. To access the Azure Repos Git repository, and Azure DevOps, provide the information that App2Container needs to authenticate to Microsoft Azure, as follows.

Step 1 Choose secret type

- **Secret type** – To store a key that App2Container uses programmatically, through API calls, choose the **Other type of secrets** option.
- Specify the following **Key/value pair** to store in the secret.

PAT key

- **Key name (box 1):** `azure-personal-access-token`
- **Key value (box 2):** Paste a copy of the token string that the Azure DevOps service generated.

Step 2 Configure secret

- Enter a name for your secret in the **Secret name** box. You can also enter optional information to help identify your secret, such as **Description**, or you can enter tags in the **Tags** panel.

Configuring application containers

When you run the [analyze](#) command, an `analysis.json` file is created for the application that is specified in the `--application-id` parameter. The **containerize** command uses this file to build the application container image and to generate artifacts.

You can configure the fields in the `containerParameters` section before running the **containerize** command to customize your application container. For configurable key/value pairs that do not apply to your container, set string values to an empty string, numeric values to zero, and Boolean values to false.

Containers running on Linux

For applications running on Linux, the application `analysis.json` file includes the following content:

Read-only data

- **Control fields** – Fields having to do with file creation, such as template version, and the file creation timestamp.
- **analysisInfo** – System dependencies for the application.

Configurable data

The `containerParameters` section contains the following fields:

- **imageRepository** (string, required) – The name of the repository where the application container image is stored.
- **imageTag** (string, required) – A tag for the build version of the application container image.
- **containerBaseImage** (string) – The base operating system (OS) image for the container build. By default, App2Container uses the operating system from the application server or worker machine where containerization runs.

Note

If specified, this must be an image name from your registry in the format `<image name>[:<tag>]`, and it must match the operating system platform and version that runs

on the application server or worker machine where containerization runs. The tag is optional if the repository supports "latest".

- **appExcludedFiles** (array of strings) – Specific files and directories to exclude from the container build.
- **appSpecificFiles** (array of strings) – Specific files and directories to include in the container build.
- **applicationPort** (number) – The application port exposed inside of the container. This port is tested for a successful HTTP response during pre-validation when the **containerize** command runs. App2Container assigns this as the default exposed port when creating a load balancer during deployment.
- **applicationMode** (Boolean, required) – The approach that App2Container uses to determine which files to include in your container image. App2Container uses application mode (value=true) for supported application frameworks, and process mode (value=false) for all other configurations.

You can override this value if necessary. For example, if your application is running on a supported framework, but App2Container did not recognize it and therefore assigned process mode, you can override the setting to use application mode instead.

Application mode settings

- *true* (application mode): For supported application frameworks, App2Container targets only the application files and dependencies that are needed for containerization, thereby minimizing the size of the resulting container image. This is known as *application mode*. Supported application frameworks include: Tomcat, TomEE, and JBoss (standalone mode).
- *false* (process mode): If App2Container does not find a supported framework running on your application server, or if you have other dependent processes running on your server, App2Container takes a conservative approach to identifying dependencies. This is known as *process mode*. For process mode, all non-system files on the application server are included in the container image.

Tip

If your application container image includes unnecessary files, or is missing files that should be included, use the following parameters to make corrections:

- To specify files to exclude from your application container image, use the `appExcludedFiles` parameter.

- To add files that were missed, use the `appSpecificFiles` parameter.

- **logLocations** (array of strings) – Specific log files or log directories to be routed to stdout. This enables applications that write to log files on the host to be integrated with AWS services such as CloudWatch and Firehose.
- **enableDynamicLogging** (Boolean, required) – Maps application logs to stdout as they are created. *If set to true, requires log directories to be entered in logLocations.*
- **dependencies** (array of strings) – A listing of all dependent processes or applications found for the application ID by the **analyze** command. You can remove specific dependencies to exclude them from the container.

Examples

The following examples show an `analysis.json` file for an application running on Linux. Choose the tab that matches your application.

Java

This example shows an `analysis.json` file for a Java application running on Linux.

```
{
  "a2CTemplateVersion": "",
  "createdTime": "",
  "containerParameters": {
    "_comment1": "*** EDITABLE: The below section can be edited according
to the application requirements. Please see the analysisInfo section below for
details discovered regarding the application. ***",
    "imageRepository": "java-tomcat-6e6f3a87",
    "imageTag": "latest",
    "containerBaseImage": "ubuntu:18.04",
    "appExcludedFiles": [],
    "appSpecificFiles": [],
    "applicationPort": 5000,
    "applicationMode": true,
    "logLocations": [],
    "enableDynamicLogging": false,
    "dependencies": []
  },
  "analysisInfo": {
    "_comment2": "*** NON-EDITABLE: Analysis Results ***",
```

```

    "processId": 2065,
    "appId": "java-tomcat-6e6f3a87",
    "userId": "1000",
    "groupId": "1000",
    "cmdline": [
        "/usr/bin/java",
        "... list of commands",
        "start"
    ],
    "osData": {
        "BUG_REPORT_URL": "",
        "HOME_URL": "",
        "ID": "ubuntu",
        "ID_LIKE": "debian",
        "NAME": "Ubuntu",
        "PRETTY_NAME": "Ubuntu 18.04.2 LTS",
        "PRIVACY_POLICY_URL": "",
        "SUPPORT_URL": "",
        "UBUNTU_CODENAME": "",
        "VERSION": "",
        "VERSION_CODENAME": "",
        "VERSION_ID": "18.04"
    },
    "osName": "ubuntu",
    "ports": [
        {
            "localPort": 8080,
            "protocol": "tcp6"
        },
        {
            "localPort": 8009,
            "protocol": "tcp6"
        },
        {
            "localPort": 8005,
            "protocol": "tcp6"
        }
    ],
    "Properties": {
        "catalina.base": "<application directory>",
        "catalina.home": "<application directory>",
        "classpath": "<application directory>/bin/bootstrap.jar:...",
        "etc.",
        "ignore.endorsed.dirs": "",

```

```

        "java.io.tmpdir": "<application directory>/temp",
        "java.protocol.handler.pkgs":
"org.apache.catalina.webresources",
        "java.util.logging.config.file": "<application directory>/conf/
logging.properties",
        "java.util.logging.manager":
"org.apache.juli.ClassLoaderLogManager",
        "jdk.tls.ephemeralDHKeySize": "2048",
        "jdkVersion": "11.0.7",
        "org.apache.catalina.security.SecurityListener.UMASK": ""
    },
    "AdvancedAppInfo": {
        "Directories": {
            "base": "<application directory>",
            "bin": "<application directory>/bin",
            "conf": "<application directory>/conf",
            "home": "<application directory>",
            "lib": "<application directory>/lib",
            "logConfig": "<application directory>/conf/
logging.properties",
            "logs": "<application directory>/logs",
            "tempDir": "<application directory>/temp",
            "webapps": "<application directory>/webapps",
            "work": "<application directory>/work"
        },
        "distro": "java-tomee",
        "flavor": "plume",
        "jdkVersion": "11.0.7",
        "version": "8.0.0"
    },
    "env": {
        "HOME": "... Java Home directory",
        "JDK_JAVA_OPTIONS": "",
        "LANG": "C.UTF-8",
        "LC_TERMINAL": "iTerm2",
        "LC_TERMINAL_VERSION": "3.3.11",
        "LESSCLOSE": "/usr/bin/lesspipe %s %s",
        "LESSOPEN": "| /usr/bin/lesspipe %s",
        "LOGNAME": "ubuntu",
        "LS_COLORS": "",
        "MAIL": "",
        "OLDPWD": "",
        "PATH": "... server PATH",
        "PWD": "",

```

```

        "SHELL": "/bin/bash",
        "SHLVL": "1",
        "SSH_CLIENT": "",
        "SSH_CONNECTION": "",
        "SSH_TTY": "",
        "TERM": "",
        "USER": "ubuntu",
        "XDG_DATA_DIRS": "",
        "XDG_RUNTIME_DIR": "",
        "XDG_SESSION_ID": "1",
        "_": "bin/startup.sh"
    },
    "cwd": "",
    "procUID": {
        "euid": "1000",
        "suid": "1000",
        "fsuid": "1000",
        "ruid": "1000"
    },
    "procGID": {
        "egid": "1000",
        "sgid": "1000",
        "fsgid": "1000",
        "rgid": "1000"
    },
    "userNames": {
        "1000": "ubuntu"
    },
    "groupNames": {
        "1000": "ubuntu"
    },
    "fileDescriptors": [
        "<application directory>/logs/... log files",
        "<application directory>/lib/... jar files",
        "... etc.",
        "/usr/lib/jvm/.../lib/modules"
    ],
    "dependencies": {}
}

```

ASP.NET generic

This example shows an `analysis.json` file for an ASP.NET generic application running on Linux.

```
{
  "a2CTemplateVersion": "1.0",
  "createdTime": "2021-11-24 18:49:1224",
  "containerParameters": {
    "_comment1": "**** EDITABLE: The below section can be edited according
to the application requirements. Please see the analysisInfo section below for
details discovered regarding the application. ****",
    "imageRepository": "dotnet-generic-a27b2829",
    "imageTag": "latest",
    "containerBaseImage": "mcr.microsoft.com/dotnet/sdk:5.0",
    "appExcludedFiles": [
      "/root/.aws"
    ],
    "appSpecificFiles": [],
    "applicationPort": 5000,
    "applicationMode": true,
    "logLocations": [],
    "enableDynamicLogging": false,
    "dependencies": []
  },
  "analysisInfo": {
    "_comment2": "**** NON-EDITABLE: Analysis Results ****",
    "processId": 1,
    "appId": "dotnet-generic-a27b2829",
    "userId": "0",
    "groupId": "0",
    "cmdline": [
      "/usr/bin/dotnet",
      "/root/nopCommerce440/Nop.Web.dll"
    ],
    "webApp": "",
    "osData": {
      "BUG_REPORT_URL": "https://bugs.launchpad.net/ubuntu/",
      "HOME_URL": "https://www.ubuntu.com/",
      "ID": "ubuntu",
      "ID_LIKE": "debian",
      "NAME": "Ubuntu",
      "PRETTY_NAME": "Ubuntu 20.04.3 LTS",
```

```

        "PRIVACY_POLICY_URL": "https://www.ubuntu.com/legal/terms-and-
policies/privacy-policy",
        "SUPPORT_URL": "https://help.ubuntu.com/",
        "UBUNTU_CODENAME": "focal",
        "VERSION": "20.04.3 LTS (Focal Fossa)",
        "VERSION_CODENAME": "focal",
        "VERSION_ID": "20.04"
    },
    "osName": "ubuntu",
    "ports": [
        {
            "localPort": 5000,
            "protocol": "tcp"
        }
    ],
    "Properties": null,
    "applicationType": "dotnet-generic",
    "AdvancedAppInfo": {
        "Directories": {
            "dotnetApp": "/root/nopCommerce440"
        },
        "dotnetVersion": "5.0"
    },
    "env": {
        "HOME": "/root",
        "HOSTNAME": "678f90a12bc3",
        "PATH": "/usr/local/sbin:/usr/local/bin:/usr/sbin:/usr/bin:/sbin:/bin",
        "TERM": "xterm",
        "TZ": "Etc/UTC"
    },
    "cwd": "",
    "exe": "/usr/share/dotnet/dotnet",
    "procUID": {
        "euid": "0",
        "suid": "0",
        "fsuid": "0",
        "ruid": "0"
    },
    "procGID": {
        "egid": "0",
        "sgid": "0",
        "fsgid": "0",
        "rgid": "0"
    },
    },

```



```

        "userNames": {
            "0": "root"
        },
        "groupNames": {
            "0": "root"
        },
        "fileDescriptors": [
            "/dev/pts/0",
            "/root/nopCommerce440/AdvancedStringBuilder.dll",
            "/root/nopCommerce440/AutoMapper.dll",
            "... etc.",
            "/root/nopCommerce440/netstandard.dll"
        ],
        "dependencies": {}
    }
}

```

ASP.NET single file

This example shows an `analysis.json` file for an ASP.NET single file application running on Linux.

```

{
  "a2CTemplateVersion": "1.0",
  "createdTime": "2021-11-29 07:08:2929",
  "containerParameters": {
    "_comment1": "*** EDITABLE: The below section can be edited according to the application requirements. Please see the analysisInfo section below for details discovered regarding the application. ***",
    "imageRepository": "dotnet-single-c2930d3132",
    "imageTag": "latest",
    "containerBaseImage": "mcr.microsoft.com/dotnet/sdk:latest",
    "appExcludedFiles": [
      "/root/.aws"
    ],
    "appSpecificFiles": [],
    "applicationPort": 5000,
    "applicationMode": true,
    "logLocations": [],
    "enableDynamicLogging": false,
    "dependencies": []
  },
  "analysisInfo": {

```

```

    "_comment2": "*** NON-EDITABLE: Analysis Results ***",
    "processId": 1,
    "appId": "dotnet-single-c2930d3132",
    "userId": "0",
    "groupId": "0",
    "cmdline": [
        "./MyCoreWebApp.5"
    ],
    "webApp": "",
    "osData": {
        "BUG_REPORT_URL": "https://bugs.launchpad.net/ubuntu/",
        "HOME_URL": "https://www.ubuntu.com/",
        "ID": "ubuntu",
        "ID_LIKE": "debian",
        "NAME": "Ubuntu",
        "PRETTY_NAME": "Ubuntu 20.04.3 LTS",
        "PRIVACY_POLICY_URL": "https://www.ubuntu.com/legal/terms-and-
policies/privacy-policy",
        "SUPPORT_URL": "https://help.ubuntu.com/",
        "UBUNTU_CODENAME": "focal",
        "VERSION": "20.04.3 LTS (Focal Fossa)",
        "VERSION_CODENAME": "focal",
        "VERSION_ID": "20.04"
    },
    "osName": "ubuntu",
    "ports": [
        {
            "localPort": 5000,
            "protocol": "tcp"
        }
    ],
    "Properties": null,
    "applicationType": "dotnet-single",
    "AdvancedAppInfo": {
        "Directories": {
            "dotnetApp": "/root/mycorewebapp"
        },
        "dotnetVersion": "latest"
    },
    "env": {
        "HOME": "/root",
        "HOSTNAME": "a1bc23d4567e",
        "PATH": "/usr/local/sbin:/usr/local/bin:/usr/sbin:/usr/bin:/sbin:/bin",
        "TERM": "xterm",

```

```

    "TZ": "Etc/UTC"
  },
  "cwd": "/root/mycorewebapp",
  "exe": "/root/mycorewebapp/MyCoreWebApp.5",
  "procUID": {
    "euid": "0",
    "suid": "0",
    "fsuid": "0",
    "ruid": "0"
  },
  "procGID": {
    "egid": "0",
    "sgid": "0",
    "fsgid": "0",
    "rgid": "0"
  },
  "userNames": {
    "0": "root"
  },
  "groupNames": {
    "0": "root"
  },
  "fileDescriptors": [
    "/dev/pts/0",
    "/root/mycorewebapp/MyCoreWebApp.5"
  ],
  "dependencies": {}
}

```

Containers running on Windows

For applications running on Windows, the `application analysis.json` file includes the following content:

Read-only data

- **Control fields** – Fields having to do with file creation, such as template version, and the file creation timestamp.
- **analysisInfo** – System dependencies for the application.

Configurable data

The `containerParameters` section contains the following fields:

- **containerBaseImage** (string) – The base operating system (OS) image for the container build. By default, App2Container uses the operating system from the application server or worker machine where containerization runs.

Note

If specified, this must be an image name from your registry in the format `<image name>[:<tag>]`, and it must match the operating system platform and version that runs on the application server or worker machine where containerization runs. The tag is optional if the repository supports "latest".

- **enableServerConfigurationUpdates** (Boolean, required) – Provides an option in the Dockerfile to restore the application configuration of the source server.
- **imageRepositoryName** (string, required) – The name of the repository where the application container image is stored.
- **imageTag** (string, required) – A tag for the build version of the application container image.
- **additionalExposedPorts** (array of numbers) – Additional port numbers that should be exposed inside of the application container.
- **appIncludedFiles** (array of strings) – Specific files and directories to include in the container build.
- **appExcludedFiles** (array of strings) – Specific files and directories to exclude from the container build.
- **enableLogging** (Boolean, required) – Enables dynamic logging, redirecting application logs to container `stdout`.
- **includedWebApps** (array of strings) – The application IDs for web applications running under the IIS site that should be included in the container image. *Applications must have been running in IIS during inventory and analysis.*
- **additionalApps** (array of strings) – For the `analysis.json` file that describes the root application in a complex Windows .NET application, these are the additional application or service components to include in the application container. You can include up to five additional application components in the array.

Examples

The following examples show an `analysis.json` file for a .NET application running on Windows. Your `analysis.json` file configuration can vary by the type of .NET application you are migrating, its dependencies, and whether you want it to run in a single container or in multiple containers. Choose the tab that matches your .NET configuration.

Simple

The following example shows an `analysis.json` file for a simple .NET application running on Windows.

```
{
  "a2CTemplateVersion": "3.1",
  "createdTime": "",
  "containerParameters": {
    "_comment": "**** EDITABLE: The below section can be edited according to the application requirements. Please see the Analysis Results section further below for details discovered regarding the application. ****",
    "containerBaseImage": "mcr.microsoft.com/dotnet/framework/aspnet:4.7.2-windowsservercore-ltsc2019",
    "enableServerConfigurationUpdates": true,
    "imageRepositoryName": "iis-smarts-51d2dbf8",
    "imageTag": "latest",
    "additionalExposedPorts": [
    ],
    "appIncludedFiles": [

    ],
    "appExcludedFiles": [
    ],
    "enableLogging": false,
    "additionalApps": [
    ]
  },
  "analysisInfo": {
    "_comment": "**** NON-EDITABLE: Analysis Results ****",
    "hostInfo": {
      "os": "...",
      "osVersion": "...",
      "osWindowsDirectory": "...",
      "arch": "..."
    }
  },
}
```

```

"appId": "iis-smarts-51d2dbf8",
"appServerIp": "localhost",
"appType": "IIS",
"appName": "smarts",
"ports": [
  {
    "localPort": 90,
    "protocol": "http"
  }
],
"features": [
  "File-Services",
  "FS-FileServer",
  "Web-Http-Tracing",
  "Web-Basic-Auth",
  "Web-Digest-Auth",
  "Web-Url-Auth",
  "Web-Windows-Auth",
  "Web-ASP",
  "Web-CGI",
  "Web-Mgmt-Tools",
  "Web-Mgmt-Console",
  "Web-Scripting-Tools",
  "FS-SMB1",
  "User-Interfaces-Infra",
  "Server-Gui-Mgmt-Infra",
  "Server-Gui-Shell",
  "PowerShell-ISE"
],
"appPoolName": "smarts",
"poolIdentityType": "ApplicationPoolIdentity",
"dotnetVersion": "v4.0",
"iisVersion": "IIS 10.0",
"sitePhysicalPath": "<IIS web root directory>\\smarts",
"discoveredWebApps": [
],
"reportPath": "<application output directory>\\iis-smarts-51d2dbf8\\report.txt",
"isSiteUsingWindowsAuth": false,
"serverBackupFile": "<application directory>\\Web Deploy Backups\\... backup zip
file"
}
}

```

Complex – one container

In this scenario, each application or service has its own `analysis.json` file, but the root application references the application ID for the service in the `additionalApps` array. This results in a single container that includes both the root application and the service when you run the **containerize** command.

- **Root application**

The following example shows the `analysis.json` file for the root application.

```
{
  "a2CTemplateVersion": "1.0",
  "createdTime": "2021-06-25-05:18:24",
  "containerParameters": {
    "_comment": "**** EDITABLE: The below section can be edited according to the
application requirements. Please see the Analysis Results section further below
for details discovered regarding the application. ****",
    "containerBaseImage": "",
    "enableServerConfigurationUpdates": true,
    "imageRepositoryName": "iis-colormvciis-b69c09ab",
    "imageTag": "latest",
    "additionalExposedPorts": [
    ],
    "appIncludedFiles": [
    ],
    "appExcludedFiles": [
    ],
    "enableLogging": false,
    "includedWebApps": [
    ],
    "additionalApps": [
      "service-colorwindowsservice-69f90194"
    ]
  },
  "analysisInfo": {
    "_comment": "**** NON-EDITABLE: Analysis Results ****",
    "hostInfo": {
      "os": "Microsoft Windows Server 2019 Datacenter",
      "osVersion": "10.0.17763",
      "osWindowsDirectory": "C:\\\\Windows",
      "arch": "64-bit"
    }
  },
}
```

```

    "appId": "iis-colormvcIis-b69c09ab",
    "appServerIp": "localhost",
    "appType": "IIS",
    "appName": "colorMvcIIs",
    "ports": [
      {
        "localPort": 82,
        "protocol": "http"
      }
    ],
    "features": [
      "Web-Http-Redirect",
      "Web-Custom-Logging",
      "... etc."
    ],
    "appPoolName": "colorMVC",
    "poolIdentityType": "ApplicationPoolIdentity",
    "dotNetVersion": "v4.0",
    "iisVersion": "IIS 10.0",
    "sitePhysicalPath": "C:\\\\colorMvcIis",
    "discoveredWebApps": [
    ],
    "siteUsesWindowsAuth": false,
    "serverBackupFile": "<application directory>\\\\Web Deploy Backups\\\\... backup
zip file",
    "reportPath": "<application output directory>\\\\iis-colormvcIis-b69c09ab\\
\\report.txt"
  }
}

```

• Windows service

The following example shows the `analysis.json` file for the Windows service that is included in the application container.

```

{
  "a2CTemplateVersion": "1.0",
  "createdTime": "2021-07-09-04:16:58",
  "containerParameters": {
    "_comment": "**** EDITABLE: The below section can be edited according to the
application requirements. Please see the Analysis Results section further below
for details discovered regarding the application. ****",
    "containerBaseImage": "",

```



```

    "enableServerConfigurationUpdates": true,
    "imageRepositoryName": "service-colorwindowsservice-69f90194",
    "imageTag": "latest",
    "additionalExposedPorts": [
    ],
    "appIncludedFiles": [
    ],
    "appExcludedFiles": [
    ],
    "enableLogging": false,
    "additionalApps": [
    ]
  },
  "analysisInfo": {
    "_comment": "*** NON-EDITABLE: Analysis Results ***",
    "hostInfo": {
      "os": "Microsoft Windows Server 2019 Datacenter",
      "osVersion": "10.0.17763",
      "osWindowsDirectory": "C:\\\\Windows",
      "arch": "64-bit"
    },
    "appId": "service-colorwindowsservice-69f90194",
    "appServerIp": "localhost",
    "appType": "service",
    "appName": "colorwindowsservice",
    "ports": [
      {
        "localPort": 33335,
        "protocol": "TCP"
      }
    ],
    "features": [
      "Web-Http-Redirect",
      "Web-Custom-Logging",
      "... etc."
    ],
    "serviceName": "colorwindowsservice",
    "serviceBinary": "ColorWindowsService.exe",
    "serviceDir": "C:\\\\COLORCODE\\\\colorservice-master\\\\ColorWindowsService\\\\bin\\\\
\\Release\\\\",
    "cmdline": [
      "C:\\\\COLORCODE\\\\colorservice-master\\\\ColorWindowsService\\\\bin\\\\Release\\\\
\\ColorWindowsService.exe"
    ]
  }
}

```

```
}
}
```

Complex – multiple containers

In this scenario, each application or service has its own `analysis.json` file, and the `additionalApps` array is empty. To create two containers, run the **containerize** command twice – once for the root application and once for the service. For container orchestration, specify the service as a dependent application when you configure the `deployment.json` file for the root application.

- **Root application**

The following example shows the `analysis.json` file for the root application.

```
{
  "a2CTemplateVersion": "1.0",
  "createdTime": "2021-06-25-05:18:24",
  "containerParameters": {
    "_comment": "**** EDITABLE: The below section can be edited according to the
application requirements. Please see the Analysis Results section further below
for details discovered regarding the application. ****",
    "containerBaseImage": "",
    "enableServerConfigurationUpdates": true,
    "imageRepositoryName": "iis-colormvciis-b69c09ab",
    "imageTag": "latest",
    "additionalExposedPorts": [
    ],
    "appIncludedFiles": [
    ],
    "appExcludedFiles": [
    ],
    "enableLogging": false,
    "includedWebApps": [
    ],
    "additionalApps": [
    ]
  },
  "analysisInfo": {
    "_comment": "**** NON-EDITABLE: Analysis Results ****",
    "hostInfo": {
      "os": "Microsoft Windows Server 2019 Datacenter",
```

```

        "osVersion": "10.0.17763",
        "osWindowsDirectory": "C:\\\\Windows",
        "arch": "64-bit"
    },
    "appId": "iis-colormvciis-b69c09ab",
    "appServerIp": "localhost",
    "appType": "IIS",
    "appName": "colorMvcIIs",
    "ports": [
        {
            "localPort": 82,
            "protocol": "http"
        }
    ],
    "features": [
        "Web-Http-Redirect",
        "Web-Custom-Logging",
        "... etc."
    ],
    "appPoolName": "colorMVC",
    "poolIdentityType": "ApplicationPoolIdentity",
    "dotNetVersion": "v4.0",
    "iisVersion": "IIS 10.0",
    "sitePhysicalPath": "C:\\\\colorMvcIis",
    "discoveredWebApps": [
    ],
    "siteUsesWindowsAuth": false,
    "serverBackupFile": "<application directory>\\\\Web Deploy Backups\\\\... backup
zip file",
    "reportPath": "<application output directory>\\\\iis-colormvciis-b69c09ab\\
\\report.txt"
    }
}

```

• Windows service

The following example shows the `analysis.json` file for the Windows service that runs in a separate container.

```

{
    "a2CTemplateVersion": "1.0",
    "createdTime": "2021-07-09-04:16:58",
    "containerParameters": {

```

```

    "_comment": "*** EDITABLE: The below section can be edited according to the
application requirements. Please see the Analysis Results section further below
for details discovered regarding the application. ***",
    "containerBaseImage": "",
    "enableServerConfigurationUpdates": true,
    "imageRepositoryName": "service-colorwindowsservice-69f90194",
    "imageTag": "latest",
    "additionalExposedPorts": [
    ],
    "appIncludedFiles": [
    ],
    "appExcludedFiles": [
    ],
    "enableLogging": false,
    "additionalApps": [
    ]
  },
  "analysisInfo": {
    "_comment": "*** NON-EDITABLE: Analysis Results ***",
    "hostInfo": {
      "os": "Microsoft Windows Server 2019 Datacenter",
      "osVersion": "10.0.17763",
      "osWindowsDirectory": "C:\\Windows",
      "arch": "64-bit"
    },
    "appId": "service-colorwindowsservice-69f90194",
    "appServerIp": "localhost",
    "appType": "service",
    "appName": "colorwindowsservice",
    "ports": [
      {
        "localPort": 33335,
        "protocol": "TCP"
      }
    ],
    "features": [
      "Web-Http-Redirect",
      "Web-Custom-Logging",
      "... etc."
    ],
    "serviceName": "colorwindowsservice",
    "serviceBinary": "ColorWindowsService.exe",
    "serviceDir": "C:\\COLORCODE\\colorservice-master\\ColorWindowsService\\bin\\
\\Release\\"
  }
}

```

```
"cmdline": [  
    "C:\\\\COLORCODE\\\\colorservice-master\\\\ColorWindowsService\\\\bin\\\\Release\\\\  
    \\\\ColorWindowsService.exe"  
  ]  
}
```

Note

For complex Windows .NET applications, you can also use a hybrid approach, with some components running together in a single container and other components running in separate containers.

Configuring container deployment

This topic contains information about the files that are used for configuring deployment of application containers.

Container deployment files

- [deployment.json file](#)

deployment.json file

When you run the [containerize](#) command, a `deployment.json` file is created for the application specified in the `--application-id` parameter. The **generate app-deployment** command uses this file, along with others, to generate application deployment artifacts. All of the fields in this file are configurable as needed so that you can customize your application container deployment before running the **generate app-deployment** command.

Important

The `deployment.json` file includes sections for both Amazon ECS and Amazon EKS. If your application is suitable for App Runner, there is a section for that too. Set the Boolean value deployment flag for the section that matches your target container management service to **true**. Set the other flags to **false**. The flag to deploy to Amazon ECS

is `createEcsArtifacts`, the flag to deploy to Amazon EKS is `createEksArtifacts`, and the flag to deploy to App Runner is `createAppRunnerArtifacts`.

The application `deployment.json` file includes the following content. While all fields are configurable, the following fields should not be changed: `a2CTemplateVersion`, `applicationId`, and `imageName`. For key-value pairs that do not apply to your deployment, set string values to an empty string, numeric values to zero, and Boolean values to false.

- **exposedPorts** (array of objects, required) – An array of JSON objects representing the ports that should be exposed when the container is running. Each object consists of the following fields:
 - **localPort** (number) – A port to expose for container communication.
 - **protocol** (string) – The application protocol for the exposed port, for example, "http".
- **environment** (array of objects) – Environment variables to be passed on to the target container management deployment. For Amazon ECS deployments, the key-value pairs update the Amazon ECS task definition. For Amazon EKS deployments, the key-value pairs update the Kubernetes `deployment.yaml` file.
- **ecrParameters** (object) – Contains parameters needed to register application container images in Amazon ECR.
 - **ecrRepoTag** (string, required) – The version tag to use for registering an application container image in Amazon ECR.
- **ecsParameters** (object) – Contains parameters needed for deployment to Amazon ECS. *The `createEcsArtifacts` parameter is always required. Other parameters in this section that are marked as required apply only to Amazon ECS deployment.*
 - **createEcsArtifacts** (Boolean, required) – A flag that indicates if you are targeting Amazon ECS for deployment.
 - **ecsFamily** (string, required) – An ID for the Amazon ECS family in the Amazon ECS task definition. We recommend setting this value to the application ID.
 - **cpu** (number, required*) – The hard limit for the number of vCPUs to present for the task. When the task definition is registered, the number of CPU units is determined by multiplying the number of vCPUs by 1024.

** This parameter is required for Linux containers, but is not supported for Windows containers.*

- **memory** (number or string, required*) – The hard limit of memory (in MiB) to present to the task. You can express this value as an integer in the Amazon ECS task definition, using MiB, for

example, 1024. You can also express the value as a string including the unit GB, for example, 1 GB. When the task definition is registered, a GB value is converted to an integer indicating the MiB.

** This parameter is required for Linux containers, but is not supported for Windows containers.*

Note

In the Amazon ECS task definition, task size consists of the `cpu` and `memory` parameters. The configuration for task size, in part, depends on where your tasks are hosted – on an EC2 instance, or in Fargate. For more information about setting the task size for your Amazon ECS task definition, see [Task definition parameters](#) in the *Amazon Elastic Container Service Developer Guide*.

- **dockerSecurityOption** (string) – For .NET applications, this is the gMSA Credspec location value for the Amazon ECS task definition.
- **enableCloudwatchLogging** (Boolean, required*) – A flag that sets the Amazon ECS task definition to turn on CloudWatch logging for your Windows application container. *If set to true, the `enableLogging` field in the `analysis.json` file must have a valid value.*

** This parameter is required for Windows containers, but is not supported for Linux containers.*

- **publicApp** (Boolean, required) – A flag to configure the CloudFormation templates with a public endpoint for your application when it runs.
- **stackName** (string, required) – A name to use as a prefix to your CloudFormation stack Amazon Resource Name (ARN). We recommend using the application ID for this.
- **resourceTags** (array of objects) – Custom tags, expressed as key/value pairs that are added to resources during deployment. For Amazon ECS deployments, the key-value pairs update the Amazon ECS task definition.

Note

An example tag is generated when the `deployment.json` file is created. If the example tag isn't removed or changed before deployment, it's ignored by default.

- **reuseResources** (object) – Contains shared resource identifiers that can be used throughout your CloudFormation templates.

- **vpId** (string) – The VPC ID, if you want to bring your own VPC or to reuse an existing VPC that App2Container created for a prior deployment.
- **reuseExistingA2cStack** (object) – Contains references so that you can reuse AWS CloudFormation resources that App2Container has already created.
- **cfnStackName** (string) – The name or ID (ARN) of the CloudFormation stack created with App2Container for the containerized application.
- **microserviceUriPath** (string) – Used to create application forwarding rules in your load balancer.

 Note

The load balancer does not strip off this prefix when it forwards traffic. Your application must be able to handle requests coming in with the prefix.

- **sshKeyPairName** (string) – The name of the EC2 key pair to use for the instances that your container runs on.
- **acmCertificateArn** (string) – The AWS Certificate Manager certificate ARN used to provide HTTPS connectivity to your Application Load Balancer.

 Note

The certificate can be imported or provisioned as follows:

- To import an IIS certificate into ACM, see [How to import PFX-formatted certificates into AWS Certificate Manager using OpenSSL](#).
- To provision a certificate in ACM, see [Issuing and Managing Certificates](#) in the AWS Certificate Manager User Guide.

If you use an HTTPS endpoint for your load balancer, this parameter is required. For more information about ACM, see [What is AWS Certificate Manager](#) in the AWS Certificate Manager User Guide.

- **gMSAParameters** (object) – Contains parameters used by the CloudFormation template to create gMSA-related artifacts for .NET applications that are deployed on EC2 instances. The gMSAParameters are not valid for deployments to Fargate, and will generate an error when the **generate app-deployment** command runs.

- **domainSecretsArn** (string) – The Secrets Manager ARN for the domain credentials to join the Amazon ECS nodes to gMSA Active Directory.
- **domainDNSName** (string) – The DNS name of the gMSA Active Directory for Amazon ECS nodes to join.
- **domainNetBIOSName** (string) – The NetBIOS name of the Active Directory for Amazon ECS nodes to join.
- **createGMSA** (Boolean, required) – A flag to create a group Managed Service Account (gMSA) Active Directory security group and account, using the name supplied in the gMSAName field.
- **gMSAName** (string) – The name of the gMSA Active Directory that the container should use for access.
- **deployTarget** (string, required) – Identifies which Amazon ECS container launch type runs the task. Valid values depend on your application environment, as follows:
 - *.NET applications running on Windows* – ec2, fargate.
 - *Java applications running on Linux* – fargate.

Note

The default value that is generated for the **deployTarget** parameter for .NET applications running on Windows is ec2. To deploy your application to Fargate, you can edit the deployment .json file, and change that value to fargate.

If your .NET application meets the following criteria, you can deploy to Fargate.

- The base operating system for your container is Windows 2019. If you are using a worker machine for containerization, this means that the worker machine must be running Windows 2019.
 - Your application must not use gMSA.
- **dependentApps** (array of objects) – For complex Windows applications, this array of JSON objects contains identifying details for dependent applications. App2Container does not generate this array. For complex Windows applications that incorporate dependent applications, you must add details to this array for each dependent application. You can include up to two dependent applications in the array.
 - **appId** (string, required) – The application ID that App2Container generated for this dependent application.

- **privateRootDomain** (string, required) – The private domain name that's used for creating the hosted zone.
- **dnsRecordName** (string, required) – The DNS record name of the application. This is combined with the `privateRootDomain` to construct the endpoint for the dependent application.
- **fireLensParameters** (object) – Contains parameters needed to use FireLens with your Linux application to route your application logs for Amazon ECS tasks. *The `enableFireLensLogging` parameter is always required. Other parameters in this section that are marked as required apply only when FireLens is used for log routing.*

 Note

This section is not included for applications running on Windows.

- **enableFireLensLogging** (Boolean, required) – A flag for using FireLens for Amazon ECS to configure application log routing for containers.
- **logDestinations** (array of objects) – A list of unique target destinations for application log routing. If more than one destination is configured, App2Container creates a custom file that contains the FireLens configuration. Otherwise, the destination parameters are used directly in the Amazon ECS task definition and CloudFormation templates.
 - **service** (string) – The AWS service to route logs to. *Valid values are "cloudwatch", "firehose", and "kinesis".*
 - **regexFilter** (string) – A Ruby regular expression to match against log content to determine where to route the log.
 - **streamName** (string) – The name of the log delivery stream that will be created at the destination.
- **eksParameters** (object) – Contains parameters needed for deployment to Amazon EKS. *The `createEksArtifacts` parameter is always required. Other parameters in this section that are marked as required apply only to Amazon EKS deployments.*
 - **createEksArtifacts** (Boolean, required) – A flag that indicates if you are targeting Amazon EKS for deployment.
 - **stackName** (string, required) – A name to use as a prefix to your CloudFormation stack ID ARN. We recommend using the application ID for this.

- **cpu** (number, required) – The hard limit for the number of vCPUs to present for the application container. The minimum value is .25, and the maximum value is 1.5. If there are no overrides, the default value is 1.5.
- **memory** (number, required) – The hard limit of memory (in MiB) for the application container. Express this value as an integer, for example, 1024.
- **ingress** (string, required) – The type of load balancer to use for the deployment. Specify one of the following values:
 - **alb** – Provisions an Application Load Balancer in the VPC for the deployment.
 - **nginx** – Provisions a Network Load Balancer in the VPC, and an NGINX ingress in the Kubernetes cluster for the deployment.

 Note

If you upgrade from a previous App2Container deployment, the load balancer URL might change.

- **dnsRecordName** (string) – The fully qualified domain name (FQDN) for a DNS record for the deployed application, for example *hello.example.com*. If you specify this parameter, then App2Container creates the DNS record in a private hosted zone in Amazon Route 53. If you also specify the **rootDomain** parameter, then App2Container creates the DNS record in the specified root domain.
- **applicationPath** (string) – The location of the application from the root of the web server, as accessed from the public URL, for example */my-application*.
- **reuseResources** (object) – Contains shared resource identifiers that can be used throughout your CloudFormation templates.
- **vpclId** (string) – The VPC ID, if you want to bring your own VPC or to reuse an existing VPC that App2Container created for a prior deployment. *If you bring a custom VPC, you must have two or more private subnets in two or more Availability Zones. In this case, you can optionally have two or more public subnets in the same two Availability Zones.*

 Note

For each private subnet in the reused VPC, you must configure a route to the internet using a NAT gateway. For more information about cluster networking for Amazon EKS, see [De-mystifying cluster networking for Amazon EKS worker nodes](#).

- **cfnStackName** (string) – The name or ID (ARN) of the CloudFormation stack created with App2Container for the containerized application.
- **sshKeyName** (string) – The name of the Amazon EC2 key pair to use for the instances that your container runs on.
- **resourceTags** (array of objects) – Custom tags, expressed as key/value pairs that are added to resources during deployment. For Amazon EKS deployments, the key/value pairs update the Kubernetes `deployment.yml` file.

 Note

An example tag is generated when the `deployment.json` file is created. If you don't remove or change the example tag before deployment, the tag is ignored by default.

- **rootDomain** (string) – The name of a root domain (hosted zone) in Amazon Route 53, for example *example.com*. If you specify the `rootDomain`, then App2Container creates the DNS record that points to it.
- **acmCertificateArn** (string) – The AWS Certificate Manager certificate ARN used to provide HTTPS connectivity to your Application Load Balancer. If you don't specify a value for **acmCertificateArn**, App2Container can only deploy HTTP applications.

 Note

The certificate can be imported or provisioned as follows:

- To import an IIS certificate into ACM, see [How to import PFX-formatted certificates into AWS Certificate Manager using OpenSSL](#).
- To provision a certificate in ACM, see [Issuing and Managing Certificates](#) in the AWS Certificate Manager User Guide.

If you use an HTTPS endpoint for your load balancer, this parameter is required. For more information about ACM, see [What is AWS Certificate Manager](#) in the AWS Certificate Manager User Guide.

- **gMSAParameters** (object) – Contains parameters used by the CloudFormation template to create gMSA-related artifacts for .NET applications.

- **domainSecretsArn** (string) – The Secrets Manager ARN for the domain credentials to join the Amazon EKS nodes to gMSA Active Directory.
- **domainDNSName** (string) – The DNS name of the gMSA Active Directory for Amazon EKS nodes to join.
- **domainNetBIOSName** (string) – The NetBIOS name of the Active Directory for Amazon EKS nodes to join.
- **createGMSA** (Boolean, required) – A flag to create a group Managed Service Account (gMSA) Active Directory security group and account, using the name supplied in the gMSAName field.
- **gMSAAccountName** (string) – The name of the gMSA Active Directory that the container should use for access.
- **dependentApps** (array of objects) – For complex Windows applications, this array of JSON objects contains identifying details for dependent applications. App2Container does not generate this array. For complex Windows applications that incorporate dependent applications, you must add details to this array for each dependent application. You can include up to two dependent applications in the array.
 - **appId** (string, required) – The application ID that App2Container generated for this dependent application.
 - **privateRootDomain** (string, required) – The private domain name that's used for creating the hosted zone.
 - **dnsRecordName** (string, required) – The DNS record name of the application. This is combined with the privateRootDomain to construct the endpoint for the dependent application.
- **appRunnerParameters** (object) – Contains parameters needed for deployment of Linux applications to an AWS App Runner environment. *The createAppRunnerArtifacts parameter is always required. Other parameters in this section that are marked as required apply only to App Runner deployments.*

 Note

This section is not included for applications running on Windows.

- **createAppRunnerArtifacts** (Boolean, required) – A flag that indicates if you are targeting App Runner for deployment.

- **stackName** (string, required) – The name of the AWS CloudFormation stack. *We recommend including the application ID in the stack name.*
- **serviceName** (string, required) – The name of the service in App Runner. *We recommend using the application ID for the service name.*
- **autoDeploymentsEnabled** (Boolean, required) – If set to true, an update to the Amazon ECR repository also updates the service in App Runner. If set to false, you can manually update the service using the App Runner console or API, or `apprunner` commands in the AWS CLI.
- **resourceTags** (array of objects) – Custom tags, expressed as key/value pairs that are added to resources during deployment. For App Runner deployments, the key/value pairs update both of the resources that are created in the `apprunner.yml` AWS CloudFormation template.

Note

An example tag is generated when the `deployment.json` file is created. If the example tag isn't removed or changed before deployment, the tag is ignored by default.

Note

When the **containerize** command runs, it determines if your application is suitable for App Runner, and adds `appRunnerParameters` to the `deployment.json` file if it is. If your application is not suitable for App Runner, the `appRunnerParameters` are ignored.

Examples

Linux Java application deployed to Amazon ECS

The following example shows a `deployment.json` file for a Java application running on Linux, with default settings to deploy to an Amazon ECS environment.

```
{
  "a2CTemplateVersion": "3.1",
  "applicationId": "java-tomcat-6e6f3a87",
  "imageName": "java-tomcat-6e6f3a87",
  "exposedPorts": [
    {
      "localPort": 8080,
```

```

        "protocol": "tcp6"
    },
    {
        "localPort": 8009,
        "protocol": "tcp6"
    },
    {
        "localPort": 8005,
        "protocol": "tcp6"
    }
],
"environment": [],
"ecrParameters": {
    "ecrRepoTag": "latest"
},
"ecsParameters": {
    "createEcsArtifacts": true,
    "ecsFamily": "java-tomcat-6e6f3a87",
    "cpu": 2,
    "memory": 4096,
    "dockerSecurityOption": "",
    "enableCloudwatchLogging": false,
    "publicApp": true,
    "stackName": "app2container-java-tomcat-6e6f3a87-ECS",
    "resourceTags": [
        {
            "key": "example-key",
            "value": "example-value"
        }
    ],
    "reuseResources": {
        "vpcId": "",
        "reuseExistingA2cStack": {
            "cfnStackName": "",
            "microserviceUrlPath": ""
        },
        "sshKeyPairName": "",
        "acmCertificateArn": ""
    },
    "gMSAParameters": {
        "createGMSA": false,
        "domainSecretsArn": "",
        "domainDNSName": "",
        "domainNetBIOSName": "",

```

```

        "gMSAName": "",
        "ADSecurityGroupName": ""
    },
    "deployTarget": "fargate"
},
"fireLensParameters": {
    "enableFireLensLogging": true,
    "logDestinations": [
        {
            "service": "cloudwatch",
            "matchRegex": "^.*INFO.*$",
            "streamName": "Info"
        },
        {
            "service": "cloudwatch",
            "matchRegex": "^.*WARN.*$",
            "streamName": "Warn"
        }
    ]
},
"eksParameters": {
    "createEksArtifacts": false,
    "stackName": "java-tomcat-6e6f3a87",
    "reuseResources": {
        "vpcId": "",
        "cfnStackName": "",
        "sshKeyPairName": ""
    },
    "gMSAParameters": {
        "createGMSA": false,
        "domainSecretsArn": "",
        "domainDNSName": "",
        "domainNetBIOSName": "",
        "gMSAAccountName": "",
        "ADSecurityGroupName": ""
    }
},
"appRunnerParameters": {
    "createAppRunnerArtifacts": false,
    "stackName": "a2c-java-tomcat-6e6f3a87-AppRunner",
    "autoDeploymentsEnabled": true,
    "resourceTags": [
        {
            "key": "example-key",

```



```

        "value": "example-value"
      }
    ]
  }
}

```

Windows .NET application deployed to AWS Fargate

The following example shows a deployment .json file for a .NET application running on Windows. The application has been configured to deploy to an Amazon ECS Fargate environment.

```

{
  "a2CTemplateVersion": "3.1",
  "applicationId": "iis-smarts-51d2dbf8",
  "imageName": "iis-smarts-51d2dbf8",
  "exposedPorts": [
    {
      "localPort": 8080,
      "protocol": "http"
    }
  ],
  "environment": [],
  "ecrParameters": {
    "ecrRepoTag": "latest"
  },
  "ecsParameters": {
    "createEcsArtifacts": true,
    "ecsFamily": "iis-smarts-51d2dbf8",
    "cpu": 2,
    "memory": 4096,
    "dockerSecurityOption": "",
    "enableCloudwatchLogging": false,
    "publicApp": true,
    "stackName": "iis-smarts-51d2dbf8-ECS",
    "resourceTags": [
      {
        "key": "example-key",
        "value": "example-value"
      }
    ],
    "reuseResources": {
      "vpcId": "vpc-0abc1defa2345b67c",
      "reuseExistingA2cStack": {

```

```

        "cfnStackName": "",
        "microserviceUrlPath": ""
    },
    "sshKeyPairName": "",
    "acmCertificateArn": ""
},
"gMSAParameters": {
    "domainSecretsArn": "",
    "domainDNSName": "",
    "domainNetBIOSName": "",
    "createGMSA": false,
    "gMSAName": ""
},
"deployTarget": "fargate",
"dependentApps" : []
},
"eksParameters": {
    "createEksArtifacts": false,
    "stackName": "iis-smarts-51d2dbf8-EKS",
    "reuseResources": {
        "vpcId": "",
        "reuseExistingA2cStack": {
            "cfnStackName": "",
            "microserviceUrlPath": ""
        },
        "sshKeyPairName": ""
    },
    "gMSAParameters": {
        "createGMSA": false,
        "domainSecretsArn": "",
        "domainDNSName": "",
        "domainNetBIOSName": "",
        "gMSAAccountName": ""
    },
    "dependentApps" : []
}
}

```

Complex Windows .NET application deployed to Amazon ECS

For complex Windows .NET web applications that consist of a root application and up to two dependent applications, each application is defined separately. Each application has its own deployment.json file.

The following example shows the `deployment.json` file for the root application in a complex .NET web service running on Windows, followed by `deployment.json` files for the two dependent applications that it refers to. The applications are deployed to an Amazon ECS environment running together in the same VPC.

- **Root application example**

```
{
  "a2CTemplateVersion": "3.1",
  "applicationId": "iis-smarts-51d2dbf8",
  "imageName": "iis-smarts-51d2dbf8",
  "exposedPorts": [
    {
      "localPort": 8080,
      "protocol": "http"
    }
  ],
  "environment": [],
  "ecrParameters": {
    "ecrRepoTag": "latest"
  },
  "ecsParameters": {
    "createEcsArtifacts": true,
    "ecsFamily": "iis-smarts-51d2dbf8",
    "cpu": 2,
    "memory": 4096,
    "dockerSecurityOption": "",
    "enableCloudwatchLogging": false,
    "publicApp": true,
    "stackName": "iis-smarts-51d2dbf8-ECS",
    "resourceTags": [
      {
        "key": "example-key",
        "value": "example-value"
      }
    ],
    "reuseResources": {
      "vpcId": "vpc-0abc1defa2345b67c",
      "reuseExistingA2cStack": {
        "cfnStackName": "",
        "microserviceUrlPath": ""
      },
      "sshKeyPairName": "",

```

```

        "acmCertificateArn": ""
    },
    "gMSAParameters": {
        "createGMSA": false,
        "domainSecretsArn": "",
        "domainDNSName": "",
        "domainNetBIOSName": "",
        "gMSAName": ""
    },
    "deployTarget": "ec2",
    "dependentApps" : [
        {
            "appId":"iis-appB-ab800cde",
            "privateRootDomain": "dependent-app1.test1.com",
            "dnsRecordName":"appB"
        },
        {
            "appId":"service-appC-9fghi90j",
            "privateRootDomain": "dependent-app2.test1.com",
            "dnsRecordName":"appC"
        }
    ]
},
"eksParameters": {
    "createEksArtifacts": false,
    "stackName": "iis-smarts-51d2dbf8",
    "reuseResources": {
        "vpcId": "",
        "reuseExistingA2cStack": {
            "cfnStackName": "",
            "microserviceUrlPath": ""
        },
        "sshKeyPairName": "",
        "acmCertificateArn": ""
    },
    "gMSAParameters": {
        "createGMSA": false,
        "domainSecretsArn": "",
        "domainDNSName": "",
        "domainNetBIOSName": "",
        "gMSAAccountName": ""
    },
    "dependentApps" : []
}

```

```
}
```

• Dependent application B

```
{
  "a2CTemplateVersion": "3.1",
  "applicationId": "iis-appB-ab800cde",
  "imageName": "iis-appB-ab800cde",
  "exposedPorts": [
    {
      "localPort": 8080,
      "protocol": "http"
    }
  ],
  "environment": [],
  "ecrParameters": {
    "ecrRepoTag": "latest"
  },
  "ecsParameters": {
    "createEcsArtifacts": true,
    "ecsFamily": "iis-appB-ab800cde",
    "cpu": 2,
    "memory": 4096,
    "dockerSecurityOption": "",
    "enableCloudwatchLogging": false,
    "publicApp": true,
    "stackName": "iis-appB-ab800cde-ECS",
    "resourceTags": [
      {
        "key": "example-key",
        "value": "example-value"
      }
    ],
    "reuseResources": {
      "vpcId": "vpc-0abc1defa2345b67c",
      "reuseExistingA2cStack": {
        "cfnStackName": "",
        "microserviceUrlPath": ""
      },
      "sshKeyName": "",
      "acmCertificateArn": ""
    },
    "gMSAParameters": {
      "createGMSA": false,
```

```

        "domainSecretsArn": "",
        "domainDNSName": "",
        "domainNetBIOSName": "",
        "gMSAName": ""
    },
    "deployTarget": "ec2",
    "dependentApps" : []
},
"eksParameters": {
    "createEksArtifacts": false,
    "stackName": "",
    "reuseResources": {
        "vpcId": "",
        "reuseExistingA2cStack": {
            "cfnStackName": "",
            "microserviceUrlPath": ""
        },
        "sshKeyPairName": ""
    },
    "gMSAParameters": {
        "createGMSA": false,
        "domainSecretsArn": "",
        "domainDNSName": "",
        "domainNetBIOSName": "",
        "gMSAAccountName": ""
    },
    "dependentApps" : []
}
}

```

• Dependent application C

```

{
    "a2CTemplateVersion": "3.1",
    "applicationId": "service-appC-9fghi90j",
    "imageName": "service-appC-9fghi90j",
    "exposedPorts": [
        {
            "localPort": 8080,
            "protocol": "http"
        }
    ],
    "environment": [],
    "ecrParameters": {

```

```

        "ecrRepoTag": "latest"
    },
    "ecsParameters": {
        "createEcsArtifacts": true,
        "ecsFamily": "service-appC-9fghi90j",
        "cpu": 2,
        "memory": 4096,
        "dockerSecurityOption": "",
        "enableCloudwatchLogging": false,
        "publicApp": true,
        "stackName": "service-appC-9fghi90j-ECS",
        "resourceTags": [
            {
                "key": "example-key",
                "value": "example-value"
            }
        ],
        "reuseResources": {
            "vpcId": "vpc-0abc1defa2345b67c",
            "reuseExistingA2cStack": {
                "cfnStackName": "",
                "microserviceUrlPath": ""
            },
            "sshKeyName": "",
            "acmCertificateArn": ""
        },
        "gMSAParameters": {
            "createGMSA": false,
            "domainSecretsArn": "",
            "domainDNSName": "",
            "domainNetBIOSName": "",
            "gMSAName": ""
        },
        "deployTarget": "ec2",
        "dependentApps" : []
    },
    "eksParameters": {
        "createEksArtifacts": false,
        "stackName": "service-appC-9fghi90j",
        "reuseResources": {
            "vpcId": "",
            "reuseExistingA2cStack": {
                "cfnStackName": "",
                "microserviceUrlPath": ""
            }
        }
    }
}

```

```

        },
        "sshKeyPairName": ""
    },
    "gMSAParameters": {
        "createGMSA": false,
        "domainSecretsArn": "",
        "domainDNSName": "",
        "domainNetBIOSName": "",
        "gMSAAccountName": ""
    },
    "dependentApps" : []
}

```

Configuring container pipelines

This topic contains information about the files that you use to configure continuous integration and deployment (CI/CD) pipelines for your application container with CodePipeline, Jenkins, or Microsoft Azure DevOps.

Pipeline configuration files

- [pipeline.json file](#)

pipeline.json file

When you run the [generate app-deployment](#) command, App2Container creates a `pipeline.json` file for the application that the `--application-id` parameter specifies. The **generate pipeline** command uses this file, along with others, to generate pipeline deployment artifacts. Before you run the **generate pipeline** command, you can configure any of the fields in this file to customize your application container pipeline.

Important

The `pipeline.json` file includes sections for all of the types of pipelines that you can configure. This includes CodePipeline, Jenkins, and Microsoft Azure DevOps. Configure exactly one source repository, and one type of pipeline. In each section, set one Boolean value `enabled` flag to `true`, and all others to `false`. For Jenkins pipelines, you can choose to use either a CodeCommit repository, or an existing Git repository.

CodePipeline

- sourceInfo
 - CodeCommit – enabled: **true**
 - ExistingGitRepo – enabled: **false**
 - AzureRepo – enabled: **false**
- pipelineInfo
 - CodePipeline – enabled: **true**
 - Jenkins – enabled: **false**
 - AzureDevOps – enabled: **false**

Jenkins

- sourceInfo
 - CodeCommit – enabled: **false**
 - ExistingGitRepo – enabled: **true**
 - AzureRepo – enabled: **false**
- pipelineInfo
 - CodePipeline – enabled: **false**
 - Jenkins – enabled: **true**
 - AzureDevOps – enabled: **false**

Microsoft Azure DevOps

- sourceInfo
 - CodeCommit – enabled: **false**
 - ExistingGitRepo – enabled: **false**
 - AzureRepo – enabled: **true**
- pipelineInfo
 - CodePipeline – enabled: **false**
 - Jenkins – enabled: **false**

App2Container enables CodeCommit as the source repository, and CodePipeline as the pipeline by default.

The application `pipeline.json` file includes the following content. While all fields are configurable, the `a2CTemplateVersion` field should not be changed. For key/value pairs that do not apply to your pipeline, set string values to an empty string, numeric values to zero, and Boolean values to false.

- **imageInfo** (object) – Contains parameters needed for Amazon ECR configuration.
 - **image** (string, required) – The full repository name of the application container image to store in Amazon ECR. *Must be in the format <application ID>.<repository name>:<tag>.*
- **sourceInfo** (object) – Contains JSON objects for pipeline source repository configuration for CodePipeline or Jenkins pipelines. CodePipeline uses CodeCommit for its source repository, while Jenkins uses Git.
 - **CodeCommit** (object) – Contains parameters needed for AWS CodeCommit configuration.
 - **enabled** (Boolean, required) – A flag that indicates if you are targeting CodeCommit as the source repository for your pipeline.
 - **repositoryName** (string, required) – The name of the CodeCommit repository to use or create.
 - **branch** (string, required) – The name of the code branch in the CodeCommit repository to commit to.
 - **ExistingGitRepo** (object) – Contains parameters needed for Git repository configuration.
 - **enabled** (Boolean, required) – A flag that indicates if you are targeting Git as the source repository for your pipeline.
 - **repositoryUri** (string, required) – The URI of the Git repository to use for your pipeline. SSH access is required.
 - **branch** (string, required) – The name of the code branch in the Git repository to commit to.
 - **sshKeyArn** (string, required) – The ARN of the secret in Secrets Manager that is used to store the user name and SSH key for Git authentication from the Jenkins server.
 - **AzureRepo** (object) – Contains parameters to specify the Azure Repos Git repository where App2Container uploads pipeline artifacts for your application.

- **enabled** (Boolean, required) – A flag that indicates if you want to use an Azure Repos Git repository as the source repository for an Azure DevOps pipeline that you create.
- **repositoryName** (string, required) – The name of the Azure Repos Git repository that you want to use or create.
- **branch** (string, required) – The name of the code branch in the Azure Repos Git repository where App2Container commits pipeline resources.
- **releaseInfo** (object) – Contains JSON objects with parameters needed to create a pipeline for your target deployment environments.
- **ECS | EKS | AppRunner** (object) – Contains JSON objects representing the environments to target for deployment. The key name specifies the container management service that you are targeting for your application container pipeline. *Key must be "ECS", "EKS", or "AppRunner". At least one of the pipeline environments must be enabled.*
- **beta** (object) –
 - **clusterName** (string, required*) – The name of the Amazon ECS or Amazon EKS cluster to set up in the AWS CloudFormation stack.
 - **serviceName** (string, required*) – The name of the Amazon ECS service to set up in the AWS CloudFormation stack.

** Applies only to Amazon ECS pipelines.*

- **enabled** (Boolean, required) – A flag indicating whether a beta environment should be configured.

 Note

Beta environments are not supported for App Runner.

- **prod** (object) –
 - **clusterName** (string, required*) – The name of the Amazon ECS or Amazon EKS cluster to set up in the AWS CloudFormation stack.

** Does not apply to App Runner.*

 - **serviceName** (string, required*) – The name of the Amazon ECS service to set up in the AWS CloudFormation stack.
- * Applies only to Amazon ECS pipelines.*

- **enabled** (Boolean, required) – A flag indicating whether a prod environment should be configured.
- **pipelineInfo** (object) – Contains JSON objects with parameters needed to access and configure your target pipeline environments.
- **CodePipeline** (object) – Contains parameters needed for CodePipeline configuration.
 - **enabled** (Boolean, required) – A flag that indicates if you are targeting CodePipeline for your pipeline.
- **Jenkins** (object) – Contains parameters needed for Jenkins pipeline access and configuration.
 - **enabled** (Boolean, required) – A flag that indicates if you are targeting Jenkins for your pipeline.
 - **jenkinsServerUrl** (string, required) – The URL of the Jenkins server. The URL requires HTTPS protocol for secure access.
 - **nodeLabels** (array of strings, required) – A list of the labels that must be attached to the Jenkins agent node that runs the pipeline. All labels specified must be present on the agent node for it to run.
 - **apiTokenArn** (string, required) – The ARN of the secret in Secrets Manager that is used to authenticate to the Jenkins server.
 - **repoSshCredentialId** (string, required) – The ID that you create on the Jenkins server that the Jenkins agent node uses for SSH access to the Git repository. For more information about SSH credentials on Jenkins, see the [Using credentials](#) chapter in the *Jenkins User Handbook*, available online..
 - **awsCredentialId** (string, required) – The AWS profile on the Jenkins server that is used to access AWS resources from the Jenkins agent node when the pipeline runs.
- **AzureDevOps** (object) – Contains parameters that you need to access and configure your Azure DevOps pipeline.
 - **enabled** (Boolean, required) – A flag that indicates if you want App2Container to use Azure DevOps to set up your CI/CD pipeline.
 - **organizationName** (string, required) – The name of the organization that you set up under your Microsoft Azure account for Azure DevOps.
 - **projectName** (string, required) – The name of the project that you set up under your Microsoft Azure account for Azure DevOps.
 - **serviceCredName** (string, required) – The name of the service credentials that Azure DevOps uses to connect to AWS.

- **agentPoolName** (string, required) – The name of the agent pool with the Microsoft-hosted agents that your pipeline uses to build and deploy updated container images for your application.
- **personalAccessTokenARN** (string, required) – The ARN that identifies the Secrets Manager secret where you store your Microsoft Azure Personal Access Token (PAT).

Examples

The following example shows a `pipeline.json` file that uses the CodePipeline environment as the pipeline for an IIS application that runs on Windows. The application runs in a beta environment, and there is no prod environment configured yet.

```
{
  "a2CTemplateVersion": "3.1",
  "imageInfo": {
    "image": "123456789012.dkr.ecr.us-west-1.amazonaws.com/iis-smarts-51d2dbf8:latest"
  },
  "sourceInfo": {
    "CodeCommit": {
      "repositoryName": "app2container-iis-smarts-51d2dbf8-ecs",
      "branch": "master"
    }
  },
  "releaseInfo": {
    "ECS": {
      "beta": {
        "clusterName": "a2c-iis-smarts-51d2dbf8-ECS-Cluster",
        "serviceName": "a2c-iis-smarts-51d2dbf8-ECS-LBWebAppStack-1EB23FI45ZYXW-Service-1mnoPQRS2Tu3",
        "enabled": true
      },
      "prod": {
        "clusterName": "",
        "serviceName": "",
        "enabled": false
      }
    }
  }
}
```

The following example shows a pipeline .json file that uses the Jenkins environment as the pipeline for an IIS application that runs on Windows.

```
{
  "a2CTemplateVersion": "1.0",
  "imageInfo": {
    "image": "123456789012.dkr.ecr.us-west-1.amazonaws.com/iis-smarts-51d2dbf8:latest"
  },
  "sourceInfo": {
    "CodeCommit": {
      "enabled": false,
      "repositoryName": "",
      "branch": ""
    },
    "ExistingGitRepo": {
      "enabled": true,
      "repositoryUri": "git@ec2-12-34-567-890.us-west-1.compute.amazonaws.com:/~/windows.git",
      "branch": "master",
      "sshKeyArn": "arn:aws:secretsmanager:us-east-1:123456789075:secret:test-We6XCm"
    }
  },
  "releaseInfo": {
    "ECS": {
      "beta": {
        "clusterName": "a2c-iis-smarts-51d2dbf8-ECS-Cluster",
        "serviceName": "a2c-iis-smarts-51d2dbf8-ECS-LBWebAppStack-1EB23FI45ZYXW-Service-1mnoPQRS2Tu3",
        "enabled": true
      },
      "prod": {
        "clusterName": "",
        "serviceName": "",
        "enabled": false
      }
    }
  },
  "resourceTags": [
    {
      "key": "example-key",
      "value": "example-value"
    }
  ]
}
```

```

    }
  ],
  "pipelineInfo": {
    "CodePipeline": {
      "enabled": false
    },
    "Jenkins": {
      "enabled": true,
      "jenkinsServerUrl": "https://ec2-3-101-121-107.us-
west-1.compute.amazonaws.com",
      "nodeLabels": [
        "windows2019",
        "beta"
      ],
      "apiTokenArn": "arn:aws:secretsmanager:us-east-1:123456789076:secret:test-
We6XCm",
      "repoSshCredentialId": "12345678-90a1-23bc-de45-f67a123bc45d",
      "awsCredentialId": "beta-tester"
    }
  }
}

```

The following example shows a `pipeline.json` file that uses Microsoft Azure DevOps as the pipeline for a Java application that runs on Linux.

```

{
  "a2CTemplateVersion": "1.0",
  "imageInfo": {
    "image": "459632601910.dkr.ecr.us-west-1.amazonaws.com/java-tomcat-9e8e4799:latest"
  },
  "sourceInfo": {
    "CodeCommit": {
      "enabled": false,
      "repositoryName": "a2c-java-tomcat-9e8e4799-ecs",
      "branch": "master"
    },
    "ExistingGitRepo": {
      "enabled": false,
      "repositoryUri": "",
      "branch": "",
      "sshKeyArn": ""
    },
    "AzureRepo": {

```

```

    "enabled": true,
    "repositoryName": "a2c-java-tomcat-9e8e4799",
    "branch": "main"
  }
},
"releaseInfo": {
  "ECS": {
    "beta": {
      "clusterName": "a2c-java-tomcat-9e8e4799-ECS-Cluster",
      "serviceName": "a2c-java-tomcat-9e8e4799-ECS-JavaStack-1AB23CD45ZYXW-
Service-1abcPQRS2Tu3",
      "enabled": true
    },
    "prod": {
      "clusterName": "",
      "serviceName": "",
      "enabled": false
    }
  }
},
"resourceTags": [{
  "key": "example-key",
  "value": "example-value"
}],
"pipelineInfo": {
  "CodePipeline": {
    "enabled": false
  },
  "Jenkins": {
    "enabled": false,
    "jenkinsServerUrl": "",

    "nodeLabels": [],
    "apiTokenArn": "",
    "repoSshCredentialId": "",
    "awsCredentialId": ""
  },
  "AzureDevOps": {
    "enabled": true,
    "organizationName": "App2Container",
    "projectName": "a2c-java-tomcat-9e8e4799-project",
    "serviceCredName": "azure-devops-to-aws-creds",
    "agentPoolName": "Azure Pipelines",

```



```
    "personalAccessTokenARN": "arn:aws:secretsmanager:us-  
east-1:12345678:secret:APP2CONTAINER-PAT"  
  }  
}  
}
```

Product and service integrations for AWS App2Container

AWS App2Container integrates with an array of AWS services, and partner products and services. After you've deployed your application containers to run on Amazon ECS, Amazon EKS, or App Runner, you can use App2Container to choose from several different continuous integration and delivery (CI/CD) platforms to keep your images up to date. Use the information in the following sections to help you configure App2Container to integrate with the products and services that you use.

Contents

- [Automatic storage and registration using Amazon Elastic Container Registry](#)
- [Deploy application containers to Amazon Elastic Container Service with AWS App2Container](#)
- [Deploy application containers to Amazon EKS with AWS App2Container](#)
- [Deploy application containers to AWS App Runner with AWS App2Container](#)
- [Set up CI/CD pipelines with AWS CodePipeline](#)
- [Set up CI/CD pipelines with Jenkins](#)
- [Set up CI/CD pipelines with Microsoft Azure DevOps](#)
- [Setting up FireLens log file routing for containers with AWS App2Container](#)

Automatic storage and registration using Amazon Elastic Container Registry

App2Container uses the Amazon Elastic Container Registry (Amazon ECR) service to register and store container images for all of the environments it supports for application container deployment. When you run the **app2container generate app-deployment** command, App2Container creates an ECR repository and registers your application container image. The ECR repository name is the application ID that App2Container creates when you run the **app2container inventory** command on your application server or worker machine.

Amazon ECR includes the following features, which are not enabled by default when App2Container creates your repository and registers your container image.

- Lifecycle policies that help you manage the lifecycle of your images, and clean up unused images. For more information, see [Lifecycle policies](#) in the *Amazon Elastic Container Registry User Guide*.
- Image scanning that helps to identify software vulnerabilities in your container images. You can configure scan on push validation for your images. You can also run a manual scan on any of your images that are stored in Amazon ECR. For more information, see [Image scanning](#) in the *Amazon Elastic Container Registry User Guide*.
- Cross-Region and cross-account replication to help you distribute your container image to destination accounts and Regions. For more information about replication settings for your registry, see [Private image replication](#) in the *Amazon Elastic Container Registry User Guide*.

To view your ECR repository, and change settings using the AWS Management Console, follow these steps:

1. Open the Amazon ECR console at <https://console.aws.amazon.com/ecr/>.

Verify that the console is showing the Region where you want to view and change settings for your repository. The current Region is displayed in the upper right corner of the console.

2. Select the option next to the **Repository name**, where the name matches your App2Container application ID.

 Tip

You can use any part of the application ID in the search bar to filter your results.

3. Choose **Edit** to view and change the settings for your repository.
4. Choose **Save** to save settings that you have changed, or **Cancel** to exit without saving.

To learn more about Amazon ECR, see [What is Amazon Elastic Container Registry?](#) in the *Amazon Elastic Container Registry User Guide*.

Deploy application containers to Amazon Elastic Container Service with AWS App2Container

Amazon Elastic Container Service (Amazon ECS) is a fully managed container orchestration service that helps you to deploy, manage, and scale containerized applications. It provides a secure

solution for running container workloads with high availability across multiple Availability Zones within a Region. Amazon ECS offers a variety of hosting options for your container environment. For more information about Amazon ECS, see [What is Amazon Elastic Container Service?](#) in the *Amazon Elastic Container Service Developer Guide*.

AWS App2Container integrates with Amazon ECS, to deploy your application containers to the following Amazon ECS environments:

- **Amazon ECS** – In the default environment, your containers run on EC2 instances. App2Container supports Windows .NET application containers for this environment. Linux is not currently supported for this environment.
- **AWS Fargate** – Fargate is a serverless architecture. App2Container supports both Linux and Windows application containers for this environment. To learn more about Fargate, see [Amazon ECS on AWS Fargate](#) in the *Amazon Elastic Container Service Developer Guide*.

Tip

To containerize your applications with a console-based experience and deploy them on Amazon ECS on AWS Fargate, you can use the *Replatform applications to Amazon ECS* template on the [AWS Migration Hub Orchestrator console](#). For more information, see [Replatform applications to Amazon ECS](#) in the *AWS Migration Hub Orchestrator User Guide*.

Prerequisites

To configure an Amazon ECS integration for your application container with App2Container, your application must meet the following criteria.

Amazon ECS

- For deployment to the Amazon ECS default environment, App2Container supports .NET applications running on Windows. [*Linux applications are not currently supported.*]
- .NET applications running on Windows must satisfy application framework and system requirements, and meet the criteria for supported applications. For details, see [Supported applications](#), and expand the **Supported applications for Windows** section.

Fargate

- For deployment to Fargate, App2Container supports the following types of applications:
 - Java applications running on Linux.
 - .NET applications running on Windows Server 2019.
- Java applications running on Linux must satisfy Java application framework requirements, and run on a supported Linux distribution. For details, see [Supported applications](#), and expand the **Supported applications for Linux** section.
- For .NET application containers, the container operating system must be Windows Server 2019. Prior versions are not supported for deployment to Fargate. The container operating system is derived from the application server or worker machine where containerization runs, so the applicable server operating system must also be Windows Server 2019.

Additionally, .NET applications running on Windows must satisfy application framework requirements, and meet the criteria for supported applications. For details, see [Supported applications](#), and expand the **Supported applications for Windows** section.

- gMSA is not supported.

Amazon ECS integration for App2Container workflow

To set up application containers for hosting in Amazon ECS within the App2Container workflow, follow these steps:

Initial steps for App2Container are the same for all applications deploying to Amazon ECS:

1. Install and set up the App2Container environment, as described in the [Prerequisites: Set up your servers](#) section.
2. Complete the initialization phase for your App2Container environment with the **init** command, and the **remote configure** command, if applicable. To learn more about what is included in all of the App2Container containerization phases, see the [Command reference](#).
3. Complete the analyze phase for each application that you want to containerize.
 - If you are running commands directly on application servers, use the **inventory** and **analyze** commands.
 - If you are running a remote workflow on a worker machine, use the **remote inventory** and **remote analyze** commands.

4. Integration begins with the containerization step.

- When you run the **containerize** command, App2Container generates the `deployment.json` file, which provides configurable parameters for all supported container management service options that could apply to your application container.
- Parameters for Amazon ECS and Amazon EKS are always included. Parameters for App Runner are also included if your application container meets the App2Container criteria for hosting in App Runner.
- Each container management service has its own section in the `deployment.json` file, and each section has a flag to indicate which container management service is the destination for your application container. Only one section can have its flag set to **true** – all others must be set to **false**.

Amazon ECS is configured by default as the container management service for your application. However, the destination settings differ, depending on system requirements and the type of application you have.

In the `deployment.json` file, App2Container initially sets the `deployTarget` parameter as follows:

- `ec2` – App2Container targets the Amazon ECS default environment, which runs containers on EC2 instances, for .NET applications that do not meet the criteria specified in the **Fargate** section under [Prerequisites](#). Java applications are not currently supported for this deployment target.
- `fargate` – App2Container targets the Fargate environment by default for Java applications, and for .NET applications that meet the criteria specified in the **Fargate** section under [Prerequisites](#).

If you want your container to run on EC2 instances instead of running in Fargate, you can change the `deployTarget` parameter to `ec2`. However, this is currently only true for .NET applications. If you change the value for a Java application, the **generate app-deployment** command throws an error when you run it.

For more information about configuring the `deployment.json` file, see [Configure deployment](#).

 Note

The `gMSAParameters` are not valid for deployments to Fargate, and will generate an error when the **generate app-deployment** command runs.

5. The deployment step generates an ECS task definition and `pipeline.json` file that are targeted for the Amazon ECS container management service, based on the settings in the `deployment.json` file, where the `createEcsArtifacts` flag is set to **true**.
 - When you run the **generate app-deployment** command, App2Container validates the properties in the `deployment.json` file, and pushes the container image to Amazon ECR. This is the standard workflow.
 - The command generates a CloudFormation template for Amazon ECS deployment (`ecs-master.yml`) that contains the IAM role that Amazon ECS uses to pull your application container images from Amazon ECR, and the Amazon ECS service definition.
 - The command generates the `pipeline.json` file to support creating a pipeline to deploy updates to your application container in Amazon ECR.
 - If you use the `--deploy` option for the **generate app-deployment** command, App2Container deploys the CloudFormation stack that creates the Amazon ECS service for the containerized application, using the configuration values in the CloudFormation template that it generates. To customize the configuration, run the command without the `--deploy` option, and then manually deploy using the AWS CLI when you are ready.
6. The pipeline step generates a CloudFormation template for the pipeline that is targeted for the Amazon ECS container management service, based on the settings in the `pipeline.json` file.
 - When you run the **generate pipeline** command, App2Container validates the properties in the `pipeline.json` file, verifies that initial deployment to Amazon ECS has been completed, and verifies that your application is active.
 - The command generates a CloudFormation template to create a two-step pipeline:
 1. **Code commit** – Creates or updates an AWS CodeCommit repository that contains the Dockerfile and application artifacts that are required to create your application container image.
 2. **Code build** – Builds the Docker image for your application container, and pushes the updated image to the Amazon ECR repository that you configured for your application.

3. If you use the `--deploy` option for the **generate pipeline** command, App2Container deploys the pipeline with the configuration values in the CloudFormation template it generates. To customize the configuration, run the command without the `--deploy` option, and then manually deploy using the AWS CLI when you are ready.

Deploy application containers to Amazon EKS with AWS App2Container

Amazon Elastic Kubernetes Service (Amazon EKS) is a managed service that you can use to run Kubernetes on AWS. Amazon EKS streamlines the provisioning of highly available and secure clusters, and automates key maintenance tasks such as patching, node provisioning, and updates. Kubernetes is an open-source system for automating the deployment, scaling, and management of containerized applications. For more information about Amazon EKS, see [What is Amazon EKS?](#) in the [Amazon EKS User Guide](#).

Prerequisites

To configure an Amazon EKS integration for your application container with App2Container, your application must meet the following criteria.

- Java applications running on Linux must satisfy Java application framework requirements, and run on a supported Linux distribution. For details, see [Supported applications](#), and expand the **Supported applications for Linux** section.
- .NET applications running on Windows must satisfy application framework and system requirements, and meet the criteria for supported applications. For details, see [Supported applications](#), and expand the **Supported applications for Windows** section.
- Application containers that run in Amazon EKS must launch EC2 instances. App2Container does not currently support Fargate as a container launch type for Amazon EKS.

Amazon EKS integration for App2Container workflow

The process for setting up application containers for hosting in Amazon EKS is integrated smoothly with the App2Container workflow. Initial steps for App2Container are the same for all applications.

1. Install and set up the App2Container environment, as described in the [Prerequisites: Set up your servers](#) section.

2. Complete the initialization phase for your App2Container environment with the **init** command, and the **remote configure** command, if applicable. To learn more about what is included in all of the App2Container containerization phases, see the [Command reference](#).
3. Complete the analyze phase for each application that you want to containerize.
 - If you are running commands directly on application servers, use the **inventory** and **analyze** commands.
 - If you are running a remote workflow on a worker machine, use the **remote inventory** and **remote analyze** commands.
4. Integration begins with the containerization step.
 - When you run the **containerize** command, App2Container generates the `deployment.json` file, which provides configurable parameters for all supported container management service options that could apply to your application container.
 - Parameters for Amazon ECS and Amazon EKS are always included. Parameters for App Runner are also included if your application container meets the App2Container criteria for hosting in App Runner
 - Each container management service has its own section in the `deployment.json` file, and each section has a flag to indicate which container management service is the destination for your application container. Only one section can have its flag set to **true** – all others must be set to **false**.

Amazon ECS is configured as the destination by default. To deploy your application containers to Amazon EKS, you can set the `createEksArtifacts` in the `eksParameters` section to **true**, and the `createEcsArtifacts` flag in the `ecsParameters` section to **false**. For more information about configuring the `deployment.json` file, see [Configure deployment](#).

- App2Container configures HTTP-based deployments by default. To use HTTPS for your deployment, specify the Amazon Resource Name (ARN) of an AWS Certificate Manager (ACM) certificate in the `deployment.json` file. For more information, see [Configure deployment](#).
5. The deployment step creates artifacts that are targeted for the Amazon EKS container hosting service, based on the settings in the `deployment.json` file, where the `createEksArtifacts` flag is set to **true**.

- When you run the **generate app-deployment** command, App2Container validates the properties in the deployment.json file, and pushes the container image to Amazon ECR. This is the standard workflow.
- The command generates a CloudFormation template (eks-master.yml) that creates an EKS cluster, pulls your application container images from Amazon ECR, and deploys your application to the cluster.

It also generates Kubernetes manifests (eks_deployment.yaml, eks_service.yaml, and eks_ingress.yaml), for post-deployment customizations using a tool such as kubectl.

- The command generates the pipeline.json file to support creating a pipeline to deploy updates to your application container in Amazon ECR.
 - If you use the --deploy option for the **generate app-deployment** command, App2Container deploys the AWS CloudFormation stack that creates the Amazon EKS service for the containerized application, using the configuration values in the AWS CloudFormation template that it generates. To customize the configuration, run the command without the --deploy option, and then manually deploy using the AWS CLI when you are ready.
6. The pipeline step generates a CloudFormation template for the pipeline that is targeted for the Amazon EKS container management service, based on the settings in the pipeline.json file.
- When you run the **generate pipeline** command, App2Container validates the properties in the pipeline.json file, and verifies that initial deployment to Amazon EKS has been completed, and that your application is active.
 - The command generates a CloudFormation template to create a two-step pipeline:
 1. **Code commit** – Creates or updates an AWS CodeCommit repository that contains the Dockerfile and application artifacts that are required to create your application container image.
 2. **Code build** – Builds the Docker image for your application container, and pushes the updated image to the Amazon ECR repository that you configured for your application.
 3. If you use the --deploy option for the **generate pipeline** command, App2Container deploys the pipeline with the configuration values in the CloudFormation template it generates. To customize the configuration, run the command without the --deploy option, and then manually deploy using the AWS CLI when you are ready.

Deploy application containers to AWS App Runner with AWS App2Container

AWS App Runner is an AWS service that provides a way for existing container images or source code to run directly as web services in AWS. App Runner uses Fargate as its underlying environment, but has its own management layer on top. With App Runner, you can access your application through an assigned web service URL, via HTTP requests.

Considerations for deploying to App Runner using App2Container:

- App Runner is not available in all Regions. To see the Regions and service endpoints for App Runner, refer to [App Runner Service endpoints](#) in the *AWS General Reference*.
- Resources that are created by App Runner reside in the multi-tenant App Runner service account. With other container management services, you might access resources such as an Amazon EC2 instance that your container runs on, or an Amazon EBS volume attached to your container instance, using the standard access methods for those resources directly. With App Runner you access resources that App Runner creates for your application through the App Runner service, using the App Runner console, API, SDKs, or by using **apprunner** commands in the AWS CLI.
- App Runner supports continuous integration and deployment from the Amazon ECR repository that App2Container creates on your behalf. When continuous deployment is configured, an update to the container image in the Amazon ECR repository automatically initiates an update in App Runner.

You can turn this on or off in the `deployment.json` file. For more information, see [Configure deployment](#).

- App Runner integrates with Amazon CloudWatch and AWS CloudTrail to provide logging and monitoring support for your application. App Runner creates the following log groups for each App Runner service:
 - An application group, which contains `stdout` from your containers.
 - A service group, which contains high-level logs from App Runner to notify you about service-related events, such as new deployments or health check failures.

These logs can also be viewed from the App Runner console, or by using the App Runner API, SDKs, or by using **apprunner** commands in the AWS CLI.

- App Runner enforces limits for the application containers that it hosts, such as the number of concurrent requests, the size of the application, and the amount of memory it can use. To learn

more about Service Quotas for App Runner, see [App Runner Service quotas](#) in the *AWS General Reference*.

- Application state is not guaranteed to be maintained between requests.

For more information about using App Runner to host your application container, see [What is AWS App Runner](#) in the *AWS App Runner Developer Guide*.

Prerequisites

To configure an App Runner integration for your application container with App2Container, your application must meet the following criteria:

- Your application runs on Linux. [*Windows applications are not currently supported.*]
- Your application meets all of the requirements that are listed in the [Supported applications](#) section for Linux.
- Your application container size is less than 3 GB.
- Your application must not be dependent on background processing. App Runner heavily throttles container CPU when requests are not actively being processed.

App Runner integration for App2Container workflow

Setting up application containers for hosting in App Runner integrates smoothly with the App2Container workflow. Initial steps for App2Container are the same for all applications:

1. Install and set up the App2Container environment, as described in the [Prerequisites: Set up your servers](#) section.
2. Complete the initialization phase for your App2Container environment with the **init** command, and the **remote configure** command, if applicable. To learn more about what is included in all of the App2Container containerization phases, see the [Command reference](#).
3. Complete the analyze phase for each application that you want to containerize.
 - If you are running commands directly on application servers, use the **inventory** and **analyze** commands.
 - If you are running a remote workflow on a worker machine, use the **remote inventory** and **remote analyze** commands.
4. Integration begins with the containerization step.

- When you run the **containerize** command, App2Container generates the `deployment.json` file, which provides configurable parameters for all supported container management service options that could apply to your application container.
- Parameters for Amazon ECS and Amazon EKS are always included. Parameters for App Runner are also included if your application container meets the App2Container criteria for hosting in App Runner (see [Prerequisites](#)).
- Each container management service has its own section in the `deployment.json` file, and each section has a flag to indicate which container management service is the destination for your application container. Only one section can have its flag set to **true** – all others must be set to **false**.

Amazon ECS is configured as the destination by default, but if your application is suitable for App Runner, you can set the `createAppRunnerArtifacts` flag in the `appRunnerParameters` section to **true**, and the `createEcsArtifacts` flag in the `ecsParameters` section to **false**. For more information about configuring the `deployment.json` file, see [Configure deployment](#).

5. The deployment step generates a CloudFormation template and `pipeline.json` file that are targeted for the App Runner container management service, based on the settings in the `deployment.json` file, where the `createAppRunnerArtifacts` flag is set to **true**.
- When you run the **generate app-deployment** command, App2Container validates the properties in the `deployment.json` file, and pushes the container image to Amazon ECR. This is the standard workflow.
 - The command generates a CloudFormation template for App Runner deployment that contains the IAM role that App Runner uses to pull your application container images from Amazon ECR, and the App Runner service definition.
 - The command generates the `pipeline.json` file to support creating a pipeline to deploy updates to your application container in Amazon ECR.
 - If you use the `--deploy` option for the **generate app-deployment** command, App2Container deploys the AWS CloudFormation stack that creates the App Runner service for the containerized application, using the configuration values in the AWS CloudFormation template that it generates. To customize the configuration, run the command without the `--deploy` option, and then manually deploy using the AWS CLI when you are ready.

6. The pipeline step generates a CloudFormation template for the pipeline that is targeted for the App Runner container management service, based on the settings in the `pipeline.json` file.
 - When you run the **generate pipeline** command, App2Container validates the properties in the `pipeline.json` file, and verifies that initial deployment to App Runner has been completed, and that your application is active.
 - The command generates a CloudFormation template to create a two-step pipeline:
 1. **Code commit** – Creates or updates an AWS CodeCommit repository that contains the Dockerfile and application artifacts that are required to create your application container image.
 2. **Code build** – Builds the Docker image for your application container, and pushes the updated image to the Amazon ECR repository that you configured for your application.
 3. If you use the `--deploy` option for the **generate pipeline** command, App2Container deploys the pipeline with the configuration values in the CloudFormation template it generates. To customize the configuration, run the command without the `--deploy` option, and then manually deploy using the AWS CLI when you are ready.

Note

If you have automatic deployments configured for App Runner, an update to your application container image in Amazon ECR automatically kicks off an update for your application in App Runner.

To configure automatic deployments, use the following settings in the `deployment.json` file:

- Set `autoDeploymentsEnabled` to **true** to automatically deploy updates to App Runner when you deploy updates to Amazon ECR. *This is the default setting.*
- Set `autoDeploymentsEnabled` to **false** if you want to update App Runner manually, using the App Runner service console, API, SDKs, or AWS CLI.

Set up CI/CD pipelines with AWS CodePipeline

AWS CodePipeline is a continuous delivery service that you can use to model, visualize, and automate your software release process. App2Container integrates with CodePipeline to automate

a consistent release process while it gives you insights to monitor and manage your pipeline. For more information, see [What is AWS CodePipeline?](#) in the *AWS CodePipeline User Guide*.

Before you run the **generate pipeline** command, review the `pipeline.json` file that the **generate app-deployment** command creates. Configure the parameters for your CodeCommit pipeline as follows:

- Set the flags to enable CodePipeline deployment.
 - `sourceInfo`
 - CodeCommit – enabled: **true**
 - ExistingGitRepo – enabled: **false**
 - AzureRepo – enabled: **false**
 - `pipelineInfo`
 - CodePipeline – enabled: **true**
 - Jenkins – enabled: **false**
 - AzureDevOps – enabled: **false**

Important

You must set the `sourceInfo` and `pipelineInfo` flags as described or else the pipeline integration will fail.

Contents

- [Validation](#)
- [Output](#)

Validation

File validation

When you run the `generate pipeline` command, App2Container performs the following validation to ensure that your pipeline deploys successfully:

- Checks that CodeCommit is the only source repository that you've activated in the `sourceInfo` section of the `pipeline.json` file, and that this section contains all required properties.

- Checks that CodePipeline is the only pipeline that you've activated in the `pipelineInfo` section of the `pipeline.json` file, and that this section contains all required properties.

Deployment validation

If you use the `App2Container generate pipeline` command with `--deploy`, the `pipeline.json` file that App2Container creates will have the required configuration already defined. If you don't specify the `--deploy` flag for the command, or you use your own deployment, you must edit the `pipeline.json` file to specify the required configuration. For more information, see [Configuring container pipelines](#).

Output

The `generate pipeline` command generates the following artifacts for CodePipeline pipelines. If you don't use the `--deploy` option with the `generate pipeline` command, you can edit the artifacts that App2Container added to your CodeCommit repository to create your pipeline from the CodePipeline interface (AWS CLI or AWS Management Console).

Note

If you run the `generate pipeline` command with the `--deploy` option, App2Container creates the pipeline in CodePipeline, and starts the pipeline build.

App2Container generates the following artifacts:

buildspec.yml files

Used to build the application container image and uploads it to Amazon ECR.

AWS CloudFormation templates

Used to create your pipeline in CodePipeline along with other required resources.

Note

If your CodeCommit repository doesn't already exist, App2Container creates it for you.

Set up CI/CD pipelines with Jenkins

Jenkins is an open source automation server that which supports building, deploying, and automating your application with the help of Jenkins Pipeline. Jenkins Pipeline is a suite of plugins that supports implementing and integrating continuous delivery pipelines into Jenkins. These plugins can be used to integrate with AWS App2Container to automate deployments for your applications. App2Container can help configure a Jenkins pipeline in your existing Jenkins environment.

For more information about using Jenkins, see the [User Handbook overview](#) on the Jenkins website.

Prerequisites

To configure Jenkins pipeline integration for your application container from App2Container, your application must meet the following criteria.

- A fully functional Jenkins server with the following plugins installed:
 - [Pipeline](#)
 - [Pipeline: AWS Steps](#)
 - [Git](#)
- One or more agent nodes, running Linux or Windows must be configured on the Jenkins server.

Note

The application container platform must match the platform of the agent node. For example, a Java application that runs on Linux, must use a Linux agent node for Jenkins. A .NET application that runs on Windows, must use a Windows agent node.

- Agent nodes must have the following tools installed:
 - **AWS command line tool** – To install the AWS CLI or Tools for Windows PowerShell on the agent nodes, [follow the same steps](#) that you used to set up your application servers and worker machines, except that you do not need to set up an AWS profile on the agent node. Agent nodes use the AWS profile that is configured on the Jenkins server.
 - **Docker** – The Docker engine installation varies by the operating system platform for the server or instance where you install it. For more information about the variations, see [Install the Docker engine](#).

- **Git** – For more information, see the [1.5 Getting Started - Installing Git](#) chapter in the *Pro Git* guide, available free to read online.
- Agent nodes must be able to connect to AWS and run commands using the AWS CLI.
- The Jenkins server must have access to an existing Git repository for pipeline source. The following credentials and resources are required for pipeline builds:
 - Credentials created on the Jenkins server that are used to access the Git repository from the Jenkins agent node through SSH. The ID of the Jenkins credentials is required in `pipeline.json` configuration. For more information about SSH credentials on Jenkins, see the [Using credentials](#) chapter in the *Jenkins User Handbook*, available online.
 - An AWS profile on the Jenkins server that is used to access AWS resources from the Jenkins agent node when the pipeline runs.
- Credentials for App2Container to integrate with Jenkins resources must be created and stored in AWS Secrets Manager. For more information, see [Create secrets for Jenkins pipelines](#)
- The application server or worker machine where the App2Container **generate pipeline** command runs must be able to connect to the Git source repository and Jenkins server, using the secrets stored in Secrets Manager.

For more information about installing and configuring a Jenkins server, see the [Installing Jenkins](#) chapter in the *Jenkins User Handbook*, available online. The [Jenkins User Documentation](#) also includes tutorials and other reference materials.

Jenkins integration for App2Container workflow

The process for setting up Jenkins pipelines to refresh components for your application container integrates smoothly with the App2Container workflow. Applications follow all the standard steps through deployment. Jenkins integration happens in the pipeline step.

1. Before you run the **generate pipeline** command, review the `pipeline.json` file that was created by the **generate app-deployment** command. Configure the parameters for your Jenkins pipeline as follows:
 - Set the flags to enable Jenkins deployment.
 - `sourceInfo`
 - `CodeCommit` – enabled: **false**
 - `ExistingGitRepo` – enabled: **true**

- AzureRepo – enabled: **false**
- pipelineInfo
 - CodePipeline – enabled: **false**
 - Jenkins – enabled: **true**
 - AzureDevOps – enabled: **false**
- In the ExistingGitRepo object, set the following parameters:
 - **repositoryUri** (string, required) – The URI of the Git repository to use for your pipeline. SSH access is required.
 - **branch** (string, required) – The name of the code branch in the Git repository to commit to.
 - **sshKeyArn** (string, required) – The ARN of the secret in Secrets Manager that is used to store the user name and SSH key for Git authentication from the Jenkins server.
- In the pipelineInfo section Jenkins object, set the following parameters:
 - **jenkinsServerUrl** (string, required) – The URL of the Jenkins server. HTTPS is required for secure access.
 - **nodeLabels** (array of strings, required) – A list of the labels that must be attached to the Jenkins agent node that runs the pipeline. All labels specified must be present on the agent node for it to run.
 - **apiTokenArn** (string, required) – The ARN of the secret in Secrets Manager that is used to authenticate to the Jenkins server.
 - **repoSshCredentialId** (string, required) – The ID of the credential that you create on the Jenkins server, which is used to access the Git repository from the Jenkins agent node through SSH. For more information about SSH credentials on Jenkins, see the [Using credentials](#) chapter in the Jenkins *User Handbook*, available online.
 - **awsCredentialId** (string, required) – The AWS profile on the Jenkins server that is used to access AWS resources from the Jenkins agent node when the pipeline runs.
- 2. When you run the **generate pipeline** command, App2Container validates the properties in the `pipeline.json` file, and verifies that initial deployment to your container management service has been completed, and that your application is active.

The **generate pipeline** command generates the following artifacts for Jenkins pipelines:

- **Jenkinsfile** – App2Container uses the Declarative Pipeline syntax to produce the ~~Jenkinsfile. The file contains the steps and stages (code, build, release, etc.) for the~~

Jenkins pipeline. For more information about Jenkins pipeline syntax, see [Pipeline Syntax](#) on the Jenkins website.

If you are not using the `--deploy` option with the **generate pipeline** command, you can customize the Jenkinsfile, and then use it to create your pipeline using the Jenkins user interface.

- **A config.xml file** – If you are not using the `--deploy` option with the **generate pipeline** command, you can use the `config.xml` file, along with the Jenkinsfile to create your pipeline using the Jenkins REST API (JenkinsAPI). For more information, see the online documentation site: [JenkinsAPI](#).
- **Amazon EKS CloudFormation template** (for Amazon EKS deployment only) – If your application is deploying to Amazon EKS, the **generate pipeline** command generates a CloudFormation template to create a two-step pipeline. For more information about Amazon EKS deployments, see [Deploy application containers to Amazon EKS with AWS App2Container](#)

Note

If you are using CodeCommit as your source repository, App2Container creates an SSH key for the IAM user that is running the command. It provides that SSH key to the Jenkins server, so that Jenkins can access files in CodeCommit when it runs the pipeline.

If you run the **generate pipeline** command with the `--deploy` option, App2Container creates the pipeline in Jenkins, and starts the pipeline build.

Set up CI/CD pipelines with Microsoft Azure DevOps

Azure DevOps is a continuous delivery platform, orchestrator, and cloud provider from Microsoft. App2Container integrates with Azure DevOps Services to automate the build and deployment process that updates your application container images in Amazon ECR. For more information about Azure DevOps, see [What is Azure DevOps?](#) in the Microsoft documentation.

Contents

- [Prerequisites](#)

- [Azure DevOps integration for App2Container workflow](#)

These directions cover IAM users. For information that covers IAM Roles Anywhere, see [External pipeline deployments to AWS using IAM Roles Anywhere](#).

Prerequisites

To configure Azure DevOps pipeline integration for your application container from App2Container, your application must meet the following criteria.

- You must have a Microsoft Azure account with the following organization and project structure:
 - An organization that Azure DevOps services can use for your pipeline. To learn more about how to set up an organization for your Microsoft Azure account, see the [Create an organization](#) page on the *Azure DevOps Services* documentation website.
 - A project that Azure DevOps services can use for your pipeline. The project establishes a repository where your pipeline stores artifacts for your application. For more information, see [Create a project in Azure DevOps](#) on the *Azure DevOps Services* documentation website.
 - An agent pool that contains Microsoft-hosted agents. Microsoft provides a predefined agent pool called **Azure Pipelines** that contains Microsoft-hosted agents. When you create your agent pool, choose the **Azure Pipelines** default agent pool. For more information, see [Create and manage agent pools](#) on the *Azure DevOps Services* documentation website.
- To access AWS resources for your application from your Azure DevOps pipeline, install the AWS Toolkit for Azure DevOps extension into your Azure DevOps account .
 - Search for AWS toolkit for Azure DevOps in the [Azure DevOps section of the Visual Studio Marketplace](#).
 - Choose the **AWS toolkit for Azure DevOps** extension from the results.
 - Choose **Get it free** If prompted, sign in to your Azure DevOps account.
 - To install the extension into your Azure DevOps account, choose **Install**.
- Azure DevOps pipelines need permission to perform pipeline actions that access or update AWS resources. To grant access for Azure DevOps, attach or embed the policy resources and actions shown in the **IAM policy for Azure DevOps** example in the [IAM policy examples](#). For more information on how to set up your IAM resources for App2Container, see [Create IAM resources for general use](#).

- After you've installed the AWS Toolkit for Azure DevOps and set up the IAM user and policy that Azure DevOps uses to interact with AWS services, you can set up an AWS service connection under your Azure project settings, as follows:
 1. Sign in to your Azure DevOps account organization, and select your project.
 2. In the lower left of your browser window, choose **Project settings**. This opens the **Project Settings** menu.
 3. In the **Pipelines** section of the menu, choose **Service connections**.
 4. Choose **New service connection**. This displays a list of services that you can connect to.
 5. To open the **New AWS service connection** form, choose **AWS** from the list, and then choose **Next**. If there is a long list of service connections, you might need to scroll down.
 6. Enter the following information in the form:

Required

- **Access Key ID** – The access key ID for the IAM user that Azure DevOps uses to access AWS services for pipeline actions.
- **Secret Access Key** – The secret access key for the IAM user that Azure DevOps uses to access AWS services for pipeline actions.
- **Service connection name** – The name of the service connection for your project
- **Grant access permission to all pipelines** – Select this check box to ensure that all of your pipelines have permission to access AWS services.

You can fill in one or more of the optional fields, if needed, depending on how you set up your security in IAM.

7. Choose **Save** to save your settings and close the form.

For more information, see [Manage service connections](#) on the *Azure DevOps Services* documentation website.

- When App2Container runs Azure DevOps pipelines, it authenticates with a Microsoft Azure Personal Access Token (PAT). To learn more about how to create a PAT and save it as a secret in AWS Secrets Manager, see [Create secrets for Microsoft Azure DevOps pipelines](#).

Azure DevOps integration for App2Container workflow

Applications follow all of the standard App2Container workflow steps through deployment. Azure DevOps integration happens in the pipeline step. To set up integration with Microsoft Azure DevOps pipelines, to refresh components for your application container, configure the `pipeline.json` file as follows.

Before you run the **generate pipeline** command, review the `pipeline.json` file that the **generate app-deployment** command created. Configure the parameters for your Azure DevOps pipeline as follows:

- Set the flags to activate Azure DevOps deployment. Configure exactly one source repository, and one type of pipeline. In each section, set one Boolean value enabled flag to `true`, and all others to `false`.
 - `sourceInfo`
 - `CodeCommit` – enabled: **false**
 - `ExistingGitRepo` – enabled: **false**
 - `AzureRepo` – enabled: **true**
 - `pipelineInfo`
 - `CodePipeline` – enabled: **false**
 - `Jenkins` – enabled: **false**
 - `AzureDevOps` – enabled: **true**
- In the `AzureRepo` object of the `sourceInfo` section, set the following additional parameters, or leave the default values that App2Container creates:
 - **repositoryName** (string, required) – The name of the Azure Repos Git repository that you want to use or create.
 - **branch** (string, required) – The name of the code branch in the Azure Repos Git repository where App2Container commits pipeline resources.
- In the `Azure DevOps` object of the `pipelineInfo` section, set the following additional parameters, or leave the default values that App2Container creates:
 - **organizationName** (string, required) – The name of the organization that you set up under your Microsoft Azure account for Azure DevOps.
 - **projectName** (string, required) – The name of the project that you set up under your Microsoft Azure account for Azure DevOps.

- **serviceCredName** (string, required) – The name of the service credentials that Azure DevOps uses to connect to AWS.
- **agentPoolName** (string, required) – The name of the agent pool with the Microsoft-hosted agents that your pipeline uses to build and deploy updated container images for your application.
- **personalAccessTokenARN** (string, required) – The ARN that identifies the Secrets Manager secret where you store your Microsoft Azure Personal Access Token (PAT).

Validation

When you run the **generate pipeline** command, App2Container performs the following validation to ensure the success of your pipeline deployment:

File validation

App2Container ensures that the Azure DevOps sections in the `pipeline.json` file are complete, and that all required properties pass validation.

- Checks that AzureRepo is the only source repository that you have activated in the `sourceInfo` section of the `pipeline.json` file, and that this section contains all required properties.
- Checks that AzureDevOps is the only pipeline that you have activated in the `pipelineInfo` section of the `pipeline.json` file, and that this section contains all required properties.

Deployment validation

Before creating a pipeline, you must have deployed your containerized application to run on Amazon ECS, Amazon EKS, or App Runner. App2Container verifies that your application container is running in the environment you've configured before it proceeds.

Microsoft-hosted agent validation

App2Container verifies that all of the following prerequisites are installed on the Microsoft-hosted agent:

- Git
- Docker engine
- AWS CLI
- `kubect1` (only for Amazon EKS container pipelines)

Azure account tools and settings

App2Container verifies that the Microsoft Azure account has the tools and settings it needs to interact with AWS for Azure DevOps pipeline deployments, as follows:

- The AWS Toolkit for Microsoft Azure DevOps is installed in the Azure DevOps account
- The Azure DevOps service connection is configured for AWS
- The Microsoft Azure Agent Pool exists

Output

The **generate pipeline** command generates the following artifacts for Azure DevOps pipelines. If you don't use the `--deploy` option with the **generate pipeline** command, you can edit the artifacts that App2Container added to your Azure Repos Git repository to create your pipeline from the Azure DevOps interface.

Amazon ECS

Scripts to install and validate prerequisites on the Microsoft-hosted agent

- **install-pre-req-aws.sh** – Installs AWS CLI on the Microsoft-hosted agent.
- **install-pre-req-docker.sh** – Installs the Docker engine on the Microsoft-hosted agent.
- **install-pre-req-git.sh** – Installs Git on the Microsoft-hosted agent.
- **pre-requisite-validation.sh** – Checks the Microsoft-hosted agent for prerequisites, and installs any that are missing.

Note

Scripts for Windows platforms use the .ps1 file extension.

Pipeline resources (in usage order)

- **pre-requisites.yml** – Sets up a pipeline stage that runs scripts to check the Microsoft-hosted agent and install any prerequisites that are missing.
- **pipeline.json** – Contains configurable settings for your pipeline..
- **image-build.yml** – Builds the application container image and uploads it to Amazon ECR.

- **beta-ecs-release.yaml** – Updates the Amazon ECS clusters for your beta environment, if you have defined that stage.
- **prod-ecs-release.yaml** – Updates the Amazon ECS clusters for your prod environment, if you have defined that stage.

 Note

App2Container supports two stages for your pipelines: beta and prod. You must have at least one stage defined, or you can have both.

Amazon EKS

Scripts to install and validate prerequisites on the Microsoft-hosted agent

- **install-pre-req-aws.sh** – Installs AWS CLI on the Microsoft-hosted agent.
- **install-pre-req-docker.sh** – Installs the Docker engine on the Microsoft-hosted agent.
- **install-pre-req-git.sh** – Installs Git on the Microsoft-hosted agent.
- **install-pre-req-kubectl.sh** – Installs kubectl on the Microsoft-hosted agent.
- **A pre-requisite-validation.sh file** – Checks the Microsoft-hosted agent for prerequisites, and installs any that are missing.

 Note

Scripts for Windows platforms use the .ps1 file extension.

Pipeline resources (in usage order)

- **pre-requisites.yaml** – Sets up a pipeline stage that runs scripts to check the Microsoft-hosted agent and install any prerequisites that are missing.
- **pipeline.json** – Contains configurable settings for your pipeline.
- **image-build.yaml** – Builds the application container image and uploads it to Amazon ECR.
- **beta-eks-release.yaml** – Updates the Amazon EKS clusters for your beta environment, if you have defined that stage.

- **prod-eks-release.yaml** – Updates the Amazon EKS clusters for your prod environment, if you have defined that stage.

 Note

App2Container supports two stages for your pipelines: beta and prod. You must have at least one stage defined, or you can have both.

App Runner

Scripts to install and validate prerequisites on the Microsoft-hosted agent

- **install-pre-req-aws.sh** – Installs AWS CLI on the Microsoft-hosted agent.
- **install-pre-req-docker.sh** – Installs the Docker engine on the Microsoft-hosted agent.
- **install-pre-req-git.sh** – Installs Git on the Microsoft-hosted agent.
- **pre-requisite-validation.sh** – Checks the Microsoft-hosted agent for prerequisites, and installs any that are missing.

 Note

Scripts for Windows platforms use the .ps1 file extension.

Pipeline resources (in usage order)

- **pre-requisites.yaml** – Sets up a pipeline stage that runs scripts to check the Microsoft-hosted agent and install any prerequisites that are missing.
- **pipeline.json** – Contains configurable settings for your pipeline.
- **image-build.yaml** – Builds the application container image and uploads it to Amazon ECR.

 Note

If your Azure Repos Git repository doesn't already exist, App2Container creates it.

If you run the **generate pipeline** command with the `--deploy` option, App2Container creates the pipeline in Azure DevOps, and starts the pipeline build.

Setting up FireLens log file routing for containers with AWS App2Container

When you set up your application containers to use FireLens for Amazon ECS you can route your application logs to CloudWatch, Kinesis Data Streams, or Firehose for log storage and analytics. After you have configured the FireLens settings in your application analysis and deployment JSON files, App2Container creates the artifacts that you need to deploy your application to Amazon EC2 or AWS Fargate. This includes:

- Creation of initial Kinesis Data Streams or Firehose streams, if applicable
- Creation of an IAM role with the permissions needed to enable FireLens log routing to the destinations that you have specified
- Deployment artifacts that contain the FireLens parameters that you specified in your JSON configuration files, including the Amazon ECS task definition and AWS CloudFormation template files

For more information about using FireLens for Amazon ECS, see [Custom log routing](#) in the *Amazon Elastic Container Service Developer Guide*.

Note

App2Container initially supports FireLens log file routing for Amazon ECS for Linux containers only.

Contents

- [FireLens log routing for Linux](#)

FireLens log routing for Linux

Before starting these configuration steps, you should have an understanding of the App2Container containerization phases – Initialize, Analyze, Transform, and Deploy. To learn more about the

containerization phases and the commands that run during each phase, see the [App2Container command reference](#) in this user guide.

Follow these steps to set up log file routing with FireLens for Amazon ECS for your Linux application containers:

FireLens configuration

- [Prerequisites](#)
- [Step 1: Identify log locations for the container](#)
- [Step 2: Configure log deployment parameters](#)
- [Step 3: Validate deployment artifacts](#)
- [Step 4: Deploy your application to Amazon ECS](#)
- [Step 5: Verify log routing](#)

Prerequisites

Prior to setting up FireLens log routing for your application, you must have completed the following prerequisites:

- You have root access on the application server (and worker machine, if using).
- You successfully completed all of the steps from the [Prerequisites: Set up your servers](#) section of this user guide.
- You have initialized the App2Container environment by successfully running the [init](#) command.
- The application must be running on the application server, and must have a valid application ID assigned by the [inventory](#) command.

Step 1: Identify log locations for the container

Run the [analyze](#) command for your application, and then update the following parameters in your `analysis.json` file:

- Update the `logLocations` array to include a list of log files or directory locations where log files can be picked up for routing with FireLens.
- Set the `enableDynamicLogging` parameter to `true` to map application logs to `stdout` as they are created. If your application appends to specific log files such as `info.log` or `error.log`, set the `enableDynamicLogging` parameter to `false`.

The `analysis.json` file is stored in the application folder, for example: `/root/app2container/java-tomcat-9e8e4799`. For more information on `analysis.json` fields and configuration, see [Configuring application containers](#) in the **Configuring your application** section of this user guide.

Example:

The following example shows container parameters in the `analysis.json` file for logging.

```
"containerParameters": {  
  ...  
  "logFiles": ["error.log", "info.log"],  
  "logDirectory": "/var/app/logs/",  
  "logLocations": ["error.log", "info.log", "/var/app/logs/"],  
  "enableDynamicLogging": true,  
  ...  
},
```

Step 2: Configure log deployment parameters

Run the [containerize](#) command, and then edit the `deployment.json` file to set the `fireLensParameters`. The `deployment.json` file is stored in the application folder, for example: `/root/app2container/java-tomcat-9e8e4799`.

There must be at least one valid log destination defined for the `logDestinations` array, with valid values for each of the parameters it contains. For more information on `deployment.json` fields and configuration, including how to target deployment to AWS Fargate with the `deployTarget` parameter, see [Configuring container deployment](#) in the **Configuring your application** section of this user guide.

- Set `enableFirelensLogging` to `true`.
- Configure one or more valid `logDestinations` as follows:
 - **service** – the AWS service to route logs to. *Valid values are "cloudwatch", "firehose", and "kinesis".*
 - **regexFilter** (string) – the pattern to match against log content using a Ruby regular expression to determine where to route the log.

Note

Ruby regular expressions begin and end with a forward slash, with the pattern to match specified in between the slashes. Patterns often begin with a caret (^), which starts matching at the beginning of the line, and end with a dollar sign (\$), which stops matching at the end of the line.

The `regexFilter` parameter in the `deployment.json` file represents only the matching pattern. Be sure to test your matching pattern using one of the many applications available for your desktop or online, such as [Rubular](#). For more information about Ruby regular expressions, see [Mastering Ruby Regular Expressions](#).

- **streamName** (string) – the name of the log delivery stream that will be created at the destination.

Examples:

The following example shows FireLens parameters in the `deployment.json` file for logging to a single destination - CloudWatch – using a Ruby regular expression.

```
"fireLensParameters": {
  "enableFireLensLogging": true,
  "logDestinations": [
    {
      "service": "cloudwatch",
      "regexFilter": "^.*INFO.*$",
      "streamName": "Info"
    }
  ]
},
```

This example shows FireLens parameters in the `deployment.json` file for logging to a single destination – Firehose – using a Ruby regular expression.

```
"fireLensParameters": {
  "enableFireLensLogging": true,
  "logDestinations": [
    {
      "service": "firehose",
      "regexFilter": "^.*INFO.*$",

```

```
        "streamName": "Info"
      }
    ]
  },
```

This example shows FireLens parameters in the `deployment.json` file for routing separate log files to different destinations in CloudWatch, using Ruby regular expressions.

```
"fireLensParameters": {
  "enableFireLensLogging": true,
  "logDestinations": [
    {
      "service": "cloudwatch",
      "regexFilter": "^.*INFO.*$",
      "streamName": "Info"
    },
    {
      "service": "cloudwatch",
      "regexFilter": "^.*WARNING.*$",
      "streamName": "Warning"
    }
  ]
},
```

Step 3: Validate deployment artifacts

The last step before deployment is to ensure that your Amazon ECS task definitions and AWS CloudFormation templates are configured as expected after running the **generate app-deployment** command, and that your log destinations were created, if applicable.

Note

- Deployment artifacts are stored in the Amazon ECS or Amazon EKS deployment folder within the application folder that App2Container created for you. For example: `/root/app2container/java-tomcat-9e8e4799`
- If you are routing to CloudWatch, your routing destination is not created prior to deployment.

1. Run the [generate app-deployment](#) command to generate container deployment artifacts.

2. Verify that the Amazon ECS task definitions include the parameters that you specified and that the values are correct. For an example of FireLens parameters in an Amazon ECS task definition, see [Example: Amazon ECS task definition FireLens parameters](#)
3. Verify that the AWS CloudFormation template includes the parameters that you specified and that the values are correct. For an example of FireLens parameters in a CloudFormation template, expand the following section: [Example: AWS CloudFormation template FireLens parameters](#)
4. If you are routing logs to Kinesis Data Streams or Firehose, verify that the streams have been created for you by using the AWS Management Console.
 - a. Sign in to the AWS Management Console and open the Kinesis console at <https://console.aws.amazon.com/kinesis>.
 - b. From the Amazon Kinesis dashboard, choose **Data streams** or **Delivery streams** from the navigation pane.
 - c. Verify that your stream **Status** is Active.

Example: Amazon ECS task definition FireLens parameters

This example shows excerpts from an Amazon ECS task definition file that was generated for logging to CloudWatch.

```
"executionRoleArn": arn:aws:iam::
<YOUR_ACCOUNT_ID>;role/A2CEcsFirelensRole",
"containerDefinitions": [
  {
    ...
    "logConfiguration": {
      "logDriver": "awsfirelens",
      "secretOptions": null,
      "options": {
        "include-pattern": "^.*INFO.*$",
        "log_group_name": "java-tomcat-c770eed9-logs",
        "log_stream_name": "java-tomcat-c770eed9-Info",
        "auto_create_group": "true",
        "region": "us-east-1",
        "Name": "cloudwatch"
      }
    },
    ...
  },
  ...
]
```

```

    "name": "java-tomcat-c770eed9"
  },
  {
    "dnsSearchDomains": null,
    "environmentFiles": null,
    "logConfiguration": {
      "logDriver": "awslogs",
      "secretOptions": null,
      "options": {
        "awslogs-group": "/ecs/containerization",
        "awslogs-region": "us-east-1",
        "awslogs-create-group": "true",
        "awslogs-stream-prefix": "firelens"
      }
    },
    ...
    "firelensConfiguration": {
      "type": "fluentbit",
      "options": null
    },
    ...
    "name": "java-tomcat-c770eed9-log-router"
  }
],
...
"taskRoleArn": arn:aws:iam::
<YOUR_ACCOUNT_ID>;role/A2CEcsFirelensRole",
"compatibilities": [
  "EC2",
  "FARGATE"
],
...
"requiresAttributes": [
  {
    "targetId": null,
    "targetType": null,
    "value": null,
    "name": "ecs.capability.execution-role-awslogs"
  },
  ...
  {
    "targetId": null,
    "targetType": null,
    "value": null,

```

```

    "name": "com.amazonaws.ecs.capability.logging-driver.awsfirelens"
  },
  ...
  {
    "targetId": null,
    "targetType": null,
    "value": null,
    "name": "com.amazonaws.ecs.capability.logging-driver.awslogs"
  },
  ...
  {
    "targetId": null,
    "targetType": null,
    "value": null,
    "name": "ecs.capability.firelens.fluentbit"
  }
],

```

Example: AWS CloudFormation template FireLens parameters

This example shows excerpts from a CloudFormation template file that was generated for logging to CloudWatch.

```

Metadata:
  AWS::CloudFormation::Interface:
    ParameterGroups:
      ...
      - Label:
          default: Logging Parameters for the application being deployed, check ecs-lb-
webapp.yml for usage
        Parameters:
          - TaskLogDriver
          - MultipleDests
          - SingleDestName
          - IncludePattern
          - LogGrpName
          - LogStrmName
          - AutoCrtGrp
          - FirehoseStream
          - KinesisStream
          - KinesisAppendNewline
          - FirelensName
          - FirelensImage

```

```
- ConfigType
- ConfigPath
- UsingCloudwatchLogs
- UsingFirehoseLogs
- UsingKinesisLogs

...
Parameters:
...
# Firelens Parameters for the application being deployed
TaskLogDriver:
  Type: String
  Default: awsfirelens
MultipleDests:
  Type: String
  AllowedValues: [true, false]
  Default: false
SingleDestName:
  Type: String
  Default: cloudwatch
IncludePattern:
  Type: String
  Default: ^.*INFO.*$
LogGrpName:
  Type: String
  Default: java-tomcat-c770eed9-logs
LogStrmName:
  Type: String
  Default: java-tomcat-c770eed9-Info
AutoCrtGrp:
  Type: String
  Default: true
FirehoseStream:
  Type: String
  Default: ""
KinesisStream:
  Type: String
  Default: ""
KinesisAppendNewline:
  Type: String
  Default: ""
FirelensName:
  Type: String
  Default: java-tomcat-c770eed9-log-router
FirelensImage:
```

```

    Type: String
    Default: 906394416424.dkr.ecr.us-east-1.amazonaws.com/aws-for-fluent-bit:latest
  ConfigType:
    Type: String
    Default: ""
  ConfigPath:
    Type: String
    Default: ""
  UsingCloudwatchLogs:
    Type: String
    Default: true
  UsingFirehoseLogs:
    Type: String
    Default: false
  UsingKinesisLogs:
    Type: String
    Default: false
  ...
Rules:
  FirelensSingleCloudwatch:
    RuleCondition: !And
      - !Equals [ !Ref MultipleDests, 'false']
      - !Equals [ !Ref UsingCloudwatchLogs, 'true']
    Assertions:
      - AssertDescription: You cannot use any other firelens destination if a single
cloudwatch stream is desired
        Assert: !And
          - !Equals [ !Ref UsingFirehoseLogs, 'false']
          - !Equals [ !Ref UsingKinesisLogs, 'false']
          - !Equals [ !Ref SingleDestName, "cloudwatch" ]
          - !Not [ !Equals [ !Ref LogGrpName, "" ]]
          - !Not [ !Equals [ !Ref LogStrmName, "" ]]
          - !Not [ !Equals [ !Ref AutoCrtGrp, "" ]]
  FirelensSingleFirehose:
    RuleCondition: !And
      - !Equals [ !Ref MultipleDests, 'false']
      - !Equals [ !Ref UsingFirehoseLogs, 'true']
    Assertions:
      - AssertDescription: You cannot use any other firelens destination if a single
firehose stream is desired
        Assert: !And
          - !Equals [ !Ref UsingCloudwatchLogs, 'false']
          - !Equals [ !Ref UsingKinesisLogs, 'false']
          - !Equals [ !Ref SingleDestName, "firehose" ]

```

```

    - !Not [ !Equals [ !Ref FirehoseStream, "" ] ]
FirelensSingleKinesis:
  RuleCondition: !And
    - !Equals [ !Ref MultipleDests, 'false' ]
    - !Equals [ !Ref UsingKinesisLogs, 'true' ]
  Assertions:
    - AssertDescription: You cannot use any other firelens destination if a single
      kinesis stream is desired
    Assert: !And
      - !Equals [ !Ref UsingCloudwatchLogs, 'false' ]
      - !Equals [ !Ref UsingFirehoseLogs, 'false' ]
      - !Equals [ !Ref SingleDestName, "kinesis" ]
      - !Not [ !Equals [ !Ref KinesisStream, "" ] ]
      - !Not [ !Equals [ !Ref KinesisAppendNewline, "" ] ]
MultipleDestinations:
  RuleCondition: !Equals [ !Ref MultipleDests, 'true' ]
  Assertions:
    - AssertDescription: You must supply a configuration file location and filepath
      if multiple firelens destinations are being used
    Assert: !And
      - !Not [ !Equals [ !Ref ConfigType, "" ] ]
      - !Not [ !Equals [ !Ref ConfigPath, "" ] ]
      - !Equals [ !Ref SingleDestName, "" ]
      - !Equals [ !Ref IncludePattern, "" ]
      - !Equals [ !Ref LogGrpName, "" ]
      - !Equals [ !Ref LogStrmName, "" ]
      - !Equals [ !Ref AutoCrtGrp, "" ]
      - !Equals [ !Ref FirehoseStream, "" ]
      - !Equals [ !Ref KinesisStream, "" ]
      - !Equals [ !Ref KinesisAppendNewline, "" ]
...
Conditions:
...
Resources:
  PrivateAppStack:
    Type: AWS::CloudFormation::Stack
    Condition: DoNotCreatePublicLoadBalancer
    Properties:
      TemplateURL: !Sub 'https://${S3Bucket}.s3.${S3Region}.${AWS::URLSuffix}/
        ${S3KeyPrefix}/ecs-private-app.yml'
      Tags:
        - Key: "a2c-generated"
          Value: !Sub 'ecs-app-${AWS::StackName}'
    Parameters:

```

```
...
    TaskLogDriver: !Ref TaskLogDriver
    MultipleDests: !Ref MultipleDests
    SingleDestName: !Ref SingleDestName
    IncludePattern: !Ref IncludePattern
    LogGrpName: !Ref LogGrpName
    LogStrmName: !Ref LogStrmName
    AutoCrtGrp: !Ref AutoCrtGrp
    FirehoseStream: !Ref FirehoseStream
    KinesisStream: !Ref KinesisStream
    KinesisAppendNewline: !Ref KinesisAppendNewline
    FirelensName: !Ref FirelensName
    FirelensImage: !Ref FirelensImage
    ConfigType: !Ref ConfigType
    ConfigPath: !Ref ConfigPath
    UsingCloudwatchLogs: !Ref UsingCloudwatchLogs
    UsingFirehoseLogs: !Ref UsingFirehoseLogs
    UsingKinesisLogs: !Ref UsingKinesisLogs
...
```

Step 4: Deploy your application to Amazon ECS

Deploy your application using the [generate app-deployment](#) command with the `--deploy` option.

```
$ sudo app2container generate app-deployment --deploy --application-id java-
tomcat-9e8e4799
✓ AWS prerequisite check succeeded
✓ Docker prerequisite check succeeded
✓ Created ECR Repository
✓ Registered ECS Task Definition with ECS
✓ Uploaded CloudFormation resources to S3 Bucket: app2container-example
✓ Generated CloudFormation Master template at: /root/app2container/java-
tomcat-9e8e4799/EcsDeployment/ecs-master.yml
✓ Initiated CloudFormation stack creation. This may take a few minutes. Please visit
the AWS CloudFormation Console to track progress.
ECS deployment successful for application java-tomcat-9e8e4799

The URL to your Load Balancer Endpoint is:
<your endpoint>.us-east-1.elb.amazonaws.com
Successfully created ECS stack app2container-java-tomcat-9e8e4799-ECS. Check the AWS
CloudFormation Console for additional details.
```

Alternatively, you can deploy your application's AWS CloudFormation template using the AWS CLI as follows.

```
$ sudo aws cloudformation deploy --template-file /root/app2container/java-  
tomcat-9e8e4799/EcsDeployment/ecs-master.yml --capabilities CAPABILITY_NAMED_IAM --  
stack-name app2container-java-tomcat-9e8e4799-ECS
```

Step 5: Verify log routing

After you deploy your application to Amazon ECS, you can verify that your logs are routing to their intended destinations.

Security in AWS App2Container

Security at AWS is the highest priority. As an AWS customer using AWS App2Container and tools such as Amazon ECR, Amazon ECS, and Amazon EKS, you benefit from data centers and network architectures that are built to meet the requirements of the most security-sensitive organizations.

Security is a shared responsibility between AWS and you. The [shared responsibility model](#) describes this as security of the cloud and security in the cloud:

- **Security of the cloud** – AWS is responsible for protecting the infrastructure that runs AWS services in the AWS Cloud. AWS also provides you with services that you can use securely. Third-party auditors regularly test and verify the effectiveness of our security as part of the [AWS Compliance Programs](#). To learn about the compliance programs that apply to Amazon EC2, see [AWS Services in Scope by Compliance Program](#).
- **Security in the cloud** – Your responsibility is determined by the AWS service that you use. You are also responsible for other factors including the sensitivity of your data, your company's requirements, and applicable laws and regulations.

Contents

- [Data protection in App2Container](#)
- [Identity and access management in App2Container](#)
- [Update management for App2Container](#)

Data protection in App2Container

The AWS [shared responsibility model](#) applies to data protection in AWS App2Container. As described in this model, AWS is responsible for protecting the global infrastructure that runs all of the AWS Cloud. You are responsible for maintaining control over your content that is hosted on this infrastructure. You are also responsible for the security configuration and management tasks for the AWS services that you use. For more information about data privacy, see the [Data Privacy FAQ](#). For information about data protection in Europe, see the [AWS Shared Responsibility Model and GDPR](#) blog post on the *AWS Security Blog*.

For data protection purposes, we recommend that you protect AWS account credentials and set up individual users with AWS IAM Identity Center or AWS Identity and Access Management (IAM).

That way, each user is given only the permissions necessary to fulfill their job duties. We also recommend that you secure your data in the following ways:

- Use multi-factor authentication (MFA) with each account.
- Use SSL/TLS to communicate with AWS resources. We require TLS 1.2 and recommend TLS 1.3.
- Set up API and user activity logging with AWS CloudTrail. For information about using CloudTrail trails to capture AWS activities, see [Working with CloudTrail trails](#) in the *AWS CloudTrail User Guide*.
- Use AWS encryption solutions, along with all default security controls within AWS services.
- Use advanced managed security services such as Amazon Macie, which assists in discovering and securing sensitive data that is stored in Amazon S3.
- If you require FIPS 140-3 validated cryptographic modules when accessing AWS through a command line interface or an API, use a FIPS endpoint. For more information about the available FIPS endpoints, see [Federal Information Processing Standard \(FIPS\) 140-3](#).

We strongly recommend that you never put confidential or sensitive information, such as your customers' email addresses, into tags or free-form text fields such as a **Name** field. This includes when you work with App2Container or other AWS services using the console, API, AWS CLI, or AWS SDKs. Any data that you enter into tags or free-form text fields used for names may be used for billing or diagnostic logs. If you provide a URL to an external server, we strongly recommend that you do not include credentials information in the URL to validate your request to that server.

Data encryption

App2Container communicates with AWS services using standard APIs when retrieving artifacts from Amazon S3 or pushing Docker containers to service endpoints in the AWS container management suite (Amazon ECR, Amazon ECS, and Amazon EKS). It works with AWS CloudFormation and AWS CodeStar services to generate and deploy relevant container and lifecycle artifacts using their standard APIs.

Encryption at rest

- App2Container installation packages are kept in a private Amazon S3 bucket with encryption enabled.
- Application artifacts can optionally be uploaded into Amazon S3 buckets. Enable encryption for your Amazon S3 bucket to enforce data encryption.

Encryption in transit

- App2Container installation packages are kept in a private Amazon S3 bucket, which requires secure download using the HTTPS protocol using links provided for each package.
- App2Container uses standard AWS APIs for the services it interacts with, including Amazon ECR, Amazon ECS, Amazon EKS, AWS CloudFormation, CodePipeline, and Amazon S3. AWS APIs use HTTPS as their default communication protocol.

Internetwork traffic privacy

App2Container does not store passwords, keys, or other secrets or customer-sensitive material. App2Container also ensures that no sensitive fields are contained in application logs.

Identity and access management in App2Container

Your AWS security credentials identify you to AWS and grant you access to your AWS resources. For example, they can allow you to access artifacts saved to an Amazon S3 bucket. You can use features of AWS Identity and Access Management (IAM) to allow other users, services, and applications to use specific resources in your AWS account without sharing your security credentials. You can choose to allow full use or limited use of your AWS resources.

If you are the owner of the AWS account and use AWS as the root user, we strongly recommend that you create an IAM admin user to use for access to your AWS resources. See [Creating Your First IAM Admin User and Group](#) in the *IAM User Guide* to set up your own access before setting up any other IAM users who need to use App2Container.

By default, IAM users don't have permission to create or modify resources. To allow IAM users to create or modify resources and perform tasks, you must create IAM policies that grant permission to use the specific resources and API actions that they need. For more information about IAM policies, see [Policies and Permissions](#) in the *IAM User Guide*.

IAM groups and roles are a flexible way to manage permissions across multiple users. When you assign a user to a group or when your user assumes a role, that user inherits the group's or role's permissions, and is allowed or denied permission to perform the specified tasks on the specified resources. You can assign multiple users to the same group, and a role can be assumed by authorized users. While groups and roles both serve the purpose of granting access to resources, roles are more task-oriented, and assuming a role provides you with temporary security credentials for your role session.

IAM security best practices

Follow these top four security best practices when setting up your IAM resources. For more information and additional best practices, see [Security Best Practices in IAM](#) in the *IAM User Guide*.

1. Lock away your AWS account root user access keys

Protect your root user access key like you would your credit card numbers or any other sensitive secret, and only use your root user account for necessary account and service management tasks.

2. Create individual IAM users

Don't use your AWS account root user credentials to access AWS, and don't give your credentials to anyone else. Instead, create individual users for anyone who needs access to your AWS account.

3. Use groups or roles to assign permissions to IAM Users

Instead of defining permissions for individual IAM users, it's usually more convenient to create groups that relate to job functions (administrators, developers, accounting, etc.) or roles that relate to specific tasks.

4. Grant least privilege

When you create IAM policies, follow the standard security advice of granting *least privilege*, or granting only the permissions required to perform a task. Determine what users (and roles) need to do and then craft policies that allow them to perform only those tasks.

We recommend that you create a general purpose IAM group that can run all of the commands *except* commands that are run with the `--deploy` option.

If you plan to use App2Container to deploy your containers or create pipelines, then you should create a separate IAM user for deployments. The deployment user needs to be able to create or update AWS objects for container management services (Amazon ECR, Amazon ECS, Amazon EKS, and App Runner), and to create pipelines with AWS CodeStar services. This requires elevated permissions that should only be used for deployment.

Set up IAM resources for App2Container

- [Create IAM resources for general use](#)
- [Create IAM resources for deployment](#)

Create IAM resources for general use

Follow best practices by using the following steps to create an IAM group with access to perform specific tasks, using specific resources, and to assign users to the group.

Note

Alternatively, you can create an IAM role and EC2 instance profile to grant permissions to applications that run on an Amazon EC2 instance. For more information about using instance profiles, see [Using an IAM role to grant permissions to applications running on Amazon EC2 instances](#) in the *IAM User Guide*.

1. Create a customer managed IAM policy

You can create a customer managed IAM policy for your general purpose user or group, using one of the [example policies](#) on this page after you have customized the JSON to refer to your resources. To create a policy using the AWS console, see [Creating policies on the JSON tab](#) in the *IAM User Guide*. To create a policy using the AWS CLI, use the [create-policy](#) command.

Tip

Review your policy periodically, to add actions required for newer features, and to ensure that the policy continues to meet your needs.

2. Create IAM users and a group

Every user who will run **app2container** commands needs to have an IAM user created for accessing AWS resources under your account. To follow best practices, you can create an IAM group with the policy attached, and assign users to it.

To create an IAM user, see [Creating an IAM User in Your AWS Account](#) in the *IAM User Guide*. Be sure to select programmatic access to AWS when you create the IAM user.

Perform the following steps to create an IAM group and assign users to it.

- a. To create an IAM group, see [Creating IAM Groups](#) in the *IAM User Guide*.
- b. Ensure that every person who will run **app2container** commands has an IAM user defined for AWS access.
- c. To assign the users to the group that you created in step 1a, see [Adding Permissions to a User \(Console\)](#), or [Adding and Removing a User's Permissions \(AWS CLI or AWS API\)](#) in the *IAM User Guide*.

3. Save your AWS access keys

Save the access keys for your new or existing IAM user in a safe place. You'll need them to [configure your AWS profile](#) as part of getting set up for App2Container.

4. Attach or assign the policy

Use one of the following methods to assign permissions to your IAM users.

- **Attach the policy to the IAM group**

Attach the policy that you created in step 1 to the group that you created in step 2. See [Attaching a Policy to an IAM Group](#) in the *IAM User Guide*.

- **Embed the policy inline for an IAM user**

Embed the policy that you created in step 1 inline for your IAM user. See the section that begins with "To embed an inline policy" in [Adding Permissions to a User \(Console\)](#), or [Adding and Removing a User's Permissions \(AWS CLI or AWS API\)](#) in the *IAM User Guide*.

Example IAM policies

You can use one of the policy templates in this section as a starting point to configure the access that App2Container uses on your behalf to generate the deployment artifacts for your application containers.

Choose the policy resources and actions that you need

The following sections in the example policies depend on choices you've made for your containerization environment and workflow:

- **AWS CodeCommit**

SectionForCodeCommitAccess – If you use App2Container to generate a container pipeline, you must grant access to interact with your CodeCommit code repository.

- **FireLens log routing to Amazon Data Firehose**

SectionForFirelensFirehoseIAMPolicyAccess, **SectionForFirelensFirehoseIAMRoleAccess**, and **SectionForFirelensFirehoseStreamsAccess** – If you use FireLens for log file routing, and you configure FireLens to route to Firehose, you must grant access for App2Container to create a new Firehose delivery stream. You must also grant access for App2Container to create an IAM policy and role so that FireLens can access the delivery stream.

- **FireLens log routing to Amazon Kinesis Data Streams**

SectionForFirelensKinesisStreamsAccess – if you use FireLens for log file routing, and you configure FireLens to route to Kinesis Data Streams you must grant access for App2Container to create a new Kinesis data stream.

- **AWS Secrets Manager**

SectionForSecretManagerAccess – If you configured your environment to run remote workflows, App2Container requires you to use Secrets Manager for connection secrets to access application servers from the worker machine. You must grant access to retrieve secrets in the policy.

- **Amazon S3**

SectionForS3Access and **SectionForS3ReadAccess** – If you set up an S3 bucket for application or deployment artifacts, you must grant access to your bucket in the policy.

You must also ensure that only authorized users can access the bucket. We recommend that you use server-side encryption for your bucket. See [Protecting data using server-side encryption](#) in the *Amazon Simple Storage Service User Guide* for more information about how to set it up.

- **Upload support bundle**

SectionForUploadSupportBundleService – If you chose to have App2Container logs and command-generated artifacts uploaded automatically for failed commands when you ran the **init** command, you must grant access to upload the application support bundles.

- **Usage metrics**

SectionForMetricsService – If you gave consent for App2Container to collect and export application usage metrics when you ran the **init** command, you must grant access to upload the metric data.

- **Amazon VPC**

SectionForByoVPC – If you specify your own VPC or want to reuse an existing VPC that App2Container created for a prior deployment, you must grant access to associated describe actions in the policy.

Other policy sections in the examples are required for App2Container to generate application deployment artifacts, or to integrate with Jenkins pipelines.

IAM policy for Amazon ECS

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "SectionForS3Access",
      "Action": [
        "s3:DeleteObject",
        "s3:GetBucketAcl",
        "s3:GetBucketLocation",
        "s3:GetObject",
        "s3:GetObjectAcl",
        "s3:ListAllMyBuckets",
        "s3:ListBucket",
        "s3:PutObject",
        "s3:PutObjectAcl"
      ],
      "Effect": "Allow",
      "Resource": "<amzn-s3-demo-bucket-ARN>"
    },
    {
      "Sid": "SectionForS3ReadAccess",
      "Effect": "Allow",
      "Action": [
        "s3:ListBucket",
        "s3:GetBucketAcl"
      ],
      "Resource": "arn:aws:s3:::*"
    }
  ]
}
```



```

    },
    {
      "Sid": "SectionForECRAccess",
      "Action": [
        "ecr:BatchCheckLayerAvailability",
        "ecr:BatchDeleteImage",
        "ecr:BatchGetImage",
        "ecr:CompleteLayerUpload",
        "ecr:CreateRepository",
        "ecr>DeleteRepository",
        "ecr:DescribeImages",
        "ecr:DescribeRepositories",
        "ecr:GetAuthorizationToken",
        "ecr:GetDownloadUrlForLayer",
        "ecr:GetRepositoryPolicy",
        "ecr:InitiateLayerUpload",
        "ecr:ListImages",
        "ecr:PutImage",
        "ecr:TagResource",
        "ecr:UntagResource",
        "ecr:UploadLayerPart"
      ],
      "Effect": "Allow",
      "Resource": "<resource-ARNs>"
    },
    {
      "Sid": "SectionForECRAccess2",
      "Action": [
        "ecr:GetAuthorizationToken"
      ],
      "Effect": "Allow",
      "Resource": "*"
    },
    {
      "Sid": "SectionForECSWriteAccess",
      "Action": [
        "ecs:CreateCluster",
        "ecs:CreateService",
        "ecs:CreateTaskSet",
        "ecs>DeleteCluster",
        "ecs>DeleteService",
        "ecs>DeleteTaskSet",
        "ecs:DeregisterTaskDefinition",
        "ecs:Poll",

```

```

        "ecs:RegisterContainerInstance",
        "ecs:RegisterTaskDefinition",
        "ecs:RunTask",
        "ecs:StartTask",
        "ecs:StopTask",
        "ecs:SubmitContainerStateChange",
        "ecs:SubmitTaskStateChange",
        "ecs:UpdateContainerInstancesState",
        "ecs:UpdateService",
        "ecs:UpdateServicePrimaryTaskSet",
        "ecs:UpdateTaskSet"
    ],
    "Effect": "Allow",
    "Resource": "<resource-ARNs>"
},
{
    "Sid": "SectionForPassRoleToECS",
    "Effect": "Allow",
    "Action": "iam:PassRole",
    "Resource": "<ARN for ecsTaskExecutionRole>"
},
{
    "Sid": "SectionForECSReadAccess",
    "Action": [
        "ecs:DescribeClusters",
        "ecs:DescribeContainerInstances",
        "ecs:DescribeServices",
        "ecs:DescribeTaskDefinition",
        "ecs:DescribeTaskSets",
        "ecs:DescribeTasks",
        "ecs:ListClusters",
        "ecs:ListContainerInstances",
        "ecs:ListServices",
        "ecs:ListTaskDefinitionFamilies",
        "ecs:ListTaskDefinitions",
        "ecs:ListTasks"
    ],
    "Effect": "Allow",
    "Resource": "<resource-ARNs>"
},
{
    "Sid": "SectionForFirelensIAMRoleAccess",
    "Action": [
        "iam:CreateRole",

```

```

        "iam:GetRole",
        "iam:AttachRolePolicy"
    ],
    "Effect": "Allow",
    "Resource": "arn:aws:iam::<your account ID>:role/A2CEcsFirelensRole"
},
{
    "Sid": "SectionForFirelensIAMPolicyAccess",
    "Action": [
        "iam:CreatePolicy"
    ],
    "Effect": "Allow",
    "Resource": "arn:aws:iam::<your account ID>:policy/service-role/
A2CEcsFirelensPolicy"
},
{
    "Sid": "SectionForFirelensFirehoseIAMPolicyAccess",
    "Action": [
        "iam:CreatePolicy",
        "iam:GetPolicy"
    ],
    "Effect": "Allow",
    "Resource": "arn:aws:iam::<your account ID>:policy/*a2c-
KinesisFirehosePolicy-*"
},
{
    "Sid": "SectionForFirelensFirehoseIAMRoleAccess",
    "Action": [
        "iam:CreateRole",
        "iam:GetRole",
        "iam:AttachRolePolicy"
    ],
    "Effect": "Allow",
    "Resource": "arn:aws:iam::<your account ID>:role/*a2c-FirehoseRole-*"
},
{
    "Sid": "SectionForFirelensFirehoseStreamsAccess",
    "Action": [
        "firehose:DescribeDeliveryStream",
        "firehose:CreateDeliveryStream"
    ],
    "Effect": "Allow",
    "Resource": "arn:aws:firehose:*:<your account ID>:deliverystream/*"
},

```

```

{
  "Sid": "SectionForFirelensKinesisStreamsAccess",
  "Action": [
    "kinesis:CreateStream"
  ],
  "Effect": "Allow",
  "Resource": "arn:aws:kinesis:*:<your account ID>:stream/*"
},
{
  "Sid": "SectionForCodeCommitAccess",
  "Effect": "Allow",
  "Action": [
    "codecommit:GetRepository",
    "codecommit:GetBranch",
    "codecommit:CreateRepository",
    "codecommit:CreateCommit",
    "codecommit:TagResource"
  ],
  "Resource": "arn:aws:codecommit:*:*:*"
},
{
  "Sid": "SectionForByoVPC",
  "Effect": "Allow",
  "Action": [
    "ec2:DescribeInternetGateways",
    "ec2:DescribeRouteTables",
    "ec2:DescribeSubnets",
    "ec2:DescribeVpcs"
  ],
  "Resource": "<resource-ARNs>"
},
{
  "Sid": "SectionForEC2",
  "Effect": "Allow",
  "Action": [
    "ec2:DescribeKeyPairs",
    "ec2:CreateKeyPair",
    "ec2:DescribeAvailabilityZones"
  ],
  "Resource": "<resource-ARNs>"
},
{
  "Sid": "SectionForMetricsService",
  "Effect": "Allow",

```

```

        "Action": "application-transformation:PutMetricData",
        "Resource": "*"
    },
    {
        "Sid": "SectionForUploadSupportBundleService",
        "Effect": "Allow",
        "Action": "application-transformation:PutLogData",
        "Resource": "*"
    },
    {
        "Sid": "SectionForSecretManagerAccess",
        "Action": [
            "secretsmanager:GetSecretValue",
            "secretsmanager:DescribeSecret"
        ],
        "Effect": "Allow",
        "Resource": "arn:aws:secretsmanager:<your region>:<your account ID>:secret:a2c/*"
    },
    {
        "Sid": "SectionForCloudFormation",
        "Action": [
            "cloudformation:DescribeStacks"
        ],
        "Effect": "Allow",
        "Resource": "arn:aws:cloudformation:*:<your account ID>:stack/a2c-*"
    }
]
}

```

IAM policy for Amazon EKS

```

{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "SectionForS3Access",
            "Action": [
                "s3:DeleteObject",
                "s3:GetBucketAcl",
                "s3:GetBucketLocation",
                "s3:GetObject",
                "s3:GetObjectAcl",

```

```

        "s3:ListAllMyBuckets",
        "s3:ListBucket",
        "s3:PutObject",
        "s3:PutObjectAcl"
    ],
    "Effect": "Allow",
    "Resource": "<amzn-s3-demo-bucket-ARN>"
},
{
    "Sid": "SectionForS3ReadAccess",
    "Effect": "Allow",
    "Action": [
        "s3:ListBucket",
        "s3:GetBucketAcl"
    ],
    "Resource": "arn:aws:s3::*"
},
{
    "Sid": "SectionForECRAccess",
    "Action": [
        "ecr:BatchCheckLayerAvailability",
        "ecr:BatchDeleteImage",
        "ecr:BatchGetImage",
        "ecr:CompleteLayerUpload",
        "ecr:CreateRepository",
        "ecr>DeleteRepository",
        "ecr:DescribeImages",
        "ecr:DescribeRepositories",
        "ecr:GetDownloadUrlForLayer",
        "ecr:GetRepositoryPolicy",
        "ecr:InitiateLayerUpload",
        "ecr:ListImages",
        "ecr:PutImage",
        "ecr:TagResource",
        "ecr:UntagResource",
        "ecr:UploadLayerPart"
    ],
    "Effect": "Allow",
    "Resource": "<resource-ARNs>"
},
{
    "Sid": "SectionForECRAccess2",
    "Action": [
        "ecr:GetAuthorizationToken"
    ]
}

```

```

    ],
    "Effect": "Allow",
    "Resource": "*"
  },
  {
    "Sid": "SectionForEKS",
    "Effect": "Allow",
    "Action": [
      "iam:GetRole",
      "lambda:GetFunction"
    ],
    "Resource": [
      "arn:aws:iam::*:role/eks-quickstart-ResourceReader",
      "arn:aws:lambda:<target Region>:*:function:eks-quickstart-
ResourceReader"
    ]
  },
  {
    "Sid": "SectionForCodeCommitAccess",
    "Effect": "Allow",
    "Action": [
      "codecommit:GetRepository",
      "codecommit:GetBranch",
      "codecommit:CreateRepository",
      "codecommit:CreateCommit",
      "codecommit:TagResource"
    ],
    "Resource": "arn:aws:codecommit:*:*:*"
  },
  {
    "Sid": "SectionForByoVPC",
    "Effect": "Allow",
    "Action": [
      "ec2:DescribeInternetGateways",
      "ec2:DescribeRouteTables",
      "ec2:DescribeSubnets",
      "ec2:DescribeVpcs"
    ],
    "Resource": "<resource-ARNs>"
  },
  {
    "Sid": "SectionForEC2",
    "Effect": "Allow",
    "Action": [

```

```

        "ec2:DescribeKeyPairs",
        "ec2:CreateKeyPair",
        "ec2:DescribeAvailabilityZones"
    ],
    "Resource": "<resource-ARNs>"
},
{
    "Sid": "SectionForMetricsService",
    "Effect": "Allow",
    "Action": "application-transformation:PutMetricData",
    "Resource": "*"
},
{
    "Sid": "SectionForUploadSupportBundleService",
    "Effect": "Allow",
    "Action": "application-transformation:PutLogData",
    "Resource": "*"
},
{
    "Sid": "SectionForSecretManagerAccess",
    "Action": [
        "secretsmanager:GetSecretValue",
        "secretsmanager:DescribeSecret"
    ],
    "Effect": "Allow",
    "Resource": "arn:aws:secretsmanager:<your region>:<your account
ID>:secret:a2c/*"
},
{
    "Sid": "SectionForIAMAccess",
    "Action": [
        "iam:AttachRolePolicy",
        "iam:CreateRole",
        "iam:GetRole",
        "iam:ListRoles (https://docs.aws.amazon.com/IAM/latest/APIReference/
API\_ListRoles.html)",
        "iam:ListRoleTags (https://docs.aws.amazon.com/IAM/latest/APIReference/
API\_ListRoleTags.html)"
    ],
    "Effect": "Allow",
    "Resource": "<resource-ARNs>"
},
{
    "Sid": "SectionForCloudFormation",

```



```

        "Action": [
            "cloudformation:DescribeStacks"
        ],
        "Effect": "Allow",
        "Resource": "arn:aws:cloudformation:*<your account ID>:stack/a2c-*"
    }
]
}

```

IAM policy for AWS App Runner

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "SectionForAppRunnerAccess",
      "Action": [
        "apprunner:List*",
        "apprunner:Describe*"
      ],
      "Effect": "Allow",
      "Resource": "<resource-ARNs>"
    },
    {
      "Sid": "SectionForECRAccess",
      "Action": [
        "ecr:BatchCheckLayerAvailability",
        "ecr:BatchDeleteImage",
        "ecr:BatchGetImage",
        "ecr:CompleteLayerUpload",
        "ecr:CreateRepository",
        "ecr>DeleteRepository",
        "ecr:DescribeImages",
        "ecr:DescribeRepositories",
        "ecr:GetDownloadUrlForLayer",
        "ecr:GetRepositoryPolicy",
        "ecr:InitiateLayerUpload",
        "ecr:ListImages",
        "ecr:PutImage",
        "ecr:TagResource",
        "ecr:UntagResource",
        "ecr:UploadLayerPart"
      ],
    },
  ],
}

```

```

        "Effect": "Allow",
        "Resource": "<resource-ARNs>"
    },
    {
        "Sid": "SectionForECRAccess2",
        "Action": [
            "ecr:GetAuthorizationToken"
        ],
        "Effect": "Allow",
        "Resource": "*"
    },
    {
        "Sid": "SectionForCodeCommitAccess",
        "Effect": "Allow",
        "Action": [
            "codecommit:GetRepository",
            "codecommit:GetBranch",
            "codecommit:CreateRepository",
            "codecommit:CreateCommit",
            "codecommit:TagResource"
        ],
        "Resource": "arn:aws:codecommit:*:*:*"
    },
    {
        "Sid": "SectionForMetricsService",
        "Effect": "Allow",
        "Action": "application-transformation:PutMetricData",
        "Resource": "*"
    },
    {
        "Sid": "SectionForUploadSupportBundleService",
        "Effect": "Allow",
        "Action": "application-transformation:PutLogData",
        "Resource": "*"
    },
    {
        "Sid": "SectionForSecretManagerAccess",
        "Action": [
            "secretsmanager:GetSecretValue",
            "secretsmanager:DescribeSecret"
        ],
        "Effect": "Allow",
        "Resource": "arn:aws:secretsmanager:us-east-1:*:secret:a2c/*"
    },

```

```

    {
      "Sid": "SectionForCloudFormation",
      "Action": [
        "cloudformation:DescribeStacks"
      ],
      "Effect": "Allow",
      "Resource": "arn:aws:cloudformation:*:<your account ID>:stack/a2c-*"
    }
  ]
}

```

IAM policy for Azure DevOps pipelines

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AzureDevOpsAWS",
      "Effect": "Allow",
      "Action": [
        "ecr:DescribeRepositories",
        "ecr:GetAuthorizationToken",
        "ecr:UploadLayerPart",
        "ecr:PutImage",
        "ecr:CompleteLayerUpload",
        "ecr:InitiateLayerUpload",
        "ecr:BatchCheckLayerAvailability",
        "ecr:BatchGetImage",
        "ecr:GetDownloadUrlForLayer",
        "ecs:UpdateService",
        "eks:DescribeCluster"
      ],
      "Resource": "*"
    }
  ]
}

```

IAM policy for Jenkins pipelines

```

{
  "Version": "2012-10-17",
  "Statement": [
    {

```

```

        "Sid": "JenkinsAWS",
        "Effect": "Allow",
        "Action": [
            "ecr:GetDownloadUrlForLayer",
            "iam:ListRoles",
            "ecr:GetAuthorizationToken",
            "ecr:UploadLayerPart",
            "ecr:PutImage",
            "ecs:UpdateService",
            "sts:AssumeRole",
            "ecr:BatchGetImage",
            "ecr:CompleteLayerUpload",
            "eks:DescribeCluster",
            "ecr:InitiateLayerUpload",
            "ecr:BatchCheckLayerAvailability"
        ],
        "Resource": "*"
    }
]
}

```

Create IAM resources for deployment

The **AdministratorAccess** policy grants an IAM user full access to AWS. Therefore, IAM users with this policy can deploy a containerized application using any of the AWS services for deployment that are supported by App2Container.

1. Create an IAM user

You can create an IAM user with full access to AWS API actions and resources. Be sure to grant the user programmatic access to AWS and to attach the **AdministratorAccess** policy. For more information, see [Creating IAM users](#) in the *IAM User Guide*.

2. Save your AWS access keys

Save the access keys for the IAM user in a safe place. You'll need them to [configure your AWS profile](#) as part of getting set up for App2Container.

Update management for App2Container

App2Container detects what version of the CLI you are using when you run a command. It notifies you if there are published updates available. You can install the latest version of App2Container using the [upgrade](#) command.

App2Container command reference

Tip

To containerize your applications with a console-based experience and deploy them on Amazon ECS on AWS Fargate, you can use the *Replatform applications to Amazon ECS* template on the [AWS Migration Hub Orchestrator console](#). For more information, see [Replatform applications to Amazon ECS](#) in the *AWS Migration Hub Orchestrator User Guide*.

AWS App2Container is a command line tool that transforms supported legacy applications running on physical servers or virtual machines into applications that run in Docker containers on Amazon ECS, Amazon EKS, or AWS App Runner.

Important

Running App2Container commands on a Linux server requires elevated permissions. Prefix the command syntax with **sudo**, or run the **sudo su** command one time when you log in before running the commands as shown in the syntax for the commands linked below.

Containerization phases

The containerization process has several phases.

Phases

- [Initialize](#)
- [Analyze](#)
- [Transform](#)
- [Deploy](#)

Initialize

The **init** command performs one-time initialization tasks for App2Container. This interactive command prompts for the information required to set up the local App2Container environment.

Run this command before you run any other App2Container commands. If you are using a worker machine to run commands remotely on application servers, you must also run the **remote configure** command on the worker machine.

init

Run the **init** command to configure the AWS App2Container workspace on your application servers and worker machines. If you are using a worker machine, and running commands remotely, the **init** command is only required on the worker machine.

remote configure

After setting up remote access for the worker machine on your application server (see [Enable remote access for a worker machine \(optional\)](#)), run the **remote configure** command on the worker machine to configure the connections needed to run remote workflows on application servers. This interactive command prompts for the required information for each application server that you enter.

Analyze

After you have completed setup and initialization tasks on your servers, you can begin the analyze phase. Run the version of these commands that applies to your server setup:

Run commands directly on application servers

inventory

Run the **inventory** command to produce an inventory of applications that are running on your application servers, and to assign each one a unique ID to use when you run other commands.

analyze

Run the **analyze** command to analyze your running applications and to identify dependencies that are required for containerization. This command creates the `analysis.json` file that feeds into the Transform phase commands.

Run commands remotely from a worker machine

[remote inventory](#)

Run the **remote inventory** command from your worker machine to produce an inventory of applications that are running on your target application server and to assign each one a unique ID to use when you run other commands.

[remote analyze](#)

Run the **remote analyze** command from your worker machine to analyze the applications running on your target application server, and to identify dependencies that are required for containerization. This command creates the `analysis.json` file that feeds into the Transform phase commands.

Transform

The transform phase creates containers for your applications that have gone through analysis. Run the version of these commands that applies to your server setup:

Run the **extract** directly on application servers, or run the **remote extract** from a worker machine

[extract](#)

Run the **extract** command on your application server to generate an application archive based on the `analysis.json` file, created by the **analyze** command. Transfer the archive to the worker machine for the remaining steps that require the operating system to support containers.

[remote extract](#)

Run the **remote extract** command from your worker machine to generate an application archive for the applications running on your target application server, based on the `analysis.json` file that was created by the **analyze** command.

Run all remaining commands directly on application servers or on a worker machine

containerize

Run the **containerize** command for the application specified in the `--application id` parameter to do the following:

- Extract application artifacts or read from an extract archive for the specified application. For complex, multi-component Windows applications, this also applies to any additional applications or services that run in the same container.
- Generate Docker container artifacts, including a Dockerfile and container image, based on the application artifacts, and the application settings in the `analysis.json` file.
- Create the `deployment.json` file for input to the **generate app-deployment** command

Deploy

The deploy phase consists of deploying an application to your target container management environment (Amazon ECR with Amazon ECS, Amazon EKS, or AWS App Runner), and optionally creating a CI/CD pipeline to automate future deployments.

generate app-deployment

Option 1: Generate deployment artifacts and deploy directly

Run the **generate app-deployment** command with the `--deploy` option to generate container deployment artifacts and to deploy them to your target environment all in one step.

Option 2: Generate deployment artifacts and customize

- Run the **generate app-deployment** command without the deployment option to generate deployment artifacts.
- Review and customize the generated Amazon ECS, Amazon EKS, or AWS App Runner deployment artifacts.
- Deploy to your target environment using the AWS CLI or AWS console.

generate pipeline (optional)

Option 1: Generate CI/CD pipeline artifacts and deploy directly

Run the **generate pipeline** command with the `--deploy` option to generate CI/CD pipeline artifacts and to deploy them with AWS CodePipeline all in one step.

Option 2: Generate CI/CD pipeline artifacts and customize

- Run the **generate pipeline** command without the deployment option to generate pipeline artifacts.
- Review and customize the generated pipeline artifacts.
- Deploy to your target environment using the AWS CLI or AWS console.

Utility commands

The following additional commands help you maintain AWS App2Container in your environment.

[upgrade](#)

Run the **upgrade** command to upgrade your existing installation of App2Container. This command checks if there is a newer version of App2Container available, and automatically upgrades if doing so will not break backwards compatibility with previously generated container artifacts.

[upload-support-bundle](#)

Run the **upload-support-bundle** command for assistance from the AWS App2Container team for troubleshooting a command failure. This command securely uploads App2Container logs and supporting artifacts, and an optional message for your troubleshooting request to the AWS App2Container team.

app2container analyze command

Analyzes the specified application and generates a report.

Note

If the command fails, an error message is displayed in the console, followed by additional messaging to help you troubleshoot.

When you ran the **init** command, if you chose to automatically upload logs to App2Container support if an error occurs, App2Container notifies you of the success of the automatic upload of your application support bundle.

Otherwise, App2Container messaging directs you to upload application artifacts by running the [upload-support-bundle](#) command for additional support.

Syntax

```
app2container analyze --application-id id [--help]
```

Parameters and options

Parameters

--application-id *id*

The application ID (*required*). After you run the [inventory](#) command, you can find the application ID in the `inventory.json` file in one of the following locations:

- **Linux:** `/root/inventory.json`
- **Windows:** `C:\Users\Administrator\AppData\Local\.app2container-config\inventory.json`

Options

--help

Displays the command help.

Output

The **analyze** command creates files and directories for each application. Output varies slightly, depending on your application language and the application server operating system.

The application directory is created in the output location that you specified when you ran the **init** command. Each application has its own directory named for the application ID. The directory contains analysis output and editable application configuration files. The files are stored in subdirectories that match the application structure on the server.

An `analysis.json` file is created for the application that is specified in the `--application-id` parameter. The file contains information about the application found during analysis, and configurable fields for container settings. See [Configuring application containers](#), and choose the platform that your application container runs on for more information about configurable fields, and for an example of what the file looks like.

For .NET applications and Windows services, App2Container detects connection strings and produces the `report.txt` file. The report location is specified in the `analysis.json` file, in the `reportPath` attribute of the `analysisInfo` section. You can use this report to identify the changes that you need to make in application configuration files to connect your application container to new database endpoints, if needed. The report also contains the locations of other configuration files that might need changes.

Examples

Choose the operating system platform tab for the application server or worker machine where you run the command.

Linux

The following example shows the **analyze** command with the `--application-id` parameter and no additional options.

```
$ sudo app2container analyze --application-id java-tomcat-9e8e4799
✓ Created artifacts folder /root/app2container/java-tomcat-9e8e4799
✓ Generated analysis data in /root/app2container/java-tomcat-9e8e4799/analysis.json
Analysis successful for application java-tomcat-9e8e4799
Please examine the application analysis file at /root/app2container/java-tomcat-9e8e4799/analysis.json,
make appropriate edits and initiate containerization using "app2container
containerize --application-id java-tomcat-9e8e4799"
```

Windows

The following example shows the **analyze** command with the `--application-id` parameter and no additional options.

```
PS> app2container analyze --application-id iis-smarts-51d2dbf8
✓ Created artifacts folder C:\Users\Administrator\AppData\Local\app2container\iis-smarts-51d2dbf8
✓ Generated analysis data in C:\Users\Administrator\AppData\Local\app2container\iis-smarts-51d2dbf8\analysis.json
Analysis successful for application iis-smarts-51d2dbf8
Please examine the application analysis file at C:\Users\Administrator\AppData\Local\app2container\iis-smarts-51d2dbf8\analysis.json,
make appropriate edits and initiate containerization using "app2container
containerize --application-id iis-smarts-51d2dbf8"
```

app2container containerize command

When you run this command, it creates a Docker container image for your application. The is based on the parameters in the `analysis.json` file that is generated by the [analyze](#) command, along with any customizations you have made. By default, the image is pre-validated to ensure that the application container runs and returns a successful response, as expected.

See [Configuring application containers](#) for more information about configuring the `analysis.json` file.

Note

If the command fails, an error message is displayed in the console, followed by additional messaging to help you troubleshoot.

When you ran the **init** command, if you chose to automatically upload logs to App2Container support if an error occurs, App2Container notifies you of the success of the automatic upload of your application support bundle.

Otherwise, App2Container messaging directs you to upload application artifacts by running the [upload-support-bundle](#) command for additional support.

Syntax

```
app2container containerize [--application-id id | --input-archive extraction-file] [--no-validate] [--help]
```

Parameters and options

Parameters

--application-id *id*

The application ID (*required*). After you run the [inventory](#) command, you can find the application ID in the `inventory.json` file in one of the following locations:

- **Linux:** `/root/inventory.json`
- **Windows:** `C:\Users\Administrator\AppData\Local\.app2container-config\inventory.json`

--input-archive *extraction-file*

The file path or Amazon S3 key (for example, `s3://bucket/archive-key`) for the application archive. If you specify an application archive, the command downloads and opens the archive, and then builds the container image.

--profile *admin-profile*

Use this option to specify a *named profile* to run this command. For more information about named profiles in AWS, see [Named profiles](#) in the AWS Command Line Interface User Guide

Options**--build-only**

Builds container images based on the existing `Dockerfile` and artifacts.

--force

Bypasses the disk space prerequisite check.

--no-validate

Bypasses pre-validation of the generated container image.

--help

Displays the command help.

Output

This command generates a `Dockerfile`, a container image, and a `deployment.json` file that you can use with the [generate app-deployment](#) command.

It also generates a `Dockerfile.update` file that you can use to make updates to your containerized application. The [generate pipeline](#) command adds this `Dockerfile` to your CodeCommit repository and deploys updates to your CodePipeline infrastructure.

See [deployment.json file](#) for more information about configuration, and for an example of what the `deployment.json` file looks like.

Examples

Choose the operating system platform tab for the application server or worker machine where you run the command.

Linux

The following example shows the **containerize** command with the **--application-id** parameter and no additional options.

```
$ sudo app2container containerize --application-id java-tomcat-9e8e4799
✓ AWS pre-requisite check succeeded
✓ Docker pre-requisite check succeeded
✓ Extracted container artifacts for application
✓ Entry file generated
✓ Dockerfile generated under /root/app2container/java-tomcat-9e8e4799/Artifacts
✓ Generated dockerfile.update under /root/app2container/java-tomcat-9e8e4799/
Artifacts
✓ Generated deployment file at /root/app2container/java-tomcat-9e8e4799/
deployment.json
✓ Deployment artifacts generated.
✓ Pre-validation succeeded.
Containerization successful. Generated docker image java-tomcat-9e8e4799
```

You're all set to test and deploy your container image.

Next Steps:

1. View the container image with `\\"docker images\\"` and test the application.
2. When you're ready to deploy to AWS, please edit the deployment file as needed at `/root/app2container/java-tomcat-9e8e4799/deployment.json`.
3. Generate deployment artifacts using `app2container generate app-deployment --application-id java-tomcat-9e8e4799`

Please use `"docker images"` to view the generated container image.

The following example shows the **containerize** command with the **--input-archive** option.

```
$ sudo app2container containerize --input-archive /var/aws/java-tomcat-9e8e4799/  
java-tomcat-9e8e4799-extraction.tar
```

Windows

The following example shows the **containerize** command with the **--application-id** parameter and no additional options.

```
PS> app2container containerize --application-id iis-smarts-51d2dbf8
✓ AWS pre-requisite check succeeded
✓ Docker pre-requisite check succeeded
✓ Extracted container artifacts for application
✓ Entry file generated
✓ Dockerfile generated under C:\Users\Administrator\AppData\Local\app2container\iis-smarts-51d2dbf8\Artifacts
✓ Generated dockerfile.update under C:\Users\Administrator\AppData\Local\app2container\iis-smarts-51d2dbf8\Artifacts
✓ Generated deployment file at C:\Users\Administrator\AppData\Local\app2container\iis-smarts-51d2dbf8\deployment.json
Containerization successful. Generated docker image iis-smarts-51d2dbf8
```

You're all set to test and deploy your container image.

Next Steps:

1. View the container image with `"docker images"` and test the application.
 2. When you're ready to deploy to AWS, please edit the deployment file as needed at `C:\Users\Administrator\AppData\Local\app2container\iis-smarts-51d2dbf8\deployment.json`.
 3. Generate deployment artifacts using `app2container generate app-deployment --application-id iis-smarts-51d2dbf8`
- Please use `"docker images"` to view the generated container image.

The following example shows the **containerize** command with the **--input-archive** option.

```
PS> app2container containerize --input-archive archive C:\Users\Administrator\Downloads\iis-smarts-51d2dbf8.zip
✓ AWS pre-requisite check succeeded
✓ Docker pre-requisite check succeeded
✓ Dockerfile generated under C:\Users\Administrator\AppData\Local\app2container\iis-smarts-51d2dbf8\Artifacts
✓ Generated dockerfile.update under C:\Users\Administrator\AppData\Local\app2container\iis-smarts-51d2dbf8\Artifacts
✓ Generated deployment file at C:\Users\Administrator\AppData\Local\app2container\iis-smarts-51d2dbf8\deployment.json
Containerization successful. Generated docker image iis-smarts-51d2dbf8
```


You're all set to test and deploy your container image.

Next Steps:

1. View the container image with `"docker images"` and test the application.
2. When you're ready to deploy to AWS, please edit the deployment file as needed at `C:\Users\Administrator\AppData\Local\app2container\iis-smarts-51d2dbf8\deployment.json`.
3. Generate deployment artifacts using `app2container generate app-deployment --application-id iis-smarts-51d2dbf8`

To have gMSA related artifacts generated with CloudFormation, please edit `gMSAParams` inside deployment file.

Otherwise look at `C:\Users\Administrator\AppData\Local\app2container\iis-smarts-51d2dbf8\Artifacts\WindowsAuthSetupInstructions.md` for setup instructions on Windows Authentication

Please use `"docker images"` to view the generated container image.

app2container extract command

Generates an application archive for the specified application. Before you call this command, you must call the [analyze](#) command.

Note

If the command fails, an error message is displayed in the console, followed by additional messaging to help you troubleshoot.

When you ran the **init** command, if you chose to automatically upload logs to App2Container support if an error occurs, App2Container notifies you of the success of the automatic upload of your application support bundle.

Otherwise, App2Container messaging directs you to upload application artifacts by running the [upload-support-bundle](#) command for additional support.

Syntax

```
app2container extract --application-id id [--output s3] [--help]
```

Parameters and options

Parameters

--application-id *id*

The application ID (*required*). After you run the [inventory](#) command, you can find the application ID in the `inventory.json` file in one of the following locations:

- **Linux:** `/root/inventory.json`
- **Windows:** `C:\Users\Administrator\AppData\Local\.app2container-config\inventory.json`

--profile *admin-profile*

Use this option to specify a *named profile* to run this command. For more information about named profiles in AWS, see [Named profiles](#) in the AWS Command Line Interface User Guide

Options

--output *s3*

If specified, this option writes the archive file to the Amazon S3 bucket that you specified when you ran the **init** command.

--force

Bypasses the disk space prerequisite check.

--help

Displays the command help.

Output

This command creates an archive file. When you use the `--output s3` option, the archive is written to the Amazon S3 bucket that you specified when you ran the **init** command. Otherwise, the archive is written to the output location that you specified when you ran the **init** command.

Examples

Choose the operating system platform tab for the application server or worker machine where you run the command.

Linux

The following example shows the **extract** command with the `--application-id` parameter and no additional options.

```
$ sudo app2container extract --application-id java-tomcat-9e8e4799
✓ Extracted container artifacts for application
✓ Application archive file created at: /root/app2container/java-tomcat-9e8e4799/
  java-tomcat-9e8e4799-extraction.tar
Extraction successful for application java-tomcat-9e8e4799

Please transfer this tar file to your worker machine and run, "app2container
containerize --input-archive <extraction-tar-filepath>"
```

Windows

The following example shows the **extract** command with the `--application-id` parameter and no additional options.

```
PS> app2container extract --application-id iis-smarts-51d2dbf8
✓ Extracted container artifacts for application
Extraction successful for application iis-smarts-51d2dbf8
```

app2container generate app-deployment command

When you run this command, it generates the artifacts needed to deploy your application container in AWS. App2Container pre-fills key values in the artifacts based on your profile, the application analysis, your App2Container workflow, and best practices.

Note

For Windows applications, App2Container chooses the base image for your application container and Amazon ECS cluster, based on the worker machine or application server OS where you run the containerization command. Windows application containers running on Amazon EKS use Windows Server Core 2019 for the base image.

You have three options for deployment to your target container management environment, all of which use Amazon ECR as the container registry (Amazon ECS, Amazon EKS, or App Runner):

- You can use the `--deploy` option to deploy directly to your target environment. When your initial deployment uses this option, you can refresh your image by running the command with the `deploy` option again.
- You can review and customize deployment artifacts, and then deploy using the AWS CLI or AWS console.

This command accesses AWS resources to generate and deploy artifacts to your target environment. The IAM user with administrator access that you created during security setup is required to run the command with the `--deploy` option. See [Identity and access management in App2Container](#) for more information about setting up IAM users for App2Container.

The command uses the `deployment.json` file that is generated by the [containerize](#) command. You can edit the `deployment.json` file to specify parameters for your deployment, such as:

- An image repository name for Amazon ECR
- Task definition parameters for Amazon ECS
- The Kubernetes app name
- The App Runner service name

See [Configuring container deployment](#) for more information about configuring the `deployment.json` file.

Note

If the command fails, an error message is displayed in the console, followed by additional messaging to help you troubleshoot.

When you ran the **init** command, if you chose to automatically upload logs to App2Container support if an error occurs, App2Container notifies you of the success of the automatic upload of your application support bundle.

Otherwise, App2Container messaging directs you to upload application artifacts by running the [upload-support-bundle](#) command for additional support.

Syntax

```
app2container generate app-deployment --application-id id [--deploy] [--profile admin-profile] [--help]
```

Parameters and options

Parameters

--application-id *id*

The application ID (*required*). After you run the [inventory](#) command, you can find the application ID in the `inventory.json` file in one of the following locations:

- **Linux:** `/root/inventory.json`
- **Windows:** `C:\Users\Administrator\AppData\Local\.app2container-config\inventory.json`

--profile *admin-profile*

Use this option to specify a *named profile* to run this command. For more information about named profiles in AWS, see [Named profiles](#) in the AWS Command Line Interface User Guide

Options

--deploy

Use this option to deploy directly to your target container management environment (Amazon ECR with Amazon ECS, Amazon EKS, or App Runner).

Note

When you use the `--deploy` option to deploy directly to target environments, we recommend that you use the `--profile` option to specify a *named profile* that has elevated permissions.

--help

Displays the command help.

Output

You have two options for deploying a container to your target environment using the **generate app-deployment** command.

- You can use the `--deploy` option to deploy directly. This option also allows you to create deployments iteratively. For example, after you create an initial Amazon ECS deployment with the `--deploy` option, then you can rerun the command with `--deploy` to update your image in Amazon ECS.
- You can review and customize deployment artifacts, and then deploy using the AWS CLI or AWS console.

Generate container deployment artifacts for customization

```
app2container generate app-deployment --application-id id
```

To see the steps App2Container performs and the artifacts that it creates to generate an application deployment for your target container management service, choose the tab that matches your environment:

Amazon ECS

App2Container performs the following tasks and creates artifacts for deployment to Amazon ECS:

Note

You must set the `createEcsArtifacts` parameter in the `deployment.json` file to `true` to generate Amazon ECS artifacts. See [Configuring container deployment](#) for more information on how to configure the `deployment.json` file.

- Checks for AWS and Docker prerequisites.
- Creates an Amazon ECR repository.
- Pushes the container image to the Amazon ECR repository.
- Generates an Amazon ECS task definition template.
- Generates a `pipeline.json` file.

Amazon EKS

App2Container performs the following tasks and creates artifacts for deployment to Amazon EKS:

Note

You must set the `createEksArtifacts` parameter in the `deployment.json` file to `true` to generate Amazon EKS artifacts. See [Configuring container deployment](#) for more information on how to configure the `deployment.json` file.

- Checks for AWS and Docker prerequisites.
- Creates an Amazon ECR repository.
- Pushes the container image to the Amazon ECR repository.
- Generates a CloudFormation template (`eks-master.yaml`) that creates an EKS cluster, pulls your application container images from Amazon ECR, and deploys your application to the cluster.
- Generates Kubernetes manifests (`eks_deployment.yaml`, `eks_service.yaml`, and `eks_ingress.yaml`), for post-deployment customizations using a tool such as `kubectl`.
- Generates a `pipeline.json` file.

AWS App Runner

App2Container performs the following tasks and creates artifacts for deployment to AWS App Runner:

Note

You must set the `createAppRunnerArtifacts` parameter in the `deployment.json` file to `true` to generate App Runner artifacts. See [Configuring container deployment](#) for more information on how to configure the `deployment.json` file.

App Runner deployment is currently available for Linux applications only.

- Checks for AWS and Docker prerequisites.
- Creates an Amazon ECR repository.

- Pushes the container image to the Amazon ECR repository.
- Generates the `apprunner.yaml` CloudFormation template.
- Generates a `pipeline.json` file.

Deploy directly to target environments

```
app2container generate app-deployment --application-id id --deploy --profile admin-profile
```

When you run this command with the `--deploy` option, App2Container uses the same process to validate and customize your deployment resources as it does when you deploy manually. Additionally, it performs the following steps to complete the deployment:

- Uploads AWS CloudFormation resources to an Amazon S3 bucket, if configured.
- Creates a CloudFormation stack and deploys your application.

See [pipeline.json file](#) for more information about pipeline configuration, and for an example of the `deployment.json` file.

Examples

Choose the operating system platform tab for the application server or worker machine where you run the command.

Linux

The following example shows the `generate app-deployment` command with the `--application-id` parameter.

```
$ sudo app2container generate app-deployment --application-id java-tomcat-9e8e4799
✓ AWS prerequisite check succeeded
✓ Docker prerequisite check succeeded
✓ Processing application java-tomcat-9e8e4799...
✓ Created ECR Repository 123456789012.dkr.ecr.us-west-2.amazonaws.com/java-tomcat-9e8e4799
✓ Pushed docker image to ECR repository
✓ Created ECR repository 123456789012.dkr.ecr.us-west-2.amazonaws.com/java-tomcat-9e8e4799-fluent-bit
```



```

✓ Pushed docker image 123456789012.dkr.ecr.us-west-2.amazonaws.com/java-
tomcat-9e8e4799-fluent-bit:latest to ECR repository
✓ Local ECS Task Definition file created
✓ Uploaded CloudFormation resources to S3 Bucket: app2container-example
✓ Generated CloudFormation Master template at: /root/app2container/java-
tomcat-9e8e4799/EcsDeployment/ecs-master.yml
✓ Initiated CloudFormation stack creation. This may take a few minutes. To track
  progress, open the AWS CloudFormation console
✓ Deploying AWS CloudFormation Stack: <link to stack>
✓ Stack a2c-java-tomcat-9e8e4799-ECS deployed successfully!
✓ Updating service

```

Deployment successful for application java-tomcat-9e8e4799

Successfully created ECS infrastructure stack app2container-java-tomcat-9e8e4799-ECS.

The URL to your Load Balancer Endpoint is:

<your endpoint>.us-east-1.elb.amazonaws.com

The URL to your application log group on CloudWatch is: <log group link>>

Set up a pipeline for your application stack using app2container:

```
app2container generate pipeline --application-id java-tomcat-9e8e4799
```

The following example shows the **generate app-deployment** command with the **--application-id** parameter for an application that is deployed to AWS App Runner.

```

$ sudo app2container generate app-deployment --application-id java-tomcat-9e8e4799
✓ AWS pre-requisite check succeeded
✓ Docker pre-requisite check succeeded
✓ Created ECR repository 123456789012.dkr.ecr.us-west-2.amazonaws.com/java-
tomcat-9e8e4799 already
✓ Pushed docker image to 123456789012.dkr.ecr.us-west-2.amazonaws.com/java-
tomcat-9e8e4799:latest to ECR repository
✓ Generated AWS App Runner CloudFormation template at /root/app2container/java-
tomcat-9e8e4799/AppRunnerDeployment/apprunner.yml

```

CloudFormation templates and additional deployment artifacts generated successfully for application java-tomcat-9e8e4799

You're all set to use AWS CloudFormation to manage your application stack.

Next Steps:

1. Edit the CloudFormation template as necessary.

2. Create an application stack using the AWS CLI or the AWS Console. AWS CLI command:

```
aws cloudformation deploy --template-file /root/app2container/java-tomcat-9e8e4799/
AppRunnerDeployment/apprunner.yml --capabilities CAPABILITY_IAM --stack-name a2c-
java-tomcat-9e8e4799-AppRunner
```

3. Set up a pipeline for your application stack using app2container:
app2container generate pipeline --application-id java-tomcat-9e8e4799

The following example shows the **generate app-deployment** command with the **--application-id** parameter and the **--deploy** option.

```
$ sudo app2container generate app-deployment --deploy --application-id java-
tomcat-9e8e4799 --profile admin-profile
✓ AWS prerequisite check succeeded
✓ Docker prerequisite check succeeded
✓ Processing application java-tomcat-9e8e4799...
✓ Created ECR Repository 123456789012.dkr.ecr.us-west-2.amazonaws.com/java-
tomcat-9e8e4799
✓ Pushed docker image to ECR repository
✓ Created ECR repository 123456789012.dkr.ecr.us-west-2.amazonaws.com/java-
tomcat-9e8e4799-fluent-bit
✓ Pushed docker image 123456789012.dkr.ecr.us-west-2.amazonaws.com/java-
tomcat-9e8e4799-fluent-bit:latest to ECR repository
✓ Local ECS Task Definition file created
✓ Uploaded CloudFormation resources to S3 Bucket: app2container-example
✓ Generated CloudFormation Master template at: /root/app2container/java-
tomcat-9e8e4799/EcsDeployment/ecs-master.yml
✓ Updating CloudFormation stack
✓ Initiated CloudFormation stack creation. This may take a few minutes. To track
progress, open the AWS CloudFormation console
✓ Deploying AWS CloudFormation Stack: <link to stack>
✓ Stack a2c-java-tomcat-9e8e4799-ECS deployed successfully!
✓ Updating service
```

Deployment successful for application java-tomcat-9e8e4799

Successfully created ECS infrastructure stack app2container-java-tomcat-9e8e4799-ECS.

The URL to your Load Balancer Endpoint is:

<your endpoint>.us-east-1.elb.amazonaws.com

The URL to your application log group on CloudWatch is: <log group link>>

```
Set up a pipeline for your application stack using app2container:  
app2container generate pipeline --application-id java-tomcat-9e8e4799
```

The following example shows the **generate app-deployment** command with the `--application-id` parameter and the `--deploy` option for an application that is deployed to AWS App Runner.

```
$ sudo app2container generate app-deployment --application-id java-tomcat-9e8e4799  
--deploy  
✓ AWS pre-requisite check succeeded  
✓ Docker pre-requisite check succeeded  
✓ Created ECR repository 123456789012.dkr.ecr.us-west-2.amazonaws.com/java-  
tomcat-9e8e4799  
✓ Pushed docker image to 123456789012.dkr.ecr.us-west-2.amazonaws.com/java-  
tomcat-9e8e4799:latest to ECR repository  
✓ Generated AWS App Runner CloudFormation template at /root/app2container/java-  
tomcat-9e8e4799/AppRunnerDeployment/apprunner.yml  
  
Deployment successful for application java-tomcat-9e8e4799  
Access your newly deployed App Runner service at the following URL:  
https://xyz123abc4.us-west-2.awsapprunner.com  
Stack deployed successfully!  
  
Set up a pipeline for your application stack using app2container:  
app2container generate pipeline --application-id java-tomcat-9e8e4799
```

Windows

The following example shows the **generate app-deployment** command with the `--application-id` parameter.

```
PS> app2container generate app-deployment --application-id iis-smarts-51d2dbf8  
✓ AWS pre-requisite check succeeded  
✓ Docker pre-requisite check succeeded  
✓ Created ECR Repository  
✓ Registered ECS Task Definition with ECS  
✓ Uploaded CloudFormation resources to S3 Bucket: app2container\--testing  
✓ Generated CloudFormation Master template at: C:\Users\Administrator\AppData\Local  
\app2container\iis-smarts-51d2dbf8\EcsDeployment\ecs-master.yml  
CloudFormation templates and additional deployment artifacts generated successfully  
for application iis-smarts-51d2dbf8
```

You're all set to use AWS CloudFormation to manage your application stack.

Next Steps:

1. Edit the CloudFormation template as necessary.
2. Create an application stack using the AWS CLI or the AWS Console. AWS CLI command:

```
aws cloudformation deploy --template-file C:\Users\Administrator\AppData\Local\app2container\iis-smarts-51d2dbf8\EcsDeployment\ecs-master.yml --capabilities CAPABILITY_NAMED_IAM --stack-name app2container-iis-smarts-51d2dbf8-ECS
```

3. Set up a pipeline for your application stack using app2container:
app2container generate pipeline --application-id iis-smarts-51d2dbf8

The following example shows the **generate app-deployment** command with the `--application-id` parameter and the `--deploy` option.

```
PS> app2container generate app-deployment --deploy --application-id iis-smarts-51d2dbf8 --profile admin-profile
✓ AWS prerequisite check succeeded
✓ Docker prerequisite check succeeded
✓ Created ECR Repository
✓ Registered ECS Task Definition with ECS
✓ Uploaded CloudFormation resources to S3 Bucket: app2container-example
✓ Generated CloudFormation Master template at: C:\Users\Administrator\AppData\Local\app2container\iis-smarts-51d2dbf8\EcsDeployment\ecs-master.yml
✓ Initiated CloudFormation stack creation. This may take a few minutes. Please visit the AWS CloudFormation Console to track progress.
ECS deployment successful for application iis-smarts-51d2dbf8
```

The URL to your Load Balancer Endpoint is:

<your endpoint>.us-east-1.elb.amazonaws.com

Successfully created ECS stack app2container-iis-smarts-51d2dbf8-ECS. Check the AWS CloudFormation Console for additional details.

app2container generate pipeline command

When you run the **generate pipeline** command, it generates the artifacts that you need to create a CI/CD pipeline with CodePipeline, Jenkins, or Microsoft Azure DevOps services. Your application pipeline settings and deployment artifacts determine the artifacts that you create.

Note

For Windows applications, App2Container chooses the base image for your application container and Amazon ECS cluster, based on the worker machine or application server OS where you run the containerization command. Windows application containers running on Amazon EKS use Windows Server Core 2019 for the base image.

You have two options for creating your pipeline:

- You can use the `--deploy` option to create your pipeline directly.
- You can review and customize pipeline artifacts, and then create your pipeline, with the AWS CLI or the AWS Management Console for CodePipeline. You can also create your pipeline in the native environments for Jenkins or Microsoft Azure DevOps pipelines.

When the **generate pipeline** command generates artifacts and creates CI/CD pipelines, it accesses AWS resources, even if your application integrates with an external pipeline tool or service. App2Container needs administrator access to run the command with the `--deploy` option. For information on how to set up AWS Identity and Access Management (IAM) users for App2Container, see [Identity and access management in App2Container](#).

The **generate pipeline** command uses the `pipeline.json` file that App2Container generates when you run the [generate app-deployment](#) command. You can edit the `pipeline.json` file to specify your container repository and target environments for Amazon ECS, Amazon EKS, or App Runner. For more information on how to configure the `pipeline.json` file, see [Configuring container pipelines](#).

Note

If the command fails, an error message is displayed in the console, followed by additional messaging to help you troubleshoot.

When you ran the **init** command, if you chose to automatically upload logs to App2Container support if an error occurs, App2Container notifies you of the success of the automatic upload of your application support bundle.

Otherwise, App2Container messaging directs you to upload application artifacts by running the [upload-support-bundle](#) command for additional support.

Syntax

```
app2container generate pipeline --application-id id [--deploy] [--profile admin-profile] [--help]
```

Parameters and options

Parameters

--application-id *id*

The application ID (*required*). After you run the [inventory](#) command, you can find the application ID in the `inventory.json` file in one of the following locations:

- **Linux:** `/root/inventory.json`
- **Windows:** `C:\Users\Administrator\AppData\Local\.app2container-config\inventory.json`

--profile *admin-profile*

Use this option to specify a *named profile* to run this command. For more information about named profiles in AWS, see [Named profiles](#) in the AWS Command Line Interface User Guide

Options

--deploy

Use this option to create your CI/CD pipeline directly.

Note

When you use the `--deploy` option to create your CI/CD pipeline directly, we recommend that you use the `--profile` option to specify a *named profile* that has elevated permissions.

--help

Displays the command help.

Output

You have two options for creating your CI/CD pipeline using the **generate pipeline** command.

- You can use the `--deploy` option to create your pipeline directly.
- You can review and customize pipeline artifacts, and then create your pipeline, with the AWS CLI or the AWS Management Console for CodePipeline. You can also create your pipeline in the native environments for Jenkins or Microsoft Azure DevOps pipelines.

When you run the **generate pipeline** command, App2Container generates the following artifacts and performs the following tasks.

CodePipeline

Generates pipeline artifacts for customization

- Generates CI/CD artifacts **generate pipeline --application-id *id***
 - Checks for AWS and Docker prerequisites
 - Creates a CodeCommit repository, if one doesn't already exist
 - Generates a buildspec file
 - Generates CloudFormation templates for a two-step pipeline to commit and build your application

Creates pipeline directly with deploy option

- When you run this command with the `--deploy` option, App2Container uses the same process to validate and customize your pipeline resources as it does when you deploy manually. Then it uses the settings from the files that it generated to create the pipeline for you: **generate pipeline --application-id *id* --deploy --profile *admin-profile***
 - Performs all steps to validate and customize pipeline resources
 - Creates the CloudFormation stack for your pipeline

Jenkins

Generates pipeline artifacts for customization

- Generates CI/CD artifacts **generate pipeline --application-id *id***

- Checks for AWS and Docker prerequisites
- Creates a CodeCommit repository, if one doesn't exist already
- Generates the following files for your pipeline definition: the Jenkinsfile, and the config.xml file that you can use with the Jenkins REST API
- If your application runs on Amazon EKS, App2Container generates a CloudFormation template for a two-step pipeline to commit and build your application

Creates pipeline directly with deploy option

- When you run this command with the --deploy option, App2Container uses the same process to validate and customize your pipeline resources as it does when you deploy manually. App2Container then creates the pipeline for you with the settings from the files that it generates: **generate pipeline --application-id *id* --deploy --profile *admin-profile***
- Performs all steps to validate and customize pipeline resources
- Creates the pipeline in Jenkins, and starts the pipeline build

Azure DevOps

Generate pipeline artifacts for customization

- Generates CI/CD artifacts: **generate pipeline --application-id *id***
 - Checks for AWS, Microsoft Azure DevOps, and Docker prerequisites
 - Creates the Azure Repos Git repository, if it doesn't already exist
 - Commits updated files to the Azure Repos Git repository
 - Generates the following files for your pipeline definition: main.yaml, build.yaml, release.yaml, pre-req.sh (Linux) or pre-req.ps1 (Windows), and install-pre-req.sh

Creates pipeline directly with deploy option

- When you run this command with the --deploy option, App2Container uses the same process to validate and customize your pipeline resources as it does when you deploy manually. Then it uses the settings from the files that it generated to create the pipeline for you: **generate pipeline --application-id *id* --deploy --profile *admin-profile***.

- Performs all steps to validate and customize pipeline resources
- Uses the configuration in `pipeline.json` to create an Azure DevOps pipeline, and initiate an Azure DevOps pipeline build

Examples

To see examples of how to use the **generate pipeline** command, choose your target environment.

CodePipeline

Linux:

The following Linux example shows the **generate pipeline** command with the `--application-id` parameter that you use to create CodeCommit pipeline resources for your application.

```
$ sudo app2container generate pipeline --application-id java-tomcat-9e8e4799
✓ Created CodeCommit repository
✓ Generated buildspec file(s)
✓ Generated CloudFormation templates
✓ Committed files to CodeCommit repository
Pipeline resource template generation successful for application java-
tomcat-9e8e4799
```

You're all set to use AWS CloudFormation to manage your pipeline stack.

Next Steps:

1. Edit the CloudFormation template as necessary.
2. Create a pipeline stack using the AWS CLI or the AWS Console. AWS CLI command:

```
aws cloudformation deploy --template-file /root/app2container/java-tomcat-9e8e4799/
Artifacts/Pipeline/CodePipeline/ecs-pipeline-master.yml --capabilities
CAPABILITY_NAMED_IAM --stack-name app2container-java-tomcat-9e8e4799-ecs-pipeline-
stack
```

The following Linux example shows the **generate pipeline** command with the `--application-id` parameter and the `--deploy` option that you use to create a CodeCommit pipeline for your application.

```
$ sudo app2container generate pipeline --deploy --application-id java-
tomcat-9e8e4799
```

```

✓ Generated buildspec file(s)
✓ Generated CloudFormation templates
✓ Committed files to CodeCommit repository
✓ Initiated CloudFormation stack creation. This may take a few minutes. Please visit
  the AWS CloudFormation Console to track progress.
✓ Deployed pipeline through CloudFormation
Pipeline deployment successful for application --application-id java-tomcat-9e8e4799

Successfully created AWS CodePipeline stack 'app2container---application-id java-
tomcat-9e8e4799-ecs-pipeline-stack' for application. Check the AWS CloudFormation
Console for additional details.

```

The following Linux example shows the **generate pipeline** command with the `--application-id` parameter and the `--deploy` option that you use to create a pipeline for an application that runs on AWS App Runner.

```

$ sudo app2container generate pipeline --deploy --application-id java-
tomcat-9e8e4799
✓ Created CodeCommit repository
✓ Generated buildspec file(s)
✓ Generated CloudFormation templates
✓ Committed files to CodeCommit repository
✓ Initiated CloudFormation stack creation. This may take a few minutes. To track
  progress, open the AWS CloudFormation console.
✓ Deployed pipeline through CloudFormation
Pipeline deployment successful for application java-tomcat-9e8e4799

Successfully created AWS CodePipeline stack 'a2c---application-id java-
tomcat-9e8e4799-ecs-pipeline-stack' for application. Check the AWS CloudFormation
Console for additional details.

```

Windows:

The following Tools for Windows PowerShell example shows the **generate pipeline** command with the `--application-id` parameter that you use to create CodeCommit pipeline resources for your application.

```

PS> app2container generate pipeline --application-id iis-smarts-51d2dbf8
✓ Created CodeCommit repository
✓ Generated buildspec file(s)
✓ Generated CloudFormation templates
✓ Committed files to CodeCommit repository

```

Pipeline resource template generation successful for application --application-id *iis-smarts-51d2dbf8*

You're all set to use AWS CloudFormation to manage your pipeline stack.

Next Steps:

1. Edit the CloudFormation template as necessary.
2. Create a pipeline stack using the AWS CLI or the AWS Console. AWS CLI command:

```
aws cloudformation deploy --template-file C:\Users\Administrator\AppData\Local
\app2container\--application-id iis-smarts-51d2dbf8\Artifacts\Pipeline\CodePipeline
\ecs-pipeline-master.yml --capabilities CAPABILITY_NAMED_IAM --stack-name
app2container---application-id iis-smarts-51d2dbf8-652becbe-ecs-pipeline-stack
```

The following Tools for Windows PowerShell example shows the **generate pipeline** command with the --application-id parameter and the --deploy option that you use to create a CodeCommit pipeline for your application.

```
PS> app2container generate pipeline --deploy --application-id iis-smarts-51d2dbf8
√ Generated buildspec file(s)
√ Generated CloudFormation templates
√ Committed files to CodeCommit repository
√ Initiated CloudFormation stack creation. This may take a few minutes. Please visit
the AWS CloudFormation Console to track progress.
√ Deployed pipeline through CloudFormation
Pipeline deployment successful for application --application-id iis-smarts-51d2dbf8

Successfully created AWS CodePipeline stack 'app2container---application-id iis-
smarts-51d2dbf8-ecs-pipeline-stack' for application. Check the AWS CloudFormation
Console for additional details.
```

Jenkins

Linux:

The following Linux example shows the **generate pipeline** command with the --application-id parameter that you use to create Jenkins pipeline resources for your application.

```
$ sudo app2container generate pipeline --application-id java-tomcat-9e8e4799
√ Discovered existing CodeCommit repository
√ Generated Jenkins pipeline configuration file
```

```

✓ Generated Jenkinsfile
✓ Committed files to source repository
Pipeline resource template generation successful for application java-
tomcat-9e8e4799

```

You're all set to use Jenkins to manage your pipeline.

Next Steps:

1. Edit the Jenkinsfile as necessary.
2. Create a Jenkins Pipeline using the Jenkins REST API or the Jenkins Dashboard.
Jenkins API command:

```

curl -k -XPOST https://ec2-1-234-567-890.<Region>.compute.amazonaws.com:8443/
createItem?name=a2c-java-tomcat-9e8e4799-eks-pipeline-stack -u
a2c:1164afa1fe791a4c86fd3117d7bc5d93e2 --data-binary @/home/ubuntu/app2container/
java-tomcat-9e8e4799/Artifacts/Pipeline/Jenkins/config.xml -H "Content-Type:text/
xml"

```

The following Linux example shows the **generate pipeline** command with the **--application-id** parameter and the **--deploy** option that you use to create a Jenkins pipeline for your application.

```

$ sudo app2container generate pipeline --deploy --application-id java-
tomcat-9e8e4799
✓ Discovered existing CodeCommit repository
✓ Generated Jenkins pipeline configuration file
✓ Generated Jenkinsfile
✓ Committed files to source repository
✓ Initiated Jenkins pipeline creation
✓ Deployed pipeline through Jenkins
Pipeline deployment successful for application java-tomcat-9e8e4799

Successfully created Jenkins Pipeline 'a2c-java-tomcat-9e8e4799-eks-pipeline' for
application. Started a build of the pipeline.
Build link: https://ec2-1-234-567-890.<Region>.compute.amazonaws.com:8443/job/a2c-
java-tomcat-9e8e4799-eks-pipeline/1

```

Windows:

The following Tools for Windows PowerShell example shows the **generate pipeline** command with the **--application-id** parameter and the **--deploy** option that you use to create a Jenkins pipeline for your application.

```
PS> app2container generate pipeline --deploy --application-id iis-smarts-51d2dbf8
✓ Validated Jenkins Nodes and Labels
✓ Generated Jenkins pipeline configuration file
✓ Generated Jenkinsfile
✓ Committed files to source repository
✓ Initiated Jenkins pipeline creation
✓ Deployed pipeline through Jenkins
Pipeline deployment successful for application iis-smarts-51d2dbf8

Successfully created Jenkins Pipeline 'iis-smarts-51d2dbf8-eks-pipeline' for
application. Started a build of the pipeline.
Build link: https://ec2-1-234-567-890.<Region>.compute.amazonaws.com:8443/job/iis-
smarts-51d2dbf8-eks-pipeline/1
```

Azure DevOps

Linux:

The following Linux example shows the **generate pipeline** command with the **--application-id** parameter that you use to create Microsoft Azure DevOps pipeline resources for your application.

```
$ sudo app2container generate pipeline --application-id java-tomcat-9e8e4799
# Discovered existing Azure repository
# Discovered existing Azure branch
# Generated pre-requisite installation scripts
# Generated pipeline definition files
# Committed artifacts to Microsoft Azure DevOps repository
Pipeline resource template generation successful for application java-
tomcat-9e8e4799
```

You're all set to use pipeline definition files in `/root/app2container/java-tomcat-9e8e4799/Artifacts/Pipeline/AzureDevOps` to create your Azure DevOps pipeline.

Next Steps:

1. Edit the pipeline definition files as necessary.
2. Created a new Azure git repository at https://dev.azure.com/a2c-azure-org/a2c-project/_git/a2c-java-tomcat-9e8e4799
3. Go to your Microsoft Azure DevOps web console https://dev.azure.com/a2c-azure-org/a2c-project/_build and click on "New Pipeline".
4. For Repositories select "Azure Repos Git" and select the repo with name `a2c-java-tomcat-9e8e4799`
5. For "Configure your pipeline" step choose "Existing Azure Pipelines YAML file"

6. In the options for "branch" select main and for "path" select /pipeline.yaml
7. Click "continue" and then click "Run"

The following Linux example shows the **generate pipeline** command with the `--application-id` parameter and the `--deploy` option that you use to create a Microsoft Azure DevOps pipeline for your application.

```
$ sudo app2container generate pipeline --deploy --application-id java-  
tomcat-9e8e4799  
# Discovered existing Azure repository  
# Discovered existing Azure branch  
# Generated pre-requisite installation scripts  
# Generated pipeline definition files  
# Committed artifacts to Microsoft Azure DevOps repository  
# Initiated Microsoft Azure DevOps pipeline creation  
# Deployed pipeline through Microsoft Azure DevOps  
Pipeline deployment successful for application java-tomcat-9e8e4799  
  
Successfully created and ran Microsoft Azure DevOps Pipeline 'a2c-java-  
tomcat-9e8e4799-pipeline' for the application, url: https://dev.azure.com/a2c-azure-  
org/a2c-project/_build?definitionId=152
```

Windows:

The following Windows example shows the **generate pipeline** command with the `--application-id` parameter that you use to create Microsoft Azure DevOps pipeline resources for your application.

```
PS> app2container generate pipeline iis-smarts-51d2dbf8  
# Discovered existing Azure repository  
# Discovered existing Azure branch  
# Generated pre-requisite installation scripts  
# Generated pipeline definition files  
# Committed artifacts to Microsoft Azure DevOps repository  
Pipeline resource template generation successful for application iis-smarts-51d2dbf8  
  
You're all set to use pipeline definition files in C:\Users\Administrator\AppData  
\Local\app2container\iis-smarts-51d2dbf8\Artifacts\Pipeline\AzureDevOps to create  
your Azure DevOps pipeline.  
  
Next Steps:  
1. Edit the pipeline definition files as necessary.
```

2. Created a new Azure git repository at https://dev.azure.com/a2c-azure-org/a2c-project/_git/a2c-iis-smarts-51d2dbf8
3. Go to your Microsoft Azure DevOps web console https://dev.azure.com/a2c-azure-org/a2c-project/_build and click on "New Pipeline".
4. For Repositories select "Azure Repos Git" and select the repo with name a2c-iis-smarts-51d2dbf8
5. For "Configure your pipeline" step choose "Existing Azure Pipelines YAML file"
6. In the options for "branch" select main and for "path" select /pipeline.yaml
7. Click "continue" and then click "Run"

The following Windows example shows the **generate pipeline** command with the `--application-id` parameter and the `--deploy` option that you use to create a Microsoft Azure DevOps pipeline for your application.

```
PS> app2container generate pipeline --deploy iis-smarts-51d2dbf8
# Discovered existing Azure repository
# Discovered existing Azure branch
# Generated pre-requisite installation scripts
# Generated pipeline definition files
# Committed artifacts to Microsoft Azure DevOps repository
# Initiated Microsoft Azure DevOps pipeline creation
# Deployed pipeline through Microsoft Azure DevOps
Pipeline deployment successful for application iis-smarts-51d2dbf8

Successfully created and ran Microsoft Azure DevOps Pipeline 'a2c-iis-smarts-51d2dbf8-pipeline' for the application, url: https://dev.azure.com/a2c-azure-org/a2c-project/_build?definitionId=151
```

app2container help command

Lists the commands for App2Container, grouped into the phases where they would normally run.

Note

Commands are shown in alphabetical order within the phases where they run. For example, in the *Analyze* phase, you would run the **inventory** command first, then the **analyze** command. Utility commands are included after the containerization phases.

Syntax

```
app2container help
```

Parameters and options

None

Output

The list of app2container commands.

Examples

app2container help

App2Container is an application from Amazon Web Services (AWS), that provides commands to discover and containerize applications.

Commands

Getting Started

<code>init</code>	Sets up workspace for artifacts
-------------------	---------------------------------

Analyze

<code>analyze</code>	Analyzes the selected application to identify dependencies required for containerization
<code>inventory</code>	Lists all applications that can be containerized

Transform

<code>containerize</code>	Generates Dockerfile, container images, and deployment metadata
<code>extract</code>	Creates an archive of application artifacts for containerization

Deploy

<code>generate</code>	Generates ECS, EKS, or Pipeline artifacts
-----------------------	---

Settings

<code>upgrade</code>	Upgrades app2container CLI to latest version
<code>upload-support-bundle</code>	Uploads user's app2container logs and supporting artifacts to the support team

Flags

<code>--debug</code>	enable debug logging
<code>-h, --help</code>	help for app2container
<code>--version</code>	version for app2container

app2container init command

The **init** command performs one-time initialization tasks for App2Container. This interactive command prompts for the information required to set up the local App2Container environment. Run this command before you run any other App2Container commands.

Note

If the command fails, an error message is displayed in the console, followed by additional messaging to help you troubleshoot.

When you ran the **init** command, if you chose to automatically upload logs to App2Container support if an error occurs, App2Container notifies you of the success of the automatic upload of your application support bundle.

Otherwise, App2Container messaging directs you to upload application artifacts by running the [upload-support-bundle](#) command for additional support.

Syntax

```
app2container init [--advanced] [--help]
```

Parameters and options

Options

--advanced

This option allows you to use features that are in the experimental phase, if any exist.

--help

Displays the command help.

Output

The **init** command prompts you for the information that it needs for initialization.

You must provide a local directory for application containerization artifacts. Ensure that only authorized users can access the local directory. If you do not specify a local directory, one is created for you at the default output location. The default locations are as follows:

- **Linux:** /root/app2container
- **Windows:** C:\Users\Administrator\AppData\Local\app2container

You can optionally provide an Amazon S3 bucket for application containerization artifacts. If you choose to set up an Amazon S3 bucket, you must ensure that only authorized users can access the bucket. We recommend that you use server-side encryption for your bucket. See [Protecting data using server-side encryption](#) in the *Amazon Simple Storage Service User Guide* for more information about how to set it up.

You can optionally upload logs and command-generated artifacts automatically to App2Container support when an app2container command crashes or encounters internal errors. Log files are retained for 90 days.

You can optionally consent to allow App2Container to collect and export the following metrics to AWS each time that you run an **app2container** command:

- Host OS name
- Host OS version
- Application stack type
- Application stack version
- JRE version (Linux only, for Java applications)
- App2Container CLI version
- Command that ran
- Command status
- Command duration
- Command features and flags
- Command errors

- Container base image

Examples

Choose the operating system platform tab for the application server or worker machine where you run the command.

Linux

The following example shows the **init** command with no additional options.

```
$ sudo app2container init
Please enter a workspace directory path to use for artifacts[default: /root/
app2container]:
Please enter an AWS Profile to use. (The same can be configured with 'aws configure
--profile <name>')[default: default]:
Please provide an S3 bucket to store application artifacts (Optional):
Automatically upload logs and App2Container generated artifacts on crashes and
internal errors? (Y/N):
Please confirm permission to report usage metrics to AWS (Y/N)[default: y]:
Would you like to enforce the use of only signed images using Docker Content Trust
(DCT)? (Y/N)[default: n]:
All application artifacts will be created under the above workspace. Please ensure
that the folder permissions are secure.
Init configuration saved
```

The following example shows the **init** command with the **--advanced** option and default values.

```
PS> sudo app2container init --advanced
Please enter a workspace directory path to use for artifacts[default: /root/
app2container]:
Please enter an AWS Profile to use. (The same can be configured with 'aws configure
--profile <name>')[default: default]:
Please provide an S3 bucket to store application artifacts (Optional):
Automatically upload logs and App2Container generated artifacts on crashes and
internal errors? (Y/N):
Please confirm permission to report usage metrics to AWS (Y/N)[default: y]:
Would you like to enforce the use of only signed images using Docker Content Trust
(DCT)? (Y/N)[default: n]:
Would you like to enable experimental features? (Y/N)[default: n]:
```

All application artifacts will be created under the above workspace. Please ensure that the folder permissions are secure.
Init configuration saved

Windows

The following example shows the **init** command with no additional options.

```
PS> app2container init
Please enter a workspace directory path to use for artifacts[default: C:\Users
\Administrator\AppData\Local\app2container]:
Please enter an AWS Profile to use. (The same can be configured with 'aws configure
--profile <name>')[default: default]:
Please provide an S3 bucket to store application artifacts (Optional):
Automatically upload logs and App2Container generated artifacts on crashes and
internal errors? (Y/N):
Please confirm permission to report usage metrics to AWS (Y/N)[default: y]:
Would you like to enforce the use of only signed images using Docker Content Trust
(DCT)? (Y/N)[default: n]:
All application artifacts will be created under the above workspace. Please ensure
that the folder permissions are secure.
Init configuration saved
```

The following example shows the **init** command with the **--advanced** option and default values.

```
PS> app2container init --advanced
Please enter a workspace directory path to use for artifacts[default: C:\Users
\Administrator\AppData\Local\app2container]:
Please enter an AWS Profile to use. (The same can be configured with 'aws configure
--profile <name>')[default: default]:
Please provide an S3 bucket to store application artifacts (Optional):
Automatically upload logs and App2Container generated artifacts on crashes and
internal errors? (Y/N):
Please confirm permission to report usage metrics to AWS (Y/N)[default: y]:
Would you like to enforce the use of only signed images using Docker Content Trust
(DCT)? (Y/N)[default: n]:
Please enter if we can enable checking for upgrades automatically (Y/N)[default:
y]:
Would you like to enable experimental features? (Y/N)[default: n]:
All application artifacts will be created under the above workspace. Please ensure
that the folder permissions are secure.
```

Init configuration saved

app2container inventory command

Records all Java or .Net processes (Linux) or all IIS websites and Windows services (Windows) that are running on the application server.

Syntax

```
app2container inventory --type [iis | service | java | dotnet] [--nofilter] [--help]
```

Parameters and options

Parameters

--type [*iis* | *service* | *java* | *dotnet*]

Use this parameter to specify the application type (*required*), as follows.

- For .NET applications running on Windows, you can specify an IIS web application (*iis*), or a Windows service (*service*).
- For Java applications running on Linux, you must specify *java*.
- For .NET applications running on Linux, you must specify *dotnet*.

Options

--nofilter

For applications running on Windows, this option prevents App2Container from filtering out default system services when building the inventory output. This can be used for complex Windows .NET applications that have dependent web apps that need to be included in the container.

--help

Displays the command help.

Output

Information about the Java processes, .NET applications, or IIS websites is saved to the `inventory.json` file in one of the following locations:

- **Linux:** `/root/inventory.json`
- **Windows:** `C:\Users\Administrator\AppData\Local\.app2container-config\inventory.json`

The application ID that is used by other App2Container commands is the key for each application object in the JSON file. The application objects are slightly different depending on your application language and the application server operating system. Choose the operating system platform for your application in the Examples section to see the differences.

Examples

Expand the section that matches the operating system platform for the application server or worker machine where you run the command.

Linux examples

Each Java process or ASP.NET application running on Linux has a unique application ID (for example, `java-tomcat-9e8e4799`, or `dotnet-single-c2930d3132`). You can use this application ID with other AWS App2Container commands. Inventory information is saved to `/root/inventory.json`.

Java

The following example shows the **inventory** command with results for Java processes running on Linux, with no additional options.

```
$ sudo app2container inventory
{
  "java-jboss-5bbe0bec": {
    "processId": 27366,
    "cmdline": "java ...",
    "applicationType": "java-jboss"
  },
  "java-tomcat-9e8e4799": {
```

```

        "processId": 2537,
        "cmdline": "/usr/bin/java ...",
        "applicationType": "java-tomcat"
    }
}

```

ASP.NET

The following example shows the **inventory** command with results for .NET applications running on Linux, with no additional options.

```

$ sudo app2container inventory
{
  "dotnet-single-c2930d3132": {
    processId": 1,
    "cmdline": "./MyCoreWebApp.3.1 ...",
    "applicationType": "dotnet-single",
    "webApp": ""
  },
  "dotnet-generic-a27b2829": {
    processId": 2,
    "cmdline": "./MyCoreWebApp.3.1 ...",
    "applicationType": "dotnet-generic",
    "webApp": ""
  }
}

```

Windows examples

Each IIS website has a unique application ID (for example, iis-smarts-51d2dbf8). You can use this application ID with other AWS App2Container commands. Inventory information is saved to C:\Users\Administrator\AppData\Local\.app2container-config\inventory.json.

The following example shows the **inventory** command with results for .NET applications running in IIS on Windows, with no additional options.

```

PS> app2container inventory
{
  "iis-smarts-51d2dbf8": {
    "siteName": "Default Web Site",
    "bindings": "http/*:80:,net.tcp/808:*",

```

```
    "applicationType": "iis",
    "discoveredWebApps": []
  },
  "iis-smart-544e2d61": {
    "siteName": "smart",
    "bindings": "http/*:82:",
    "applicationType": "iis",
    "discoveredWebApps": []
  },
  "service-colorwindowsservice-69f90194": {
    "serviceName": "colorwindowsservice",
    "applicationType": "service"
  }
}
```

app2container remote analyze command

Run this command from a worker machine to analyze the specified application on the target application server, and generate a report. The target application server is specified by its IP address or Fully Qualified Domain Name (FQDN).

Note

If the command fails, an error message is displayed in the console, followed by additional messaging to help you troubleshoot.

When you ran the **init** command, if you chose to automatically upload logs to App2Container support if an error occurs, App2Container notifies you of the success of the automatic upload of your application support bundle.

Otherwise, App2Container messaging directs you to upload application artifacts by running the [upload-support-bundle](#) command for additional support.

Syntax

```
app2container remote analyze --application-id id --target IP/FQDN [--help]
```


Parameters and options

Parameters

--application-id *id*

The application ID (*required*). After you run the [remote inventory](#) command, you can find the application ID in the `inventory.json` file in one of the following locations:

- **Linux:** `<workspace>/remote/<target server IP or FQDN>/inventory.json`
- **Windows:** `<workspace>\remote\<target server IP or FQDN>\.app2container-config\inventory.json`

--target *IP/FQDN*

Specifies the IP address or FQDN of the application server targeted for the inventory (*required*).

Options

--help

Displays the command help.

Output

The **remote analyze** command creates files and directories on the worker machine for the specified application on the target application server. Each application has its own directory, named for the application ID. Output varies slightly, depending on your application language and the application server operating system.

The application directory contains analysis output and editable application configuration files. The files are stored in subdirectories that match the application structure on the application server.

An `analysis.json` file is created for the application that is specified in the `--application-id` parameter. The file contains information about the application found during analysis, and configurable fields for container settings. See [Configuring application containers](#), and choose the platform that your application container runs on for more information about configurable fields, and for an example of what the file looks like.

For .NET applications and Windows services, App2Container detects connection strings and produces the `report.txt` file. The report location is specified in the `analysis.json` file, in

the `reportPath` attribute of the `analysisInfo` section. You can use this report to identify the changes that you need to make in application configuration files to connect your application container to new database endpoints, if needed. The report also contains the locations of other configuration files that might need changes.

Examples

Choose the operating system platform tab for the application server or worker machine where you run the command.

Linux

The following example shows the **remote analyze** command with the `--target` and `--application-id` parameters and no additional options.

```
$ sudo app2container remote analyze --target 192.0.2.0 --application-id java-tomcat-9e8e4799
```

Analysis successful for application java-tomcat-9e8e4799

Next Steps:

1. View the application analysis file at `<workspace>/remote/<target server IP or FQDN>/java-tomcat-9e8e4799/analysis.json`.
2. Edit the application analysis file as needed.
3. Start the extraction process using the following command: `app2container remote extract --target 192.0.2.0 java-tomcat-9e8e4799`

Windows

The following example shows the **remote analyze** command with the `--target` and `--application-id` parameters and no additional options.

```
PS> app2container remote analyze --target 192.0.2.0 --application-id iis-smarts-51d2dbf8
```

Analysis successful for application iis-smarts-51d2dbf8;

Next Steps:

1. View the application analysis file at `<workspace>\remote\<target server IP or FQDN>/iis-smarts-51d2dbf8/analysis.json`.
2. Edit the application analysis file as needed.
3. Start the extraction process using the following command: `app2container remote extract --target 192.0.2.0 iis-smarts-51d2dbf8`

app2container remote configure command

Run this command from a worker machine to configure the connections needed to run remote workflows on application servers. This interactive command prompts for the required information for each application server that you enter, or you can provide a JSON input file with your connection information by specifying the `--input-json` parameter when you run the command.

Note

For the remote configure command prompts, if you specify the Fully Qualified Domain Name (FQDN), the server IP address is optional and is not used by App2Container.

Syntax

```
app2container remote configure [--input-json myhosts.json] [--help]
```

Parameters and options

Parameters

`--input-json`

Uses the provided JSON file as input to configure connections to application servers for the worker machine to run remote commands.

Options

`--help`

Displays the command help.

Input

To see the input file format, choose the system platform that matches your configuration. For key/value pairs that do not apply to your configuration, set string values to an empty string.

Linux remote hosts file

The Linux `remote_hosts.json` file contains an array of Linux platform hosts, with connection information. The key for each host is the host IP address or FQDN, with an array of strings for the connection information. Each host includes the following content:

- **Fqdn** (string, conditionally required) – the fully qualified domain name of the host, used as the identifier for connecting. *If an IP address is used as the host identifier, this must be empty. If the FQDN has a value, the IP address is ignored.*
- **Ip** (string, conditionally required) – the IP address of the host, used as the identifier for connecting. *Required if the FQDN is empty.*
- **SecretArn** (string, required) – the Amazon Resource Name (ARN) that identifies the Secrets Manager secret to use for credentials.
- **AuthMethod** (string, required) – the authentication method used to connect to the host. *Valid values include "cert" and "key".*

The following example shows a `remote_hosts.json` file for a Java application running on Linux.

```
{
  "10.10.10.10": {
    "Fqdn": "",
    "Ip": "10.10.10.10",
    "SecretArn": "arn:aws:secretsmanager:us-west-2:123456789012:secret:linux-cert-Abcdef",
    "AuthMethod": "cert"
  },
  "myhost.mydomain.com": {
    "Fqdn": "myhost.mydomain.com",
    "Ip": "",
    "SecretArn": "arn:aws:secretsmanager:us-west-2:987654321098:secret:linux-cert-Ghijkl",
    "AuthMethod": "key"
  }
}
```

Windows remote hosts file

The Windows `remote_hosts.json` file contains an array of Windows Server platform hosts, with connection information. The key for each host is the host IP address or FQDN, with an array of strings for the connection information. Each host includes the following content:

- **fqdn** (string, conditionally required) – the fully qualified domain name of the host, used as the identifier for connecting. *If an IP address is used as the host identifier, this must be empty. If the FQDN has a value, the IP address is ignored.*
- **ip** (string, conditionally required) – the IP address of the host, used as the identifier for connecting. *Required if the FQDN is empty.*
- **secretArn** (string, required) – the Amazon Resource Name (ARN) that identifies the Secrets Manager secret to use for credentials.

The following example shows a `remote_hosts.json` file for a .NET application running on Windows Server.

```
{
  "10.10.10.10": {
    "fqdn": "",
    "ip": "10.10.10.10",
    "secretArn": "arn:aws:secretsmanager:us-
west-2:123456789012:secret:windows-cred-Abcdef"
  }
}
```

Output

This command does not produce a configurable output file. For troubleshooting purposes, or if you need to verify what was entered during the interactive command dialog, you can find the entries in the `remote_hosts.json` file by searching the folder structure on the server where you ran the command.

Examples

Choose the operating system platform tab for the application server or worker machine where you run the command.

Linux

The following example shows the **remote configure** command with no additional options.

```
$ sudo app2container remote configure
Server IP address: 10.10.10.10
```

```
Server FQDN (Fully Qualified Domain Name):
Authentication method to be used key/cert: cert
Secret ARN for remote connection credentials: arn:aws:secretsmanager:us-
west-2:123456789012:secret:linux-cert-Abcdef
Continue to configure servers? (y/N)[default: n]: y
Server IP address:
Server FQDN (Fully Qualified Domain Name): fqdn2
Authentication method to be used key/cert: key
Secret ARN for remote connection credentials: arn:aws:secretsmanager:us-
west-2:987654321098:secret:linux-cert-Ghijkl
Continue to configure servers? (y/N)[default: n]: n
```

Windows

The following example shows the **remote configure** command with no additional options.

```
PS> app2container remote configure
Server IP address: 10.10.10.10
Server FQDN (Fully Qualified Domain Name):
Authentication method to be used key/cert: cert
Secret ARN for remote connection credentials: arn:aws:secretsmanager:us-
west-2:123456789012:secret:windows-cred-Abcdef
Continue to configure servers? (y/N)[default: n]: y
Server IP address:
Server FQDN (Fully Qualified Domain Name): fqdn2
Authentication method to be used key/cert: key
Secret ARN for remote connection credentials: arn:aws:secretsmanager:us-
west-2:987654321098:secret:windows-cred-Ghijkl
Continue to configure servers? (y/N)[default: n]: n
```

app2container remote extract command

Run this command from a worker machine to generate an application archive for the specified application on the target application server. The target application server is specified by its IP address or Fully Qualified Domain Name (FQDN). Before you call this command, you must call the [remote analyze](#) command.

Note

If the command fails, an error message is displayed in the console, followed by additional messaging to help you troubleshoot.

When you run the **init** command, if you chose to automatically upload logs to App2Container support if an error occurs, App2Container notifies you of the success of the automatic upload of your application support bundle. Otherwise, App2Container messaging directs you to upload application artifacts by running the [upload-support-bundle](#) command for additional support.

Syntax

```
app2container remote extract --application-id id --target IP/FQDN [--help]
```

Parameters and options

--application-id *id*

The application ID (*required*). After you run the [remote inventory](#) command, you can find the application ID in the `inventory.json` file in one of the following locations:

- **Linux:** `<workspace>/remote/<target server IP or FQDN>/inventory.json`
- **Windows:** `<workspace>\remote\<target server IP or FQDN>\.app2container-config\inventory.json`

--target *IP/FQDN*

Specifies the IP address or FQDN of the application server targeted for the inventory (*required*).

Options

--help

Displays the command help.

Output

This command creates an archive file. The archive is written to the output location that you specified when you ran the **init** command.

Examples

Choose the operating system platform tab for the application server or worker machine where you run the command.

Linux

The following example shows the **remote extract** command with the `--target` and `--application-id` parameters and no additional options.

```
$ sudo app2container extract --target 192.0.2.0 --application-id java-  
tomcat-9e8e4799  
Extraction successful for application java-tomcat-9e8e4799  
Next Steps:  
1. Please initiate containerization using "app2container containerize --input-  
archive <workspace>/remote/<target server IP or FQDN>/java-tomcat-9e8e4799/java-  
tomcat-9e8e4799-extraction.tar"
```

Windows

The following example shows the **remote extract** command with the `--target` and `--application-id` parameters and no additional options.

```
PS> app2container extract --target 192.0.2.0 --application-id iis-smarts-51d2dbf8  
Extraction successful for application iis-smarts-51d2dbf8  
  
Next Steps:  
1. Please initiate containerization using "app2container containerize --input-  
archive <workspace>\remote\<target server IP or FQDN>/iis-smarts-51d2dbf8/iis-  
smarts-51d2dbf8.zip"
```

app2container remote inventory command

Run this command from a worker machine to retrieve an inventory of all Java or .Net processes (Linux) or all IIS websites and Windows services (Windows) that are running on the application server specified in the `--target` parameter. The target application server is specified by its IP address or Fully Qualified Domain Name (FQDN). The inventory details are captured in the `inventory.json` file and stored on the worker machine under the target server folder.

Syntax

```
app2container remote inventory --target IP/FQDN --type [iis | service | java | dotnet]  
[--nofilter] [--help]
```

Parameters and options

Parameters

--target *IP/FQDN*

Specifies the IP address or FQDN of the application server targeted for the inventory (*required*).

--type [*iis* | *service* | *java* | *dotnet*]

Use this parameter to specify the application type (*required*), as follows.

- For .NET applications running on Windows, you can specify an IIS web application (*iis*), or a Windows service (*service*).
- For Java applications running on Linux, you must specify *java*.
- For .NET applications running on Linux, you must specify *dotnet*.

Options

--nofilter

For applications running on Windows, this option prevents App2Container from filtering out default system services when building the inventory output. This can be used for complex Windows .NET applications that have dependent web apps that need to be included in the container.

--help

Displays the command help.

Output

Information about the Java processes, .NET applications, or IIS websites is saved to `inventory.json` in one of the following locations:

- **Linux:** `<workspace>/remote/<target server IP or FQDN>/inventory.json`

- **Windows:** <workspace>\remote*<target server IP or FQDN>*\.app2container-config\inventory.json

The application ID that is used by other App2Container commands is the key for each application object in the JSON file. The application objects are slightly different depending on your application language and the application server operating system. Choose the operating system platform for your application in the Examples section to see the differences.

Examples

Expand the section that matches the operating system platform for the worker machine where you run the command.

Linux examples

Each Java process or ASP.NET application running on Linux has a unique application ID (for example, java-tomcat-9e8e4799, or dotnet-single-c2930d3132). You can use this application ID with other AWS App2Container commands. Inventory information is saved to /root/inventory.json.

Java

The following example shows the **remote inventory** command with results for Java processes running on Linux, with no additional options.

```
$ sudo app2container remote inventory --target IP/FQDN
: Retrieving inventory from remote server 192.0.2.0
√ Server inventory has been stored under <workspace>/remote/<target server IP or FQDN>/inventory.json
Remote inventory retrieved successfully
```

Sample inventory data:

```
{
  "java-jboss-5bbe0bec": {
    "processId": 27366,
    "cmdline": "java ...",
    "applicationType": "java-jboss"
  },
  "java-tomcat-9e8e4799": {
    "processId": 2537,
```

```

        "cmdline": "/usr/bin/java ...",
        "applicationType": "java-tomcat"
    }
}

```

ASP.NET

The following example shows the **remote inventory** command with results for .NET applications running on Linux, with no additional options.

```

$ sudo app2container remote inventory --target IP/FQDN
: Retrieving inventory from remote server 192.0.2.0
√ Server inventory has been stored under <workspace>/remote/<target server IP or FQDN>/inventory.json
Remote inventory retrieved successfully

```

Sample inventory data:

```

{
  "dotnet-single-c2930d3132": {
    "processId": 1,
    "cmdline": "./MyCoreWebApp.3.1 ...",
    "applicationType": "dotnet-single",
    "webApp": ""
  },
  "dotnet-generic-a27b2829": {
    "processId": 2,
    "cmdline": "./MyCoreWebApp.3.1 ...",
    "applicationType": "dotnet-generic",
    "webApp": ""
  }
}

```

Windows examples

Each IIS website has a unique application ID (for example, iis-smarts-51d2dbf8). You can use this application ID with other AWS App2Container commands. Inventory information is saved to C:\Users\Administrator\AppData\Local\.app2container-config\inventory.json.

The following example shows the **remote inventory** command with results for .NET applications running in IIS on Windows, with no additional options.

```
PS> app2container remote inventory --target IP/FQDN
```

```
: Retrieving inventory from remote server 192.0.2.0
```

```
✓ Server inventory has been stored under <workspace>\remote\<target server IP or FQDN>\inventory.json
```

```
Remote inventory retrieved successfully
```

Sample inventory data:

```
{
  "iis-smarts-51d2dbf8": {
    "siteName": "Default Web Site",
    "bindings": "http/*:80:,net.tcp/808:*",
    "applicationType": "iis",
    "discoveredWebApps": []
  },
  "iis-smart-544e2d61": {
    "siteName": "smart",
    "bindings": "http/*:82:",
    "applicationType": "iis",
    "discoveredWebApps": []
  },
  "service-colorwindowsservice-69f90194": {
    "serviceName": "colorwindowsservice",
    "applicationType": "service"
  }
}
```

app2container upgrade command

Run this command to upgrade your existing installation of App2Container.

If a newer version of AWS App2Container will break backwards compatibility with previously generated container artifacts when you do an upgrade, the upgrade command notifies you and requests permission to continue. If you choose to continue with the upgrade, you will be required to restart any ongoing analysis and containerization workflows for your applications.

Syntax

```
app2container upgrade [--help]
```

Options

--help

Displays the command help.

Output

Console output is included in the Examples section for this command.

Examples

Choose the operating system platform tab for the application server or worker machine where you run the command.

Linux

Run the command shown below to upgrade your existing App2Container for Linux.

```
$ sudo app2container upgrade
Using version 1.0
Version 1.1 available for download
Starting Download...
Starting Installation...
Installation successful!
```

Windows

Run the command shown below to upgrade your existing App2Container for Windows.

```
PS> app2container upgrade
Using version 0.0
Version 2.0 available for download Starting Download...
Starting Installation...Installation successful!
```

app2container upload-support-bundle command

For assistance with troubleshooting, run this command to securely upload App2Container logs and supporting artifacts to the AWS App2Container support team. The following list shows the types of files that you can upload with the **upload-support-bundle** command:

- App2Container logs
- The `analysis.json` file
- The `Dockerfile`
- The `deployment.json` file
- The `EcsDeployment.yml` or `ecs-master.yml` deployment artifacts

Syntax

```
app2container upload-support-bundle [--application-id id] [--support-message "message"]  
[--help]
```

Options

`--application-id` *id*

The application ID (*required*). After you run the [inventory](#) command, you can find the application ID in the `inventory.json` file in one of the following locations:

- **Linux:** `/root/inventory.json`
- **Windows:** `C:\Users\Administrator\AppData\Local\.app2container-config\inventory.json`

`--support-message`

Include a message for the App2Container support team with your bundle.

`--help`

Displays the command help.

Output

Console output is included in the Examples section for this command.

Examples

Choose the operating system platform tab for the application server or worker machine where you run the command.

Linux

Run the following command to upload a support bundle from a Linux operating system, including the application ID and a message for the support team.

```
$ sudo app2container upload-support-bundle --application-id java-tomcat-9e8e4799 --  
support-message "I ran into an issue during deployment ..."  
Support Message: I ran into an issue during deployment ...  
[displays while bundle is uploading] Uploading logs and supporting artifacts to  
App2Container support  
Support bundle upload successful
```

Windows

Run the following command to upload a support bundle from a Windows operating system, including the application ID and a message for the support team.

```
PS> app2container upload-support-bundle --application-id iis-smarts-51d2dbf8 --  
support-message "I ran into an issue during deployment ..."  
Support Message: I ran into an issue during deployment ...  
[displays while bundle is uploading] Uploading logs and supporting artifacts to  
App2Container support  
Support bundle upload successful
```

Troubleshooting App2Container issues

The following documentation can help you troubleshoot problems that you might have with the App2Container CLI.

Contents

- [Access App2Container logs on your server](#)
- [Access application logs inside of a running container](#)
- [AWS resource creation fails for the generate command](#)
- [Troubleshoot Java applications on Linux](#)
- [Troubleshoot .NET applications on Windows](#)
- [Troubleshoot generate pipeline build for Jenkins](#)

Access App2Container logs on your server

A common first step in troubleshooting issues with any application is reviewing application logs. App2Container logs contain a history of the information and error messages that are produced by the commands that you run. If you opted out of metrics during initialization, the metrics messages are also logged in the local application log file.

Review log files in one of the following locations, depending on where you are running the command that needs troubleshooting:

Application logs

- **Linux:** /root/app2container/log/app2container.log
- **Windows:** C:\Users\Administrator\AppData\Local\app2container\log\app2container.log

Upgrade logs

- **Linux:** /usr/local/app2container/log/app2container_upgrade.log
- **Windows:** C:\Users\Administrator\app2container\log\app2container_upgrade.log

If there is more than one log file, it means that the first log file reached its maximum size, and a new log file was created to continue logging. Choose the most recent log file to troubleshoot.

Access application logs inside of a running container

You can access application logs on your running container by running a command shell from the container host that attaches to your container. Choose the tab that matches your container operating system to see the command.

Linux

From the host server, run an interactive bash shell on your running container.

```
$ docker exec -it container-id bash
```

Using the bash shell, you can then navigate to the location where your application logs are stored.

Windows

From the host server, run an interactive PowerShell session attached to your running container.

```
PS> docker exec -it container-id powershell.exe
```

Using the PowerShell session, you can then navigate to the location where your application logs are stored.

To look up Docker commands, use the Docker command line reference. See [Use the Docker command line](#).

AWS resource creation fails for the generate command

Description

When you run the **generate app-deployment** or **generate pipeline** command, you receive an error message saying AWS resource creation has failed.

Cause

App2Container requires permission to access and create AWS resources when it generates and deploys application containers or pipelines. If the permission has not been configured in your IAM policy, or if you are using the default AWS profile for a command using the `--deploy` option, the command will fail.

Solution

Verify your IAM resources and AWS profile settings and adjust as necessary, depending on the command that failed and the details shown in your error message. For more information and instructions about how to set up IAM resources for App2Container, see [Identity and access management in App2Container](#).

Troubleshoot Java applications on Linux

This section contains issues you might have with using App2Container for Java applications running on Linux servers.

Note

The App2Container **containerize** command creates a Dockerfile, along with other deployment artifacts. To reduce container sizes for Java applications, the Dockerfile installs the Java Runtime Environment (JRE) on your container by default. If your application requires the Java Development Kit (JDK) instead, you can edit your Dockerfile to change it. Your container size will be affected.

Edit the Dockerfile in your application directory (*<app2container workspace>/<app ID>/Artifacts/Dockerfile*) as follows:

Configure Dockerfile to use JDK

1. Locate the line that installs the JRE and change it to install the JDK. The change that you make depends on your container base image. For example, if your container uses an Ubuntu or Debian base image, you would change the package name from `openjdk-<version>-jre` to `openjdk-<version>-jdk`.
2. Re-run the **containerize** command, using the `--build-only` option, which instructs App2Container to recreate the container using the existing build artifacts.

```
$ sudo app2container containerize --application-id java-tomcat-9e8e4799 --  
build-only
```

Application container image size is very large

Description

Your application container image is much larger than expected.

Cause

The application container image includes a kernel image with the application bits layered on top. The size of the image depends on both the size of the container operating system and the size of the application.

To catch all potential dependencies for Java applications on Linux that are not using JBoss or Tomcat frameworks, the container initially includes everything except the files that are already included in the kernel image.

Solution

Follow these steps to reduce the size of your application container image.

1. Use the `appExcludedFiles` section in your `analysis.json` file to exclude specific file and directory paths from the containerization process, and save the file when you are done.
2. Run the [containerize](#) command again to create a new application container image with the updates that you specified.

You can repeat this process as needed to further reduce the size.

Error: Insufficient disk space

Description

When you run the `containerize` command, it fails with the following error message: Error: Insufficient disk space.

Cause

For Java applications running on Linux, App2Container calculates the disk space that is required to generate the application container, and produces this error message if there is not enough free space. The calculation includes the space needed for the application archive (including all non-system files on the server), plus the space needed for **docker build** actions.

Solution

The error message generated by the **containerize** command includes the estimated space it needs to run successfully. There are many ways to address an insufficient space issue on Linux.

One way to ensure that your **containerize** command runs successfully is to reduce the size of the container that you are creating. Follow these steps to reduce the size of your application container image.

1. Use the `appExcludedFiles` section in your `analysis.json` file to exclude specific file and directory paths from the containerization process, and save the file when you are done.
2. Run the [containerize](#) command again to create a new application container image with the updates that you specified.

You can repeat this process as needed to further reduce the size.

Troubleshoot .NET applications on Windows

This section contains issues you might have with using App2Container for .NET applications running in IIS on Windows servers.

Application container image size is very large

Description

Your application container image is much larger than expected.

Cause

The application container image includes a kernel image with the application bits layered on top. The size of the image depends on both the size of the container operating system and the size of the application. The Windows Server Core image can be quite large, especially for versions prior to Windows Server Core 2019.

Solution

We recommend that you use Windows Server Core 2019 for your container operating system to create the smallest base container size possible.

Follow these steps to reduce the size of your application container image if you are not currently using Windows Server Core 2019 as your base image. To ensure that you get the correct version, specify the version tag as shown below. The repository for Windows base images does not support the concept of "latest" to target the most recent image version.

1. Use the `containerBaseImage` section in your `analysis.json` file to target the Windows Server Core 2019 base image tagged as `ltsc2019` and save the file when you are done.

The `containerBaseImage` value includes both the image name and the `ltsc2019` tag, separated by a colon (:). For example: `"containerBaseImage": "mcr.microsoft.com/dotnet/framework/aspnet:4.7.2-windowsservercore-ltsc2019"`.

2. Run the [containerize](#) command again to create a new application container image. It will use the container operating system image that you specify in the `containerBaseImage` of your `analysis.json` file to build a new application container image.

Troubleshoot generate pipeline build for Jenkins

This section contains issues you might have for your App2Container pipeline build that is configured for Jenkins. As with any other troubleshooting scenario, the first step should be to review your application logs. For more information, see [Access App2Container logs on your server](#).

Unable to negotiate with **x.x.x.x** port 22: no matching host key type found. (Windows)

Description

Your Jenkins Windows agent, or the Jenkins server, if it's running on Windows is not able to connect to your CodeCommit repository to perform Git operations.

Cause

Pre-conditions:

- You are running the **generate pipeline** command with the `--deploy` option to deploy a Jenkins pipeline.

- You are deploying the pipeline for a Windows application.
- You have CodeCommit configured as your Jenkins code repository.

When the **generate pipeline** command runs with the `--deploy` option, App2Container detects the code repository that you have configured for Jenkins. If CodeCommit is your Jenkins code repository, App2Container generates an SSH-RSA key for the Jenkins server or Windows agent to connect to the CodeCommit repository for Git operations.

If OpenSSH on your Jenkins server or Windows agent is not configured to accept RSA-encrypted keys, your **generate pipeline** build fails with an error message that is similar to this example:

```
Unable to negotiate with 11.22.333.444 port 22: no matching host key type found. Their offer: ssh-rsa
```

Solution

To configure OpenSSH in your Jenkins environment, add the following configuration to your user profile `%userprofile%/.ssh/config` on Jenkins Windows agents, and also on the Jenkins server, if it is running on Windows.

```
Host git-codecommit.*.amazonaws.com
  HostkeyAlgorithms +ssh-rsa
  PubkeyAcceptedKeyTypes +ssh-rsa
```

Note

- If your Jenkins server is running on Windows, update the user profile that you ran the Jenkins setup with.
- For Jenkins Windows agents, update the user profile that has your connection to the Jenkins server configured.

Release notes for AWS App2Container

The following table describes the release history for AWS App2Container in descending date order.

Release date	Version	Details
August 29, 2025	1.51	Updated the AWS Lambda function for Amazon Elastic Container Service (Amazon ECS) deployment with gMSA to use Node.js 22.
May 28, 2025	1.50	- Changed the default deployed version of Amazon Elastic Kubernetes Service (Amazon EKS) clusters to 1.32. - Removed support for Windows Server 2019-based applications in the <code>generate pipeline</code> command for Azure DevOps due to Microsoft's deprecation of Windows Server 2019 hosted agents. - Updated allowed Amazon Elastic Compute Cloud (Amazon EC2) instance types when deploying to Amazon Elastic Container Service (Amazon ECS) and Amazon Elastic Kubernetes Service (Amazon EKS). - Added bug fixes, including the following: - Fixed an issue where deploying to Amazon Elastic Container Service (Amazon ECS) would fail when using float values (for example, 2.0) for vCPU.
April 21, 2025	1.49	Added bug fixes, including the following: - Fixed an issue where ECS tasks were not tagged - Fixed an issue where remote configuration failed when using wildcard certificates.

Release date	Version	Details
December 17, 2024	1.48	Starting with App2Container version 1.48, only Windows Server 2016 and later Windows Server operating systems are supported. For Windows Server 2008 and 2012, you can use App2Container version 1.47. - Added bug fixes, including the following: - Fixed an issue where the <code>containerize</code> command would fail due to spaces in Windows service names.
November 22, 2024	1.47	This is the last version that is compatible with Windows Server 2008 and Windows Server 2012. If necessary for your operating system version, you can download the App2Container version 1.47 binary file for Windows and the corresponding file hash to verify the integrity of the download. For more information on installing App2Container, see Step 1: Install App2Container. - Removed IIS as a prerequisite for Windows services inventory. - Added support for .NET 9 applications. - Added bug fixes, including the following: - Improved error handling when integrating with FireLens for deployments to Amazon Elastic Container Service (Amazon ECS).

Release date	Version	Details
October 10, 2024	1.46	- Updated the Lambda function for Amazon Elastic Container Service (Amazon ECS) deployments with gMSA to use Node.js 18. - Added bug fixes, including the following: - Improved the handling of file system changes during the creation of application archives. - Improved error handling for updating the Amazon ECS service when you deploy to Amazon ECS.
September 11, 2024	1.45	- Launch templates are now used when you deploy to Amazon Elastic Compute Cloud (Amazon EC2) container instances for Amazon Elastic Container Service (Amazon ECS). - Added bug fixes, including the following: - Addressed an issue when files aren't found with the <code>extract</code> command. - Improved detection for subnets when you deploy to Amazon ECS in an existing VPC. - Improved error handling for invalid CPU and memory configurations when you deploy to Amazon ECS on AWS Fargate. - Addressed an issue where removing some sections from the <code>pipeline.json</code> file could cause App2Container to crash.
July 26, 2024	1.44	Enhanced logging for Jenkins input validation during pipeline deployment.

Release date	Version	Details
July 11, 2024	1.43	Improved Jenkins plugin detection during pipeline deployment.
July 3, 2024	1.42	Improved connection string detection for Windows applications.
June 14, 2024	1.41	Added support for Windows Server 2022 as a container host operating system.
June 10, 2024	1.40	• Changed the default deployed version of Amazon Elastic Kubernetes Service (Amazon EKS) clusters to 1.30. • Updated the analyze command to use CentOS Stream 9 as the default base image on CentOS application server. • Added bug fixes, including the following: • Improved error handling when customized base images are used in the containerization process.
May 21, 2024	1.39	This release improves the detection of connection strings and configuration files for IIS .NET applications and Windows services.
May 2, 2024	1.38	Added bug fixes, including the following: • Improved error handling for insufficient disk space errors encountered during the containerization phase. • Improved detection for Windows authentication during the analysis phase.

Release date	Version	Details
March 8, 2024	1.37	Updated additional AWS Lambda functions for Amazon Elastic Kubernetes Service (Amazon EKS) deployments to use Python 3.9 and runtimes provided by Amazon Linux 2 (AL2). • Added bug fixes, including the following: • Improved error handling when customized base images are used in the containerization process.
February 22, 2024	1.36	Added bug fixes, including the following: • Fixed an issue where App2Container version 1.35 couldn't run or be installed on Windows Server 2008.

Release date	Version	Details
February 14, 2024	1.35	• Updated the AWS CloudFormation custom resources to use Node.js 18. • Updated the AWS Lambda functions for Amazon Elastic Container Service (Amazon ECS) and Amazon Elastic Kubernetes Service (Amazon EKS) deployments to use Python 3.9. • Clarified error messaging when containerization fails due to the host operating system and container operating system being incompatible. • Added detection for the command used to invoke Windows services when you use the analyze command. • Added bug fixes, including the following: • Fixed an issue where Windows deployments to Amazon ECS on Amazon EC2 could fail. • Fixed an issue where the analyze command would fail due to having wildcard expressions which contain square bracket characters ([or]) in the application's name. For more information, see about_Wildcards in the Microsoft documentation.
November 16, 2023	1.34	Added support for .NET 8 applications.

Release date	Version	Details
October 20, 2023	1.33	- Added bug fixes, including the following: - Fixed an issue where the <code>generate app-deployment</code> command could fail if the container image was untagged. - Fixed an issue where the <code>inventory</code> command could encounter a runtime exception.
August 28, 2023	1.32	- Adopted a new telemetry endpoint to collect metrics, logs, and command-generated artifacts. - Updated the <code>SectionForMetricsService</code> and <code>SectionForUploadSupportBundleService</code> statements in the example IAM policies for Amazon EKS, Amazon ECS, and App Runner to support the new telemetry endpoint. For more information, see Example IAM policies. - Added bug fixes, including the following: - Fixed an issue where generating a pipeline in AWS CodePipeline for Windows applications could fail. - Fixed an issue where containerization doesn't wait long enough for health checks to pass. - Improved error handling for missing IAM permissions during deployment. - Fixed a bug with handling quotation marks in the paths for Windows services.

Release date	Version	Details
August 1, 2023	1.31	• Changed the Windows build instances for AWS CodePipeline integration to use ECS-optimized AMIs. • Added bug fixes, including the following: • Fixed an issue with generated pipelines in AWS CodePipeline where builds could fail. • Fixed an issue where deployments to Amazon Elastic Kubernetes Service would fail in some newer AWS regions. • Fixed an issue where Application Load Balancers would not be created when deploying to Amazon EKS.

Release date	Version	Details
June 30, 2023	1.30	• Added support to deploy Microsoft Azure DevOps pipelines from servers with operating systems other than Ubuntu. • Changed the default cluster instance type to <code>c5.4xlarge</code> when deploying to Amazon ECS with EC2 instances. • Changed the instance type to <code>t3.medium</code> for Windows container pipelines created in AWS CodePipeline. • Changed the default AWS CodeBuild environment type to Amazon Linux 2023 for all container pipelines created in CodePipeline. • Added bug fixes, including the following: • Fixed a bug where the <code>remote analyze</code> command was unable to access files on the application server if they required a second hop to a network drive. • Fixed a bug where remote workflows weren't successfully completing with Microsoft Azure DevOps pipelines, even with Windows Server 2019 worker machines. • Fixed a bug where deployments to Amazon EKS could indicate they were successful when the application container didn't function properly. • Improved error handling for when the operating system or architecture of the container image doesn't match that of the server running the Docker engine.

Release date	Version	Details
May 26, 2023	1.29	• Changed the default deployed version of Amazon Elastic Kubernetes Service (Amazon EKS) to 1.26. • Added bug fixes, including the following: • Fixed a bug where the <code>containerize</code> command failed to handle Linux applications with the <code>\n</code>, <code>\r</code>, <code>\t</code>, <code>\b</code>, <code>\f</code>, or <code>\v</code> characters defined in environment variables. • Fixed a bug where deploying to Amazon EKS was failing when the existing stack could not be updated. • Fixed a bug where some App2Container commands would not complete for machines without internet connectivity.

Release date	Version	Details
May 4, 2023	1.28	• Changed the Amazon Simple Storage Service (Amazon S3) object key prefix that App2Container uses when you upload AWS CloudFormation templates to Amazon S3 with either the <code>generate app-deployment</code> or <code>generate pipeline</code> command. • Changed the instance metadata version of Windows build instances to IMDSv2 when they integrate with AWS CodePipeline. • Added support for Amazon Linux 2023. • Changed the default installation from Java Development Kit (JDK) 11 to Amazon Corretto 11 when App2Container generates a <code>Dockerfile</code>. • Added bug fixes, including the following: • Removed the unnecessary creation of an Amazon EC2 key pair when you deploy to Amazon Elastic Container Service (Amazon ECS) on AWS Fargate. • Improved how errors are handled when AWS CloudFormation fails to create and update stacks. • Fixed a bug that caused the analysis phase to fail during discovery of database connection files over networked drives. • Fixed a bug that prevented detection of all ports for Windows services. • Changed how quotation marks are handled in the paths for Windows services. •

Release date	Version	Details
		Fixed a bug that caused Windows services to use an incorrect health check. • Changed how arguments are handled for the executable that a Windows services is using. • Fixed a bug that caused the generated name of S3 buckets for CodePipeline integration to be too long. • Fixed a bug that caused Azure DevOps pipelines for Linux applications to use an unsupported agent image. • Fixed a bug where a named profile would be required for the <code>containerize</code> command, even when you don't use an input archive from Amazon S3. • Fixed a bug so that you can use the <code>generate app-deployment</code> command while you specify the <code>--deploy</code> flag to prevent failure of S3 bucket validation.
April 20, 2023	1.27	Announcement App2Container version 1.27 has been rolled back due to an issue with analyzing applications on Linux. We recommend that you run the <code>app2container upgrade</code> command to install the latest generally available version. For more information, see app2container upgrade command.

Release date	Version	Details
February 23, 2023	1.26	- Changed the default instance metadata version to IMDSv2 for cluster instances when you create Amazon ECS and Amazon EKS clusters. - Introduced the generation of machine-based application IDs to ensure uniqueness of the application. - Added a bug fix, as follows: - Fixed a containerization issue when the application runs as a non-root user.
January 11, 2023	1.25	- Changed default .NET Framework base image to version 4.8. - Added bug fixes, including the following: - Clarified error messaging when the containerize command fails while using the --build-only flag. - Clarified error messaging related to analyze command failures in Windows. - Fixed an issue causing a duplicate s3 pipeline bucket name across different accounts with same app id.
December 11, 2022	1.24	- Added support for .NET 7 application. - Added bug fixes, including the following: - Fixed an issue when containerizing an application running as a different user on Linux.

Release date	Version	Details
November 15, 2022	1.23	- Enhanced Amazon EKS deployment features to help you configure the <code>deployment.json</code> file with existing clusters more conveniently, as follows: - Added properties for CPU and memory limit configuration. - Added properties for ingress configuration with an AWS Application Load Balancer or using NGINX with an Network Load Balancer. - Added properties to set up DNS records for the deployed application. - Added properties to use with an AWS Certificate Manager (ACM) certificate for HTTPS deployments. - To better reflect inclusive language, the default git branch name is now main. - Added bug fixes, including the following: - Fixed a duplicate port issue with the analyze command. - Fixed an issue with uploading large log files with the upload-support-bundle command. - Fixed an issue with Firelens logging for Linux ECS deployments.

Release date	Version	Details
October 10, 2022	1.22	Added bug fixes, including the following: • Fixed an issue with the analyze command for complex Windows applications that include an IIS application that resides on a shared network drive. • Changed the containerize command so that it no longer adds IIS web applications to the default application pool when it generates containers for complex Windows applications. • The remote analyze and remote extract commands no longer require AWS Tools for Windows PowerShell to run on the remote application server. • Fixed a deployment validation issue for Amazon EKS application containers that specify an existing VPC. • Fixed an issue where App2Container doesn't update an application deployment if the AWS CloudFormation stack is not in a complete state.
September 16, 2022	1.21	• Changed new Amazon EKS deployment version default to Amazon EKS version 1.22. • Added bug fixes, including the following: • Fixed an Amazon EKS deployment issue where Windows applications could be stuck in the ContainerCreating state. • Fixed an issue that can happen when App2Container deploys new application containers to Amazon EKS within an existing VPC.

Release date	Version	Details
September 9, 2022	1.20	Added bug fixes, including the following: Fixed an analyze command issue for Java Tomcat applications on Linux that have a non-default web application directory.
August 30, 2022	1.19	Added bug fixes, including the following: - Fixed an issue with Windows applications where the local input archive path fails validation and causes an internal error. - Fixed an issue with the <code>--input-json</code> parameter for the remote configure command that affected Linux and Windows platforms. - Added validation for the generate pipeline command that returns an error if the <code>Dockerfile.update</code> file is missing.

Release date	Version	Details
August 2, 2022	1.18	- Enhanced Azure DevOps support so that Amazon EKS deployments work without a kubectl-specific IAM user. - Added bug fixes, including the following: - Improved run time exception handling for AWS Secrets Manager secret retrieval. - Fixed a containerization issue for complex applications that use a cooperating application. - Fixed a Windows application issue that occurs when the file path is longer than 260 characters. - Fixed an issue that can occur when App2Container archives input from S3 URLs during containerization. - Fixed a containerization issue for Java 17 applications on Amazon Linux 2.
June 20, 2022	1.17	- Support Microsoft Azure DevOps as a deployment pipeline. - Added bug fixes, including the following: - Added an error message response when you attempt to install App2Container from a version of the PowerShell command line interface that App2Container doesn't support. To run commands in PowerShell, App2Container requires version 5.0 or above. - Fixed an issue when Kubernetes authenticates with the latest version of the AWS CLI.

Release date	Version	Details
May 9, 2022	1.16	• Added bug fixes, including the following: • Addressed an issue that occurs when analyzing ports used by the Windows Service. • Addressed an issue with application analysis when some Linux applications might not progress during a port connectivity test. • Addressed an issue with large Windows applications that caused the Windows CodePipeline to fail.
April 14, 2022	1.15	• Optimized AWS App2Container installer size. • Deprecated MD5 checksum validation for the App2Container installer. • Made AWS profile setup optional during init. • Added bug fixes, including the following: • Added validation for S3 bucket name during init.

Release date	Version	Details
March 31, 2022	1.14	Added bug fixes, including the following: • Fixed an issue that caused the containerize command with the <code>--build-only</code> option specified to fail. • Fixed issue with missing path references that caused analyze command to fail. • Corrected upload issue for the support bundle when a panic error occurs. • Fixed null pointer issues for the following commands: • generate app-deployment • remote configure • Clarified messaging for AWS CloudFormation access denied error.
March 2, 2022	1.13	Added bug fixes, including the following: • Fixed an issue that caused the containerize command to fail when the <code>--build-only</code> option was specified. • Fixed JDK version check failures. • Fixed an issue with environment variables used in the application path for Windows servers. • Fixed incorrect container image tag assignment for some operating systems. • Fixed an issue that reported success metrics when the <code>remote</code> command failed.

Release date	Version	Details
February 16, 2022	1.12	Fixed an issue that caused remote execution commands to fail in Windows.
February 09, 2022	1.11	Added bug fixes, including the following: • Use CentOS Stream as the base image for applications containerized on the CentOS platform. • Added an explicit check for the tar command on Linux. • Fixed an issue related to checking application images in Amazon ECR. • Fixed an issue related to early container removal for pre-validation.
January 14, 2022	1.10	• Removed dependencies on Amazon ECR and Docker plugins for Jenkins pipelines. • Added bug fixes, including the following: • Added validation for PowerShell Version less than 5.0. • Fixed analyze command to handle paths with %SystemDrive% .

Release date	Version	Details
December 8, 2021	1.9	• Added end-to-end workflow support for ASP.NET Core applications running on Linux, including single file applications. • Added bug fixes, including the following: • Fixed <code>publicApp</code> parameter issue in the <code>deployment.json</code> file. • Added user validation for existing Jenkins pipelines. • Fixed an issue with incremental deployments on failed CloudFormation stacks. • Fixed intermittent failure of Amazon EKS deployments on Windows.

Release date	Version	Details
November 24, 2021	1.8	- Added support for Jenkins pipeline deployment. - Added support for incremental deployments of service and infrastructure changes to Amazon ECS, Amazon EKS, and App Runner auto-deployments. - Added automatic pre-validation for the containerize command, with the <code>--no-validate</code> option to skip that step. - Added bug fixes, including the following: - Fixed an issue with the upgrade command that caused the download to time-out. - Fixed an Amazon ECS deployment issue caused by the wrong version of Windows 2016 being used as the base AMI. - Fixed an issue with selecting the container image for Amazon ECS when Firelens logging is enabled. - Fixed issue with application analysis that missed application ports bound to a loopback.
November 4, 2021	1.7	Upgraded AWS Lambda Node.js runtime to 14.x in CloudFormation templates to fix Amazon ECS deployment issue.

Release date	Version	Details
October 28, 2021	1.6	- Added support for Windows application deployment to ECS Fargate. - Optimized Docker image sizes on Linux by using the Java JRE instead of the JDK. If your application needs to use the JDK, you can edit the Dockerfile that is produced by the containerize command. For more information, see Troubleshoot Java applications on Linux. - Updated default version for new Amazon EKS deployments to Amazon EKS version 1.19. - Added automatic filtering for the inventory command on Linux, to suppress reporting of standard Java processes that are running on the application server. - Added bug fixes, including the following: - Improved tagging for EC2 instances created by SSM documents. - Fixed containerize command issue with lowercase Windows drive names. - Fixed deployment issue for Docker image tags. - Fixed repeated ports in the <code>analysis.json</code> file. - Fixed server backup issue for IIS site with encrypted password. - Fixed Windows Active Directory issue for containerization and deployment.

Release date	Version	Details
July 26, 2021	1.5	• Added support for containerizing complex Windows applications. • Added support for Amazon EKS tagging. • Added support for pipeline tagging. • Added bug fixes, including the following: • Reclassified DockerInvalidImageError for clarity. • Fixed App Runner deployment error that occurs when the local App Runner container is still running at deployment time. • Fixed containerization error when analysis.json has escaped characters.
May 20, 2021	1.4	• Added support for deployments to AWS App Runner. • Added support for reuse of existing Active Directory security groups with gMSA. • Added bug fixes, including the following: • Fixed containerization of Windows applications that use a secondary drive mount. • Clarified some messaging on actionable errors. • Fixed an issue that resulted in incorrect analysis of RHEL Java processes. • Fixed issues related to symbolic links in the application server that resulted in larger than necessary image sizes.

Release date	Version	Details
March 29, 2021	1.3	• Added support for Amazon EC2 instance profiles. • Added support for Amazon EC2 Nitro instance types in App2Container deployments for Windows applications. • Added support for Windows Server Core Version 2004 base images for containerized Windows applications. • The container base image for Windows applications now defaults to match the OS version for the server that runs containerization. • The base image for Windows Amazon ECS deployment artifacts matches the container base image. • Added bug fixes, including the following: • Fixed issue related to symbolic links during containerization. • Fixed issue related to spaces in paths in Dockerfiles. • Fixed issues related to the Windows remote setup script.

Release date	Version	Details
December 21, 2020	1.2	• Added capability to run commands remotely. • Added custom tag support for deployment resources. • Added support for HTTPS endpoints and ACM-based certificate management for Amazon ECS deployments. • Added bug fixes, including the following: • Fixed containerization issue on Linux when unidentified base image uses default image. • Added exclusion for AWS credentials when containerizing Linux applications.

Release date	Version	Details
November 24, 2020	1.1	• Added support for Active Directory authenticated Windows application deployments to Amazon EKS using gMSA. • Added support for named profile overrides to commands that interact with AWS. • Enabled automatic log upload and adjusted console messaging when errors occur (requires IAM policy update). • Added capability to manually upload logs and other artifacts with the upload-support-bundle command (requires IAM policy update). • Added bug fixes, including the following: • Updated CloudFormation templates for Amazon EKS deployments to address previous issues. • Fixed internal/user error classification. • Fixed upgrade errors to point to the correct log file. • Added an explicit check to validate use of Docker version 17.07 or above.

Release date	Version	Details
September 15, 2020	1.0.2	• Added FireLens logging support. • Added container image validation to pipeline generation. • Added bug fixes, including the following: • Removed execution role in template if Windows is specified. • Fixed template to reflect CloudFormation API change. • Fixed autocomplete installation bug. • Fixed dark font for Windows errors. • Improved error messaging for command execution errors. • Fixed containerize error where included file is not valid.

Release date	Version	Details
August 5, 2020	1.0.1	- Improved memory usage while archiving in Windows. - Added support for containerizing individual applications running in Tomcat and JBoss standalone frameworks. - Added schema version and unhealthy version checks. - Added bug fixes, including the following: - Fixed handling for .NET Windows app running on alternative drives (not C). - Fixed COPY command failure in DockerFile. - Access denied error now throws user error. - Added automatic removal of characters that are not allowed in AppId. - Optimized Windows container image size for websites with multiple apps. - Fixed error handling for input arguments validation. - Fixed Dockerfile generation failure when dynamic logging is enabled. - EKS CloudFormation templates are now compatible with the new CloudFormation custom resource API.
June 30, 2020	1.0.0	Initial release.

Document history for AWS App2Container

The following table describes important changes to the documentation by date. For detailed updates to AWS App2Container, see [Release notes for AWS App2Container](#).

You can subscribe to the RSS feed on this page to receive notifications about updates to the documentation.

Change	Description	Date
App2Container version 1.51	This release updated the AWS Lambda function for Amazon Elastic Container Service (Amazon ECS) deployment with gMSA to use Node.js 22.	August 29, 2025
App2Container version 1.50	This release changed the default deployed version of Amazon Elastic Kubernetes Service (Amazon EKS) clusters to 1.32, removed support for Windows Server 2019-based applications in the generate pipeline command for Azure DevOps, and updated allowed Amazon EC2 instance types for Amazon ECS and Amazon EKS deployments. This release also includes miscellaneous bug fixes.	May 28, 2025
App2Container version 1.48	This release revised the supported operating systems for installed App2Container to Windows Server 2016 and later and fixed an issue where the container	December 17, 2024

ize command could fail.
This release also includes
miscellaneous bug fixes.

[App2Container version 1.47](#)

This release removed IIS as a prerequisite for Windows services inventory, added support for .NET 9 applications, and is the last version compatible with Windows Server 2008 and Windows Server 2012. This release also includes miscellaneous bug fixes.

November 22, 2024

[App2Container version 1.46](#)

Updated the Lambda function used for Amazon ECS deployments. This release also includes miscellaneous bug fixes.

October 10, 2024

[App2Container version 1.45](#)

Launch templates are now used when you deploy to Amazon Elastic Compute Cloud (Amazon EC2) container instances for Amazon Elastic Container Service (Amazon ECS). This release also includes miscellaneous bug fixes.

September 11, 2024

[App2Container version 1.44](#)

Enhanced logging for Jenkins input validation during pipeline deployment.

July 26, 2024

[App2Container version 1.43](#)

Improved Jenkins plugin detection during pipeline deployment.

July 11, 2024

App2Container version 1.42	This release improved connection string detection for Windows applications.	July 3, 2024
App2Container version 1.41	This release added a new supported operating system for container hosts.	June 14, 2024
App2Container version 1.40	This release changed the default deployment version of Amazon Elastic Kubernetes Service (Amazon EKS) and updated the analyze command to a different default base image on CentOS applications servers. This release also includes miscellaneous bug fixes.	June 10, 2024
App2Container version 1.39	This release improves the detection of connection strings and configuration files for IIS .NET applications and Windows services.	May 21, 2024
App2Container version 1.38	This release includes miscellaneous bug fixes.	May 2, 2024
App2Container version 1.37	This release includes updates for additional AWS Lambda functions for Amazon Elastic Kubernetes Service (Amazon EKS) deployments. This release also includes miscellaneous bug fixes.	March 8, 2024

App2Container version 1.36	This release includes a bug fix that affects Windows Server 2008.	February 22, 2024
App2Container version 1.35	This release changed the Node.js version used by AWS CloudFormation, updated the AWS Lambda functions to use Python 3.9, improved error messaging, and added detection for the command line used to invoke Windows services. This release also includes miscellaneous bug fixes.	February 14, 2024
App2Container version 1.34	This release added support for .NET 8 applications.	November 16, 2023
AWS Systems Manager automation runbook	The AWSApp2Container-ReplatformApplications Automation runbook is available for use on Amazon EC2 instances. The automation performs the installation of App2Container as well as the initialize, analyze, and transform phases for replatforming supported applications. It can also push the containerized application to Amazon Elastic Container Registry (Amazon ECR). For more information, see App2Container Automation runbook .	November 2, 2023

App2Container version 1.33	This is a maintenance release that contains miscellaneous bug fixes.	October 20, 2023
App2Container version 1.32	This release adopted a new telemetry endpoint to collect metrics, logs, and command-generated artifacts. The example IAM policies were also revised to support this new telemetry endpoint. This release also includes miscellaneous bug fixes.	August 28, 2023
App2Container version 1.31	This release changed to using ECS-optimized AMIs for Windows build instances for AWS CodePipeline integration. This release also includes miscellaneous bug fixes.	August 1, 2023
App2Container version 1.30	This release includes support to deploy Microsoft Azure DevOps for some operating systems, changed the default instance type used when deploying to Amazon ECS on EC2 instances, changed the instance type for Windows container pipelines created in AWS CodePipeline, and changed the default AWS CodeBuild environment type. This release also includes miscellaneous bug fixes.	June 30, 2023

App2Container version 1.29	This release changed the default deployed version of Amazon Elastic Kubernetes Service (Amazon EKS) to 1.26 and includes miscellaneous bug fixes.	May 26, 2023
Documentation addition – compatibility guide	Added a section to the documentation to detail App2Container compatibility for operating systems, software, and tooling. For more information, see App2Container compatibility .	May 22, 2023
App2Container version 1.28	This release causes App2Container to change the Amazon Simple Storage Service (Amazon S3) object key prefix when you use certain commands to upload AWS CloudFormation templates to Amazon S3. The release also changed the instance metadata version of Windows build instances to IMDSv2 when they integrate with AWS CodePipeline and includes miscellaneous bug fixes.	May 4, 2023
App2Container version 1.27	This release was rolled back due to an issue with analyzing applications on Linux.	April 20, 2023

Documentation addition – AWS CodePipeline	Added a section to the documentation for how to integrate AWS CodePipeline with AWS App2Container.	April 18, 2023
App2Container version 1.26	This release changed the instance metadata version to IMDSv2 for cluster instances when you create Amazon Elastic Container Service and Amazon Elastic Kubernetes Service clusters. This release also includes the creation of machine-based application IDs to ensure uniqueness of the application, and miscellaneous bug fixes.	February 23, 2023
App2Container version 1.25	This release causes App2Container to default to .NET Framework version 4.8 for base images, and miscellaneous bug fixes.	January 23, 2023
App2Container version 1.24	This release includes support for the .NET 7 application, and miscellaneous bug fixes.	December 11, 2022
App2Container version 1.23	This release includes support for clusters created by EKS Blueprints, and miscellaneous bug fixes.	November 15, 2022
App2Container version 1.22	This is a maintenance release that contains miscellaneous bug fixes.	October 10, 2022

App2Container version 1.21	This release includes miscellaneous bug fixes and changes the new Amazon EKS deployment version default to Amazon EKS version 1.22.	September 16, 2022
App2Container version 1.20	This is a maintenance release that contains miscellaneous bug fixes.	September 9, 2022
App2Container version 1.19	This is a maintenance release that contains miscellaneous bug fixes.	August 30, 2022
App2Container version 1.18	This release includes miscellaneous bug fixes and enhanced Azure DevOps so that Amazon EKS deployments can work without a kubectl-specific IAM user.	August 2, 2022
App2Container version 1.17	This release includes miscellaneous bug fixes and adds support for Microsoft Azure DevOps as a deployment pipeline.	June 20, 2022
Documentation update – Manage secrets	Updated information about storing secrets using Secrets Manager to reflect changes to the console.	May 31, 2022
App2Container version 1.16	This is a maintenance release that contains miscellaneous bug fixes.	May 9, 2022

Documentation improvement – getting started	Reorganized setup and getting started sections to reduce confusion for people who are just getting started with App2Container. Steps and prerequisites are now clearly labeled in the navigation panel.	April 14, 2022
App2Container version 1.15	This release includes the following changes, along with miscellaneous bug fixes: optimized AWS App2Container installer size, deprecated MD5 checksum validation for the App2Container installer , made AWS profile setup optional during init.	April 14, 2022
App2Container version 1.14	This is a maintenance release that contains miscellaneous bug fixes.	March 31, 2022
App2Container version 1.13	This is a maintenance release that contains miscellaneous bug fixes.	March 2, 2022
App2Container version 1.12	Fixed an issue that caused remote execution commands to fail on Windows.	February 16, 2022
App2Container version 1.11	This is a maintenance release that contains miscellaneous bug fixes.	February 9, 2022

App2Container version 1.10	This release includes the following changes, along with miscellaneous bug fixes: removed dependencies on Amazon ECR and Docker plugins for Jenkins pipelines.	January 14, 2022
Docs-only: IAM policy adjustments	Updated example policy to include FireLens permissions. Updated permissions list order to alphabetical for more intuitive search.	January 3, 2022
Docs-only: Jenkins troubleshooting	Added troubleshooting scenario for Jenkins.	December 20, 2021
App2Container version 1.9	This release includes the following changes, along with miscellaneous bug fixes: workflow support for ASP.NET Core applications running on Linux, including single file applications.	December 8, 2021

[App2Container version 1.8](#)

This release includes the following changes, along with miscellaneous bug fixes: support for Jenkins pipeline deployment, support for incremental deployments of service and infrastructure changes for automatic deployments to Amazon ECR, Amazon EKS, and App Runner, plus automatic pre-validation for the **containerize** command, with the `--no-validate` option to skip that step.

November 26, 2021

[App2Container version 1.7](#)

Upgraded AWS Lambda Node.js runtime to 14.x in CloudFormation templates to fix Amazon ECS deployment issue.

November 4, 2021

[App2Container version 1.6](#)

This release includes the following changes, along with miscellaneous bug fixes: Windows application deployment to AWS Fargate, optimization of Docker image sizes on Linux, improved filtering for standard Java processes reported by the `inventory` command, and updated Amazon EKS default version to 1.19.

October 28, 2021

[App2Container version 1.5](#)

This release includes the following changes, along with miscellaneous bug fixes: support for containerizing complex Windows applications, and tagging support for Amazon EKS and pipelines.

July 26, 2021

[App2Container version 1.4](#)

This release includes the following changes, along with miscellaneous bug fixes: support for deployments to AWS App Runner, and support for reuse of existing Active Directory security groups with gMSA.

May 20, 2021

[App2Container version 1.3](#)

This release includes the following changes, along with miscellaneous bug fixes: support for Amazon EC2 instance profiles, and enhancements for Windows application containers (support for Amazon EC2 Nitro instance types in App2Container deployments, support for Windows Server Core Version 2004 base images, container base image defaults to match the OS version for the server that runs containerization, and Amazon ECS deployment artifacts to match the container base image).

March 29, 2021

Docs-only: applicationMode settings	Describe container configuration applicationMode settings in more detail.	March 19, 2021
Docs-only: IAM policy sections	Add content to describe optional sections of the IAM policy templates.	February 18, 2021
Docs-only: IAM update	Update IAM policy examples for Amazon EKS and Amazon ECS to reflect recent changes and adjust S3 section to remove problematic permission.	January 12, 2021
App2Container version 1.2	This release includes the following changes, along with miscellaneous bug fixes: capability to run commands remotely, custom tag support for deployment resources, support for HTTPS endpoints and ACM-based certificate management for Amazon ECS deployments, and exclusion of AWS credentials when containerizing Linux applications. Note: remote command capability requires an IAM policy update.	December 21, 2020

App2Container version 1.1	This release includes the following changes, along with miscellaneous bug fixes: Added support for Amazon EKS gMSA, introduced named profile overrides for commands that interact with AWS, enabled automatic log uploads for command failures, and added a command to upload a support bundle for help with troubleshooting from App2Container support. Note: uploads for log and support file bundles require an IAM policy update.	November 24, 2020
App2Container version 1.0.2	Added FireLens logging support, plus patches for AWS App2Container version 1.0.2.	September 15, 2020
App2Container version 1.0.1	Added Release notes page with version 1.0.1 changes for AWS App2Container.	August 5, 2020
Docs-only: configuration and IAM updates	A chapter was added to describe configurable fields in files generated by App2Container commands, and the security section was updated with an IAM best practices summary and guidance for setting up IAM general use resources for App2Container.	August 1, 2020

[Initial release](#)

This release introduces AWS App2Container.

June 30, 2020