JON BONSO AND LERVIN JOHN OBANDO

AWS CERTIFIED SYSOPS ADMINISTRATOR ASSOCIATE

VERSION 2.0

Tutorials Dojo
Study Guide and Cheat Sheets

TABLE OF CONTENTS

INTRODUCTION

Today, we live in a world of fast-paced innovation and the invention of new technologies, where competitors race to develop the next disruptive product in the market. Companies with on-premises environments are quickly shifting to the cloud, such as AWS, for the many advantages that it brings. Furthermore, AWS has been the leading cloud service provider for the past few years and is continually maturing. Millions of users and businesses have already adopted the AWS platform for their operations, but not all can capitalize on the benefits that AWS brings to its customers. It takes specialized individuals to operate on the AWS cloud successfully.

AWS is built and managed by highly experienced engineers who offer their expertise to deliver the best products and solutions. That is why you can almost always find a function or service in AWS that would fulfill whatever need or requirement you have. A lot of the heavy-lifting is offloaded from you as the customer so that you can dedicate your efforts and resources to your business operations. Another significant benefit of the AWS cloud is that it is cost-effective. Resources can be quickly provisioned for a very low price and can be decommissioned as quickly once you don't need them anymore. The cloud is an essential piece in DevOps and SysOps since you can quickly spin up test environments and simplify deployment processes that are usually difficult and expensive to do in traditional data center setups.

The AWS Certified SysOps Administrator Associate is a well-recognized certificate in the IT industry and is a major booster for career opportunities and salary increases. Having this certificate in your portfolio means that you have the knowledge and skills in deployment, management, and operations on AWS. Once you have gained more experience with AWS, you can also aim for higher-level certifications, such as the AWS Certified DevOps Engineer Professional certificate. Previously, it was mandatory that you first pass the SysOps Associate exam since the Professional level certificate is very tough and requires much hands-on experience from the exam taker. So if you are planning to pursue a career in Cloud DevOps, passing the AWS Certified SysOps Administrator Associate is a great way to start the journey.

Note: We took extra care to come up with these concise articles and cheat sheets, however, this is meant to be just a supplementary resource when preparing for the exam. We highly recommend working on **hands-on sessions** and **practice exams** to further expand your knowledge and improve your test-taking skills.

AWS CERTIFIED SYSOPS ADMINISTRATOR ASSOCIATE EXAM - STUDY GUIDE AND TIPS

If you are a Systems Administrator or a DevOps Engineer, then this certification will test your knowledge on these areas in AWS. Your experience in these fields will come in handy in passing the exam, but this should be complemented by actual AWS SysOps knowledge. In the AWS Certified SysOps Administrator Associate Exam (or AWS SOA for short), the questions including the exam labs will test your ability to perform the following:

- Deploy, manage, and operate scalable, highly available, and fault-tolerant systems on AWS
- Implement and control the flow of data to and from AWS
- Select the appropriate AWS service based on compute, data, or security requirements
- Identify appropriate use of AWS operational best practices
- Estimate AWS usage costs and identify operational cost control mechanisms
- Migrate on-premises workloads to AWS

Given the scope of the questions and exam labs, you should learn the concepts of the AWS architecture, the AWS Operational Framework, as well as the AWS CLI and AWS SDK/API tools. Having prior knowledge of fundamental networking and security will also be very valuable. This guide aims to provide you a straightforward guide when reviewing for this exam.

Study Materials

The FREE AWS Exam Readiness video course, official AWS sample questions, whitepapers, AWS Documentation, AWS cheat sheets, and AWS practice exams will be your primary study materials for this exam. There are multiple papers that you should read and familiarize yourself with as a SysOps Administrator.

Having an AWS account you can use will help ingest the different concepts within these whitepapers. Since the exam itself contains multiple scenario questions and exam labs, using the services and applying them in practice yourself will allow you to determine the types of situations they are applied in.

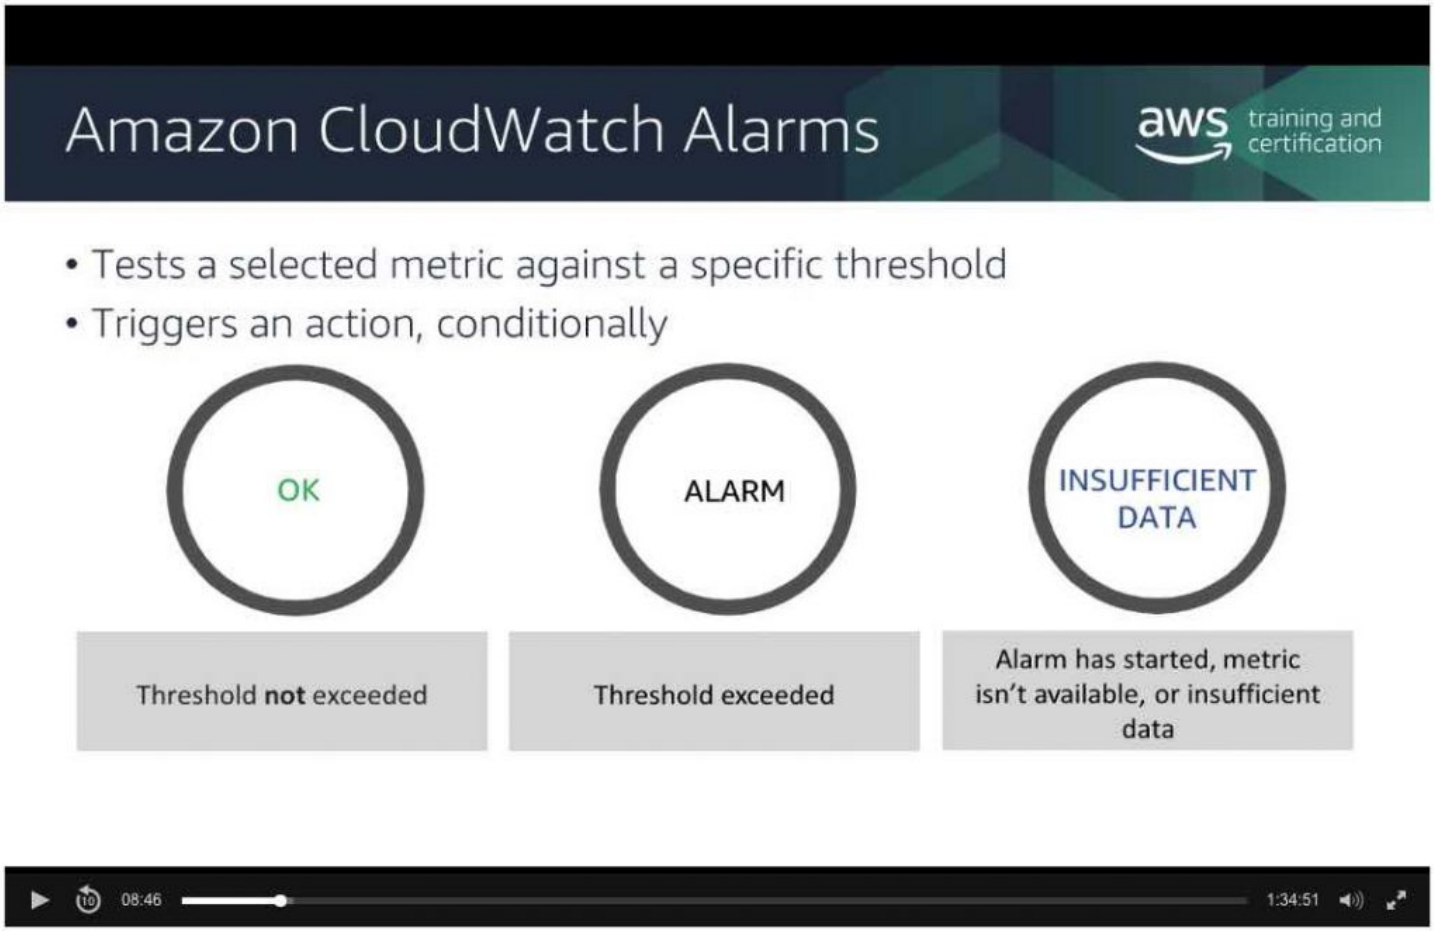

Additional details regarding your AWS SOA exam can be seen in this AWS exam blueprint.

The whitepapers listed below are arranged in such a way that you will learn the concepts first, before proceeding to application and best practices. If you need a refresh on your AWS fundamentals, go check out our guide on the AWS Certified Cloud Practitioner Exam before proceeding below.

1. Amazon Virtual Private Cloud Connectivity Options - Study how you can connect different VPCs together, your VPCs to your on-premises network, and vice versa.
2. Development and Test on AWS - Study how you can leverage AWS to create development and test environments, implement pipelines and automation, and perform different validation tests for your applications.
3. Backup and Recovery Approaches Using AWS - Learn which AWS services offer backup and restore features. It is also important to know how these backups are stored and secured, and selecting the correct storage options for them.
4. How AWS Pricing Works - Study the fundamental drivers of cost in AWS, the pricing models of commonly used services in compute, storage, and database, and how to optimize your costs.

5. Amazon Web Services: Overview of Security Processes - You should study the different security features in AWS - including infrastructure, account, network, application, and data security. Determine which aspects of security are your responsibilities, and which are AWS'.
6. AWS Security Best Practices - This whitepaper complements the previous. Understand the security best practices and their purpose in your environment. Some services offer more than one form of security feature, such as multiple key management schemes for encryption. It is important that you can determine which form is most suitable to the given scenarios in your exam.
7. Architecting for the Cloud: AWS Best Practices - Be sure to understand the best practices in AWS since exam questions will focus their scenarios around these best practices. The whitepaper contains a number of design principles with examples for each.
8. AWS Well-Architected Framework - This whitepaper is one of the most important papers that you should study for the SOA-C02 exam. It discusses the different pillars that make up a well-architected cloud environment.

Optional whitepapers:

1. Overview of Deployment Options on AWS - This is an optional whitepaper that you can read to be aware of your deployment options in AWS. There is a chance that this might come up in the exam.
2. AWS Disaster Recovery Plans - As a SysOps Administrator, you should be familiar with your DR options when outages occur. Having knowledge of DR will determine how fast you can recover your infrastructure.

AWS Services to Focus On

AWS offers extensive documentation and well-written FAQs for all of their services. These two will be your primary source of information when studying. Furthermore, as an AWS SysOps Administrator, you need to be well-versed in a number of AWS products and services since you will almost always be using them in your work. I recommend checking out Tutorials Dojo's AWS Cheat Sheets which provides a summarized but highly informative set of notes and tips for your review on these services.

Core services to study:

1. EC2 - As the most fundamental compute service offered by AWS, you should know about EC2 inside out.
2. Elastic Load Balancer - Load balancing is very important for a highly available system. Study the different types of ELBs, and the features each of them supports.
3. Auto Scaling - Study what services in AWS can be auto-scaled, what triggers scaling, and how auto scaling increases/decreases the number of instances.
4. Elastic Block Store - As the primary storage solution of EC2, study the types of EBS volumes available. Also study how to secure, backup, and restore EBS volumes.

5. S3 / Glacier - Study the S3 storage types are and what differs between them. Also review the capabilities of S3 such as hosting a static website, securing access to objects using policies, lifecycle policies, etc. Learn as much about S3 as you can.
6. VPC - Study every service that is used to create a VPC (subnets, route tables, internet gateways, nat gateways, VPN gateways, etc). Also, review the differences between network access control lists and security groups, and during which situations they are applied.
7. Route 53 - Study the different types of records in Route 53. Study also the different routing policies. Know what hosted zones and domains are.
8. RDS - Know how each RDS database differs from one another, and how they are different from Aurora. Determine what makes Aurora unique, and when it should be preferred from other databases (in terms of function, speed, cost, etc). Learn about parameter groups, option groups, and subnet groups.
9. DynamoDB - Consider how DynamoDB compares to RDS, Elasticache, and Redshift. This service is also commonly used for serverless applications along with Lambda.
10. Elasticache - Familiarize yourself with Elasticache redis and its functions. Determine the areas/services where you can place a caching mechanism to improve data throughput, such as managing the session state of an ELB, optimizing RDS instances, etc.
11. SQS - Gather info on why SQS is helpful in decoupling systems. Study how messages in the queues are being managed (standard queues, FIFO queues, dead letter queues). Know the differences between SQS, SNS, SES, and Amazon MQ.
12. SNS - Study the function of SNS and what services can be integrated with it. Also, be familiar with the supported recipients of SNS notifications.
13. IAM - Services such as IAM Users, Groups, Policies, and Roles are the most important to learn. Study how IAM integrates with other services and how it secures your application through different policies. Also, read on the best practices when using IAM.
14. CloudWatch - Study how monitoring is done in AWS and what types of metrics are sent to CloudWatch. Also read upon CloudWatch Logs, CloudWatch Alarms, and the custom metrics made available with CloudWatch Agent.
15. CloudTrail - Familiarize yourself with how CloudTrail works, and what kinds of logs it stores as compared to CloudWatch Logs.
16. Config - Be familiar with the situations where AWS Config is useful.
17. CloudFormation - Study how CloudFormation is used to automate infrastructure deployment. Learn the basic makeup of a CloudFormation template, stack, and stack set.
18. KMS – Familiarize how KMS integrates with other services in storing encryption keys.
19. Secrets Manager – Understand how Secrets Manager stores secrets and how you can use them with other AWS services.
20. Parameter Store – Know when to use Parameter store and how compute services like EC2, ECS, and Lambda utilize it.
21. DataSync – Familiarize which AWS services can be used to migrate data from an on-premises data center.

Some additional services we recommend to review:

1. Trusted Advisor
2. Systems Manager
3. CodeDeploy
4. CodePipeline
5. CloudFront
6. Cost and Billing Management Console
7. OpsWorks
8. Direct Connect

For the exam version (SOA-C02), you should also know the following services:

1. Amazon FSx for Windows File Server and Amazon FSx for Lustre
2. AWS Backup
3. EC2 Image Builder
4. S3 Transfer Acceleration
5. AWS Global Accelerator
6. RDS Proxy
7. IAM Access Analyzer

Exam Labs

The SOA-CO2 includes an exam labs section where you have to perform SysOps related tasks on the AWS Management Console. To prepare for this, make sure to play around with the different AWS services covered in the exam. You don't need to memorize all the configurations for each service. But you have to be really good at navigating the AWS Management console to understand where you can configure the requirements in each exam lab. Focus on preparing for exam labs on setting up a VPC, CloudWatch, Load Balancer, Auto Scaling, CloudFormation, and S3.

View our sample exam lab **here**.

Common Exam Scenarios

Scenario	Solution
Monitoring, Reporting, and Remediation	
You need to set up an alert that notifies the IT manager about EC2 instances service limits.	Use Amazon CloudWatch Events to detect and react to changes in the status of Trusted Advisor checks

You need to track the deletion and rotation of CMKs.	Use AWS CloudTrail to log AWS KMS API calls
You need to investigate if the traffic is reaching the EC2 instance.	Use VPC flow logs
You need to ensure that the SSH protocol is always disabled on private servers.	Use AWS Config Rules
You need to retrieve the instance metadata of an EC2 instance.	`http://169.254.169.254/latest/`
You have to monitor the CPU usage of a single process in your EC2 instance.	Use the CloudWatch Agent procstat plugin to monitor system utilization.
You need to generate a report on the replication and encryption status of all of the objects stored in the S3 bucket.	Use S3 Inventory
Metric to use to alarm when all instances behind an ALB becomes unhealthy	`AWS/ApplicationELB HealthyHostCount <= 0`
Monitor restricted CIDR changes on a security group and remove them automatically.	Use AWS Config to evaluate the security group and AWS Systems Manager Automation document to remove the unwanted CIDR range.
Monitor CreateUser API call via email	Utilize Amazon EventBridge, declare CloudTrail as a source, and CreateUser as an event pattern. Create an SNS topic and set it as an event target on Amazon EventBridge.
Reliability and Business Continuity	
When the incoming message traffic increases, the EC2 instances fall behind and it takes too long to process the messages.	Create an Auto Scaling group that can scale out based on the number of messages in the queue.
You need to log the client's IP address, latencies, request paths, and server responses that go through your Application Load Balancer.	Enable access logging in ALB and store the logs on an S3 bucket.
You need to determine which cipher is used for the SSL connection in your ELB.	Enable Server Order Preference
You need to monitor the total number of requests or connections in your load balancer.	Monitor the `SurgeQueueLength` metric

You need to ensure that the backups of an Amazon Redshift cluster are always available.	Configure the Amazon Redshift cluster to automatically copy snapshots of a cluster to another region.
Highly available File Server that supports SMB and manages file permissions using Windows Access Control List (A.	Multi-AZ Amazon FSx for Windows File Server
Slow load time when uploading objects to S3	S3 Transfer Acceleration
PercentIOLimit metric hits 100% on EFS	Create a new Max I/O performance mode EFS file system and migrate data to the new file system using AWS DataSync.
Must ensure data integrity when performing EBS backups	Build a Lambda function that uses CreateImage API to generate AMI of the EC2 instance and include a reboot parameter. Create an Amazon EventBridge rule to execute the Lambda function daily.
Deployment, Provisioning, and Automation	
You must remotely execute shell scripts and securely manage the configuration of EC2 instances.	Use Systems Manager Run Command
You need to identify the configuration changes in the CloudFormation resources.	Use drift detection
Requires a CloudFormation template that can be reused for multiple environments. If the template has been updated, all the stack that is referencing it will automatically use the updated configuration.	Use Nested Stacks
You need to automate the process of updating the CloudFormation templates to map to the latest AMI IDs.	Use CloudFormation with Systems Manager Parameter Store
The eviction count in Amazon ElastiCache for Memcached has exceeded its threshold.	Scale the cluster by increasing the number of nodes.
You need to provide each department a new AWS account with governance guardrails and a defined baseline in place.	Set up AWS Control Tower

An S3 bucket must be configured to move the objects older than 60 days to the Infrequent Access storage class.	Set up a lifecycle policy
You need to monitor all the COPY and UNLOAD traffic in the Redshift cluster.	Enable Enhanced VPC routing on the Redshift cluster.
A total of 500 TB of data needs to be transferred to Amazon S3 in the fastest way.	Use multiple AWS Snowball devices
TLS certificate should be renewed automatically	Request a public certificate via AWS Certificate Manager (ACM)
Get cost expenses of each AWS user account	Enable the createdBy tag in the Billing and Management console
Provisioning instances on ASG takes time because of software dependencies installed via the UserData script.	EC2 Image Builder
Get cost expenses of each AWS user account	Enable the createdBy tag in the Billing and Management console
Security and Compliance	
You have to rotate an existing CMK with imported key material every 6 months	Create a new CMK with imported key material and update the key ID to point to the new CMK
A company needs to restrict access to the data in an S3 bucket.	Use S3 ACL and bucket policy
Mitigate malicious attacks such as SQL injection and DDoS attacks from unknown origins.	Use AWS WAF and Shield
You need to define an IAM policy to enable the user to pass a role to an AWS service.	Define iam:PassRole in the IAM policy
You need to create a solution that allows multiple EC2 instances in a private subnet to use AWS KMS and the traffic must not pass through the public Internet.	Configure a VPC endpoint
You need to encrypt all the objects at rest in your S3 bucket	Use SS3-S3, SSE-KMS or SSE-C

Enable authentication to AWS services using Active Directory Federation Services.	Amazon Cognito user pool
Create a bucket policy to only allow AWS accounts in the organization to access an S3 bucket.	Set principal to (*) and create a condition for PrincipalOrgId
Read, update, delete messages from SQS queues from an instance.	Create a policy with sqs:SendMessage, sqs:ReceiveMessage, sqs:DeleteMessage, and attach the policy to a new role that can perform API calls to AWS. Associate the new role to the instance.
RDS credentials should not be hardcoded on Lambda functions	Use Secrets Manager to store credentials
Networking and Content Delivery	
You need to allow the EC2 instances in your VPC that support IPv6 to connect to the Internet but block any incoming connection.	Set up an egress-only Internet gateway
You have to establish a dedicated connection between their on-premises network and their Amazon VPC.	Set up a Direct Connect connection
You need to increase the cache hit ratio for a CloudFront web distribution.	Add a `Cache-Control max-age` and increase the TTL by specifying the longest value for max-age
You need to ensure that users are consistently directed to the AWS region nearest to them.	Set up a Route 53 Geoproximity routing policy
A company plans to implement a hybrid cloud architecture. You need to allow your resources on AWS the connectivity to external networks.	Assign an Internet Gateway to the VPC Create a Virtual Private Gateway
Users being served desktop version on mobile phones	Add a `User-Agent` header to the list of origin custom header on CloudFront.
DNS record at the apex domain	ALIAS record
Cost and Performance Optimization	
You have to automate the process of patching managed instances with security-related updates.	Use AWS Systems Manager Patch Manager

You need to analyze the data hosted in Amazon S3 using standard SQL.	Use Amazon Athena
Improving the site speed of a static S3 web hosting with customers around the globe	Create a CloudFront web distribution and set Amazon S3 as the origin.
You need to implement a solution to enforce the tagging of all instances that will be launched in the VPC.	Use AWS Service Catalog TagOption library
You need to get billing alerts once it reaches a certain limit.	Enable billing alerts in Account Preferences of the AWS Console.
Resize an Amazon Elasticache for Redis cluster.	Use online resizing for Amazon Elasticache Redis cluster
No sharing of Reserved Instance (RI) discounts between AWS accounts in the Organization	Disable RI discount sharing via management account and provision instances using individual AWS accounts.

Validate Your Knowledge

Once you have finished your review and you are more than confident of your knowledge, test yourself with some practice exams available online. AWS offers a practice exam that you can try out at their aws.training portal. **Tutorials Dojo** also offers a top-notch set of **AWS Certified SysOps Administrator Associate practice tests**. Each test contains unique questions that will surely help verify if you have missed out on anything important that might appear on your exam. You can pair our practice exams with this study guide eBook to further help in your exam preparations.

Sample Practice Test Questions:

Question 1

A company is heavily using AWS CloudFormation templates to automate the deployment of their cloud resources. The SysOps Administrator needs to write a template that will automatically copy objects from an existing S3 bucket into the new one.

Which of the following is the most suitable configuration for this scenario?

1. Set up an AWS Lambda function and configure it to perform the copy operation. Integrate the Lambda function to the CloudFormation template as a custom resource.
2. Configure the CloudFormation template to modify the existing S3 bucket to allow cross-origin requests.
3. Set up the CloudFormation template to use the AWS Data Pipeline CopyActivity object to copy the files from the existing S3 bucket to the new S3 bucket.
4. Configure the CloudFormation template to enable cross-region replication on the existing S3 bucket and select the new S3 bucket as the destination.

Correct Answer: 1

AWS CloudFormation gives you an easy way to model a collection of related AWS and third-party resources, provision them quickly and consistently, and manage them throughout their lifecycles, by treating infrastructure as code. A CloudFormation template describes your desired resources and their dependencies so you can launch and configure them together as a stack. You can use a template to create, update, and delete an entire stack as a single unit, as often as you need to, instead of managing resources individually. You can manage and provision stacks across multiple AWS accounts and AWS Regions.

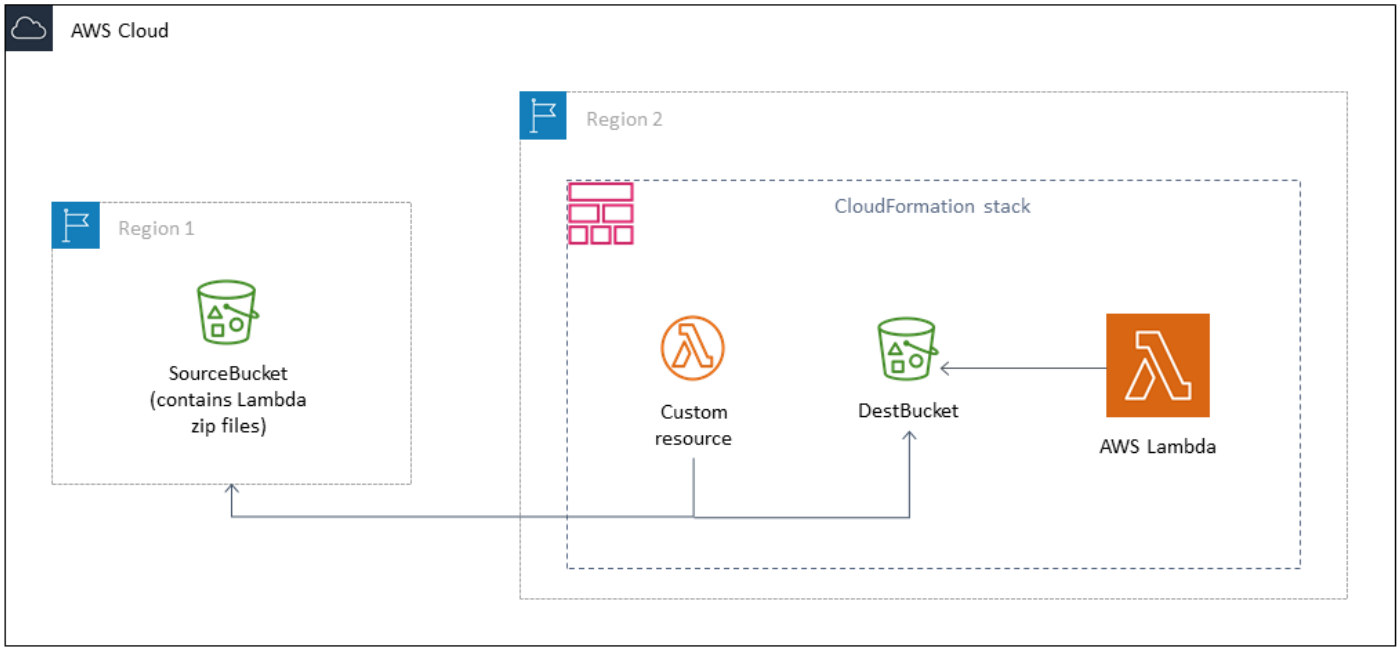

In an AWS CloudFormation template, you can specify a Lambda function as the target of a custom resource. Use custom resources to process parameters, retrieve configuration values, or call other AWS services during stack lifecycle events. When you associate a Lambda function with a custom resource, the function is invoked whenever the custom resource is created, updated, or deleted. AWS CloudFormation calls a Lambda API to invoke the function and to pass all the request data (such as the request type and resource properties) to the function. The power and customizability of Lambda functions in combination with AWS CloudFormation enable a wide range of scenarios, such as dynamically looking up AMI IDs during stack creation, or implementing and using utility functions, such as string reversal functions.

The requirement for this scenario is to copy all the objects from an existing S3 bucket to a new S3 bucket created by the CloudFormation template. To accomplish this requirement, you need to create a custom Lambda function that can copy the objects from the source bucket to the new S3 bucket. You can also define the options you want Amazon S3 to apply during replication, such as server-side encryption, replica ownership, and transitioning replicas to another storage class.

Hence, the correct answer is: **Set up an AWS Lambda function and configure it to perform the copy operation. Integrate the Lambda function to the Cloudformation template as a custom resource.**

The option that says: **Configure the Cloudformation template to enable cross-region replication on the existing S3 bucket and select the new S3 bucket as the destination** is incorrect because this option won't be able to copy the existing objects to the new S3 bucket. For this configuration, you need to invoke Lambda first to copy the objects in the S3 bucket.

The option that says: **Set up the CloudFormation template to CopyActivity object to copy the files from the existing S3 bucket to the new S3 bucket** is incorrect because CopyActivity does not support copying multipart Amazon S3 files. The most suitable configuration to copy the objects from an existing bucket to a new S3 bucket is to use a custom Lambda resource in CloudFormation.

The option that says: **Configure the CloudFormation template to modify the S3 bucket to allow cross-origin requests** is incorrect because the scenario did not state anything about allowing cross-origin access to your Amazon S3 resources. Also, this option does not have the capability to copy all the objects from an existing S3 bucket to a new S3 bucket.

References:
https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/template-custom-resources-lambda.html
https://aws.amazon.com/blogs/infrastructure-and-automation/deploying-aws-lambda-functions-using-aws-cloudformation-the-portable-way/
https://aws.amazon.com/blogs/devops/use-aws-cloudformation-to-automate-the-creation-of-an-s3-bucket-with-cross-region-replication-enabled/

Check out this AWS CloudFormation Cheat Sheet:
https://tutorialsdojo.com/aws-cloudformation/

Question 2

A company plans to expand its use of AWS services across its product portfolios. To ensure separation of business processes for billing, security, and compliance, the SysOps Administrator must provide each department with new AWS accounts having governance guardrails and a defined baseline in place. An efficient and scalable provisioning process is required to optimize the workflow and save time.

Which of the following options can satisfy the given requirement?

1. Use AWS Batch and AWS Organizations to automatically provision new resources and accounts.
2. Use AWS Control Tower to generate templates in the Account Factory and to provision new accounts in AWS Service Catalog.

3. Use AWS OpsWorks for Chef Automate and bootstrapping scripts to handle the configuration management and provisioning tasks.
4. Use AWS Service Catalog and AWS Config to automate account creation and configuration.

Correct Answer: 2

AWS Control Tower provides a single location to easily set up your new well-architected multi-account environment and govern your AWS workloads with rules for security, operations, and internal compliance. You can automate the setup of your AWS environment with best-practices blueprints for multi-account structure, identity, access management, and account provisioning workflow. For ongoing governance, you can select and apply pre-packaged policies enterprise-wide or to specific groups of accounts.

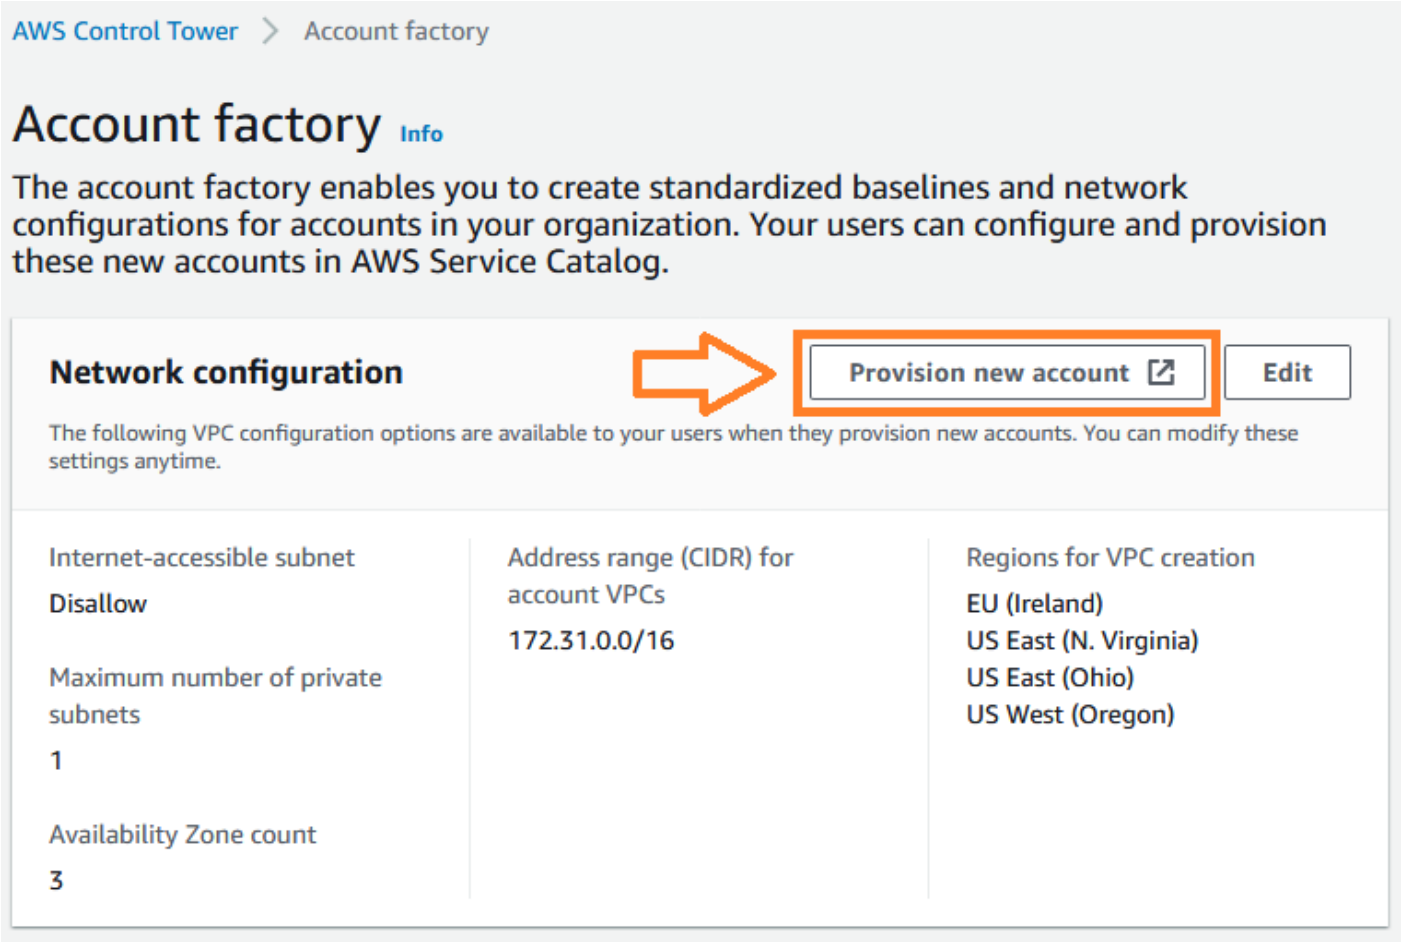

AWS Control Tower provides three methods for creating member accounts:

- Through the Account Factory console that is part of AWS Service Catalog.

- Through the Enroll account feature within AWS Control Tower.

- From your AWS Control Tower landing zone's management account, using Lambda code and appropriate IAM roles.

AWS Control Tower offers "guardrails" for ongoing governance of your AWS environment. Guardrails provide governance controls by preventing deployment of resources that don't conform to selected policies or detecting non-conformance of provisioned resources. AWS Control Tower automatically implements guardrails using multiple building blocks such as AWS CloudFormation to establish a baseline, AWS Organizations service control policies (SCPs) to prevent configuration changes, and AWS Config rules to continuously detect non-conformance.

In this scenario, the requirement is to provide each department with AWS accounts that have governance guardrails and defined baseline in place. To save time and resources, you can use AWS Control Tower to automate the account creation. With the appropriate user group permissions, you can specify standardized baselines and network configurations for all accounts in the organization.

Hence, the correct answer is: **Use AWS Control Tower to generate templates in the Account Factory and to provision new accounts in AWS Service Catalog.**

The option that says: **Use AWS Service Catalog and AWS Config to automate account creation and configuration** is incorrect. Although you can use AWS Service Catalog to create and manage catalogs of your IT services, it still does not offer "guardrails" for ongoing governance of your AWS environment. Moreover, AWS Config is mainly used to evaluate the configuration of various AWS services in an AWS environment and not to provision new accounts.

The option that says: **Use AWS Batch and AWS Organizations to automatically provision new resources and accounts** is incorrect because AWS Batch can only provision compute resources. Furthermore, AWS Organizations do not fully provide governance guardrails unlike AWS Control Tower.

The option that says: **Use AWS OpsWorks for Chef Automate and bootstrapping scripts to handle the configuration management and provisioning tasks** is incorrect because AWS OpsWorks is not a suitable service to be used in provisioning new AWS accounts. The common use case for AWS OpsWorks for Chef Automate is to automate operational tasks on Amazon EC2 instances and on-premises servers.

References:
https://docs.aws.amazon.com/controltower/latest/userguide/account-factory.html
https://aws.amazon.com/blogs/mt/how-to-automate-the-creation-of-multiple-accounts-in-aws-control-tower/
https://aws.amazon.com/blogs/aws/aws-control-tower-set-up-govern-a-multi-account-aws-environment/

Check out these AWS Cheat Sheets:
https://tutorialsdojo.com/aws-cheat-sheets/

Click **here** for more **AWS Certified SysOps Administrator Associate practice exam questions**.

It is best to get some rest before the day of your exam, and review any notes that you have written down. If you have done well in the **practice tests**, go over the questions where you made a mistake and understand why so. If you are not feeling so confident after trying the practice tests, you can just reschedule your exam and take your time preparing. The AWS SOA certification is one of the most sought after certifications in the SysOps Administration field. The exam will not be easy to pass, but it'll be worth it when you do.

AWS Deep Dives

EC2 Placement Groups

Placement Groups is a logical grouping of your interdependent instances in AWS. This logical grouping affects how your instances are placed on the underlying hardware. Having the instances in a placement group has particular benefits in terms of network latency, throughput, and minimizing correlated hardware failure. By default, AWS automatically spreads out your instances across underlying hardware to reduce this correlated hardware failure.

AWS offers different placement strategies which can suit the placement requirements of your application hosted in Amazon EC2.

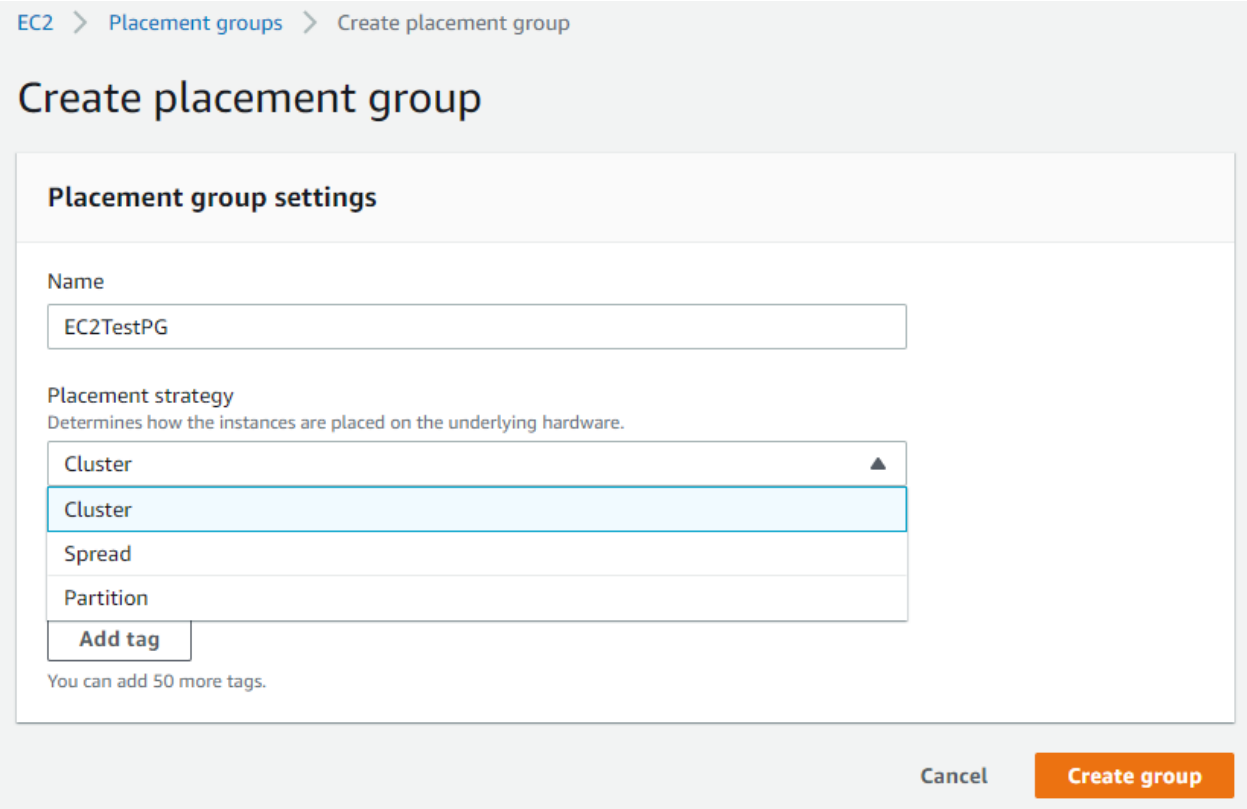

Cluster Placement Group

A cluster placement group is a logical group of instances within a single Availability Zone and instances from peered VPC in the same region. Through VPC Peering, you can still add instances from different Availability Zones to your cluster placement group.

Instances on the same cluster group have a low network latency and high throughput. A cluster placement group is beneficial to applications with a high volume of network traffic between their instances. To further maximize these network performance benefits, you can choose instance types with enhanced networking for your cluster placement group.

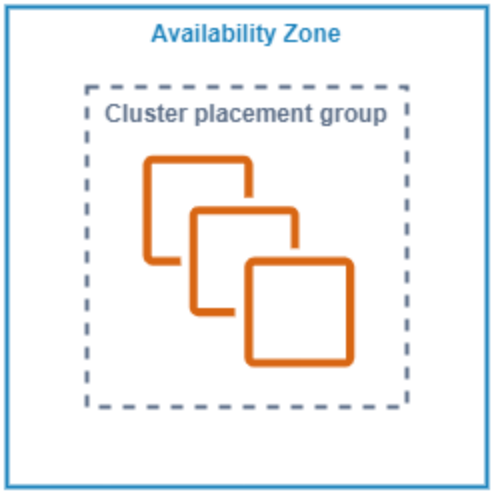

AWS recommends launching the instances for the cluster placement group through a single launch request. They also recommend using the same instance type for all the instances in the placement group to minimize the chance of getting an insufficient capacity error. This error comes out when there is not enough hardware capacity to launch an instance. For example, when adding more instances to an existing placement group or adding instances with a different instance type. The capacity error can also be encountered when you stop and then start an instance again in a placement group.

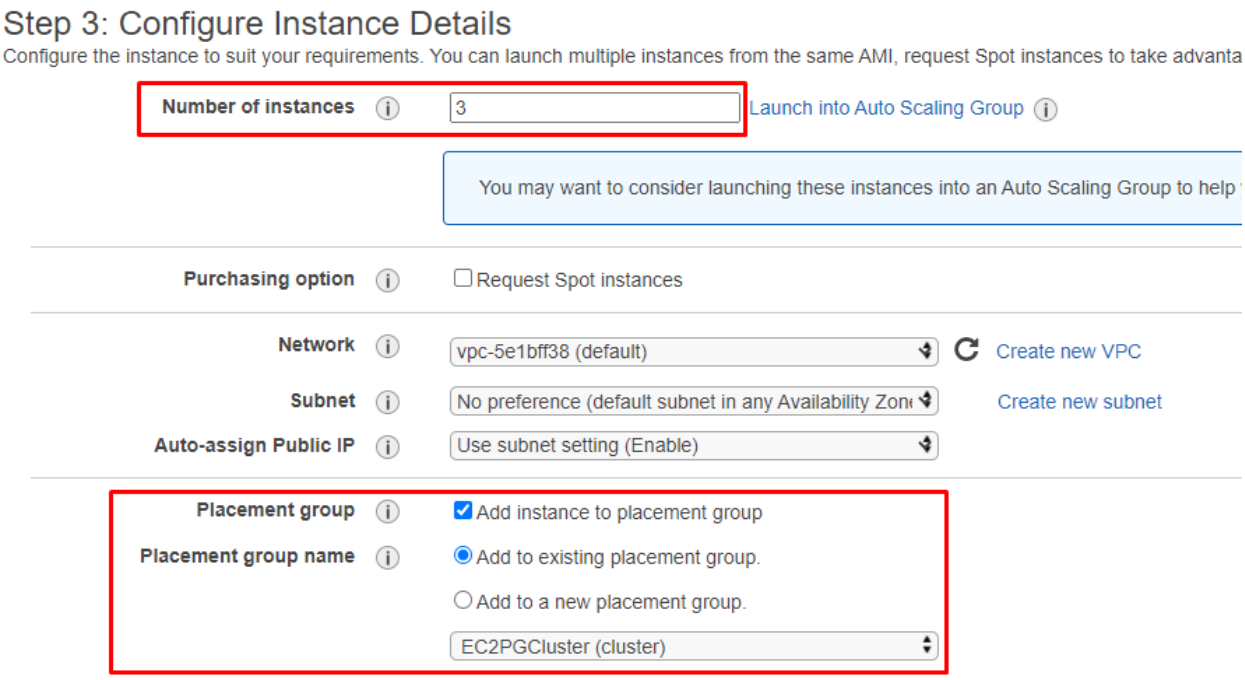

However, if you need to launch an instance to an existing placement group with running instances and encounter an insufficient capacity error, try to stop and start all of the running instances inside the placement group, then relaunch the instance. Doing so may force the instances to boot into new hardware capable of accommodating all the instance requests for the placement group.

Cluster placement groups are commonly used for High-Performance Computing (HPC) applications, like genomics, computational chemistry, financial risk modeling, machine learning, deep learning, etc.

Partition Placement Group

A partition placement group spreads all instances into logical segments called partitions. Each partition has a dedicated rack with its network and power source. This placement strategy ensures that all partitions are isolated from each other, reducing the risk of correlated hardware failures.

Also, partition placement groups can have partitions from different Availability Zones in the same Region with a limit of seven partitions per AZ. The account limit determines the maximum number of instances. However, a maximum of two partitions is allowed for the partition placement group with Dedicated Instance.

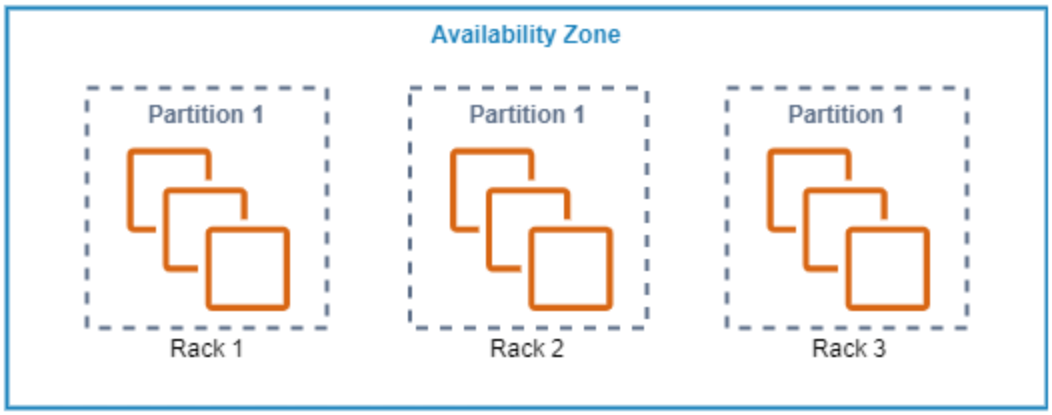

When launching instances to the partition placement group, you can either let Amazon EC2 handle the distribution or specify the specific partition.

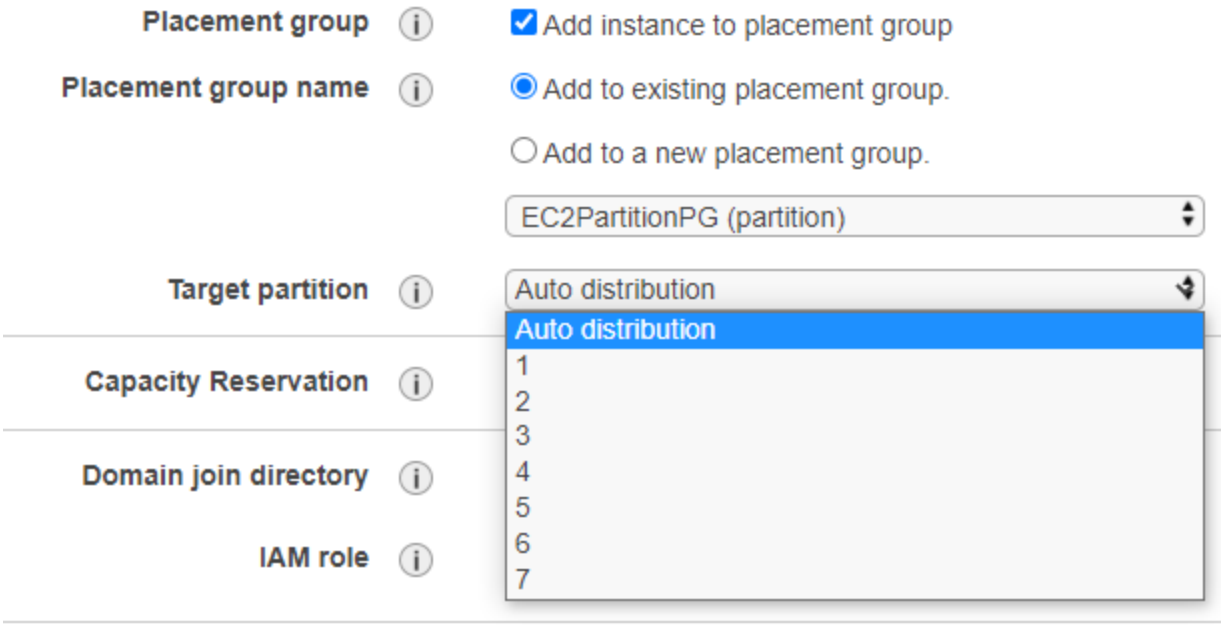

To achieve high availability for the application, we often go to multi-AZ deployment, but some applications are dependent on internode latency, thus making it unavailable for multi-AZ deployment. With a partition placement group, you can deploy this kind of application in a single Availability Zone but with improved performance and less chance for correlated hardware faults.

Applications like HDFS, HBase, and Cassandra are benefiting from this kind of placement strategy. Because they are topology-aware applications, they can use the topology information to make intelligent data storage decisions.

Spread Placement Group

A Spread placement group is a placement strategy that strictly hosts instances separately on a distinct rack that has an individual network and power source. Since all instances are hosted on distinct racks, you can freely have multiple instance types or add instances over time on your spread placement group.

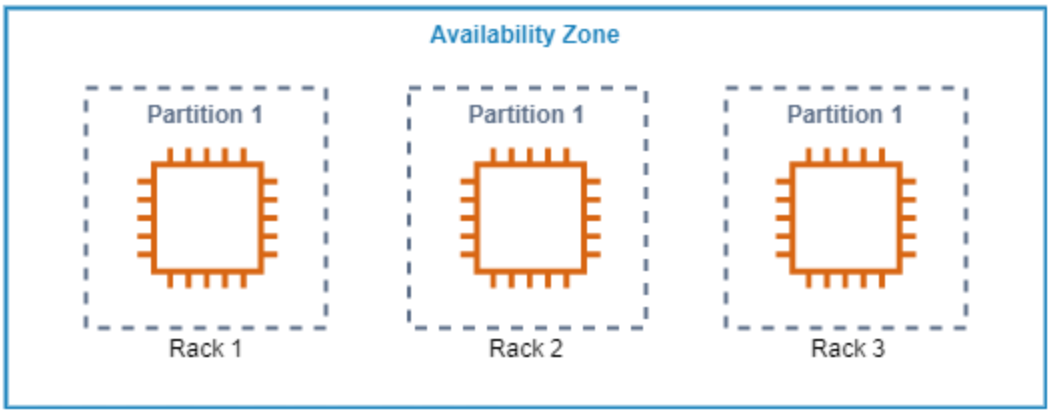

Since instances on the spread placement group are isolated from each other, the chance of having hardware faults is reduced when compared to instances sharing the same rack.

Like partition placement groups, spread placement groups can also span on different Availability Zones with a maximum of seven running instances per AZ.

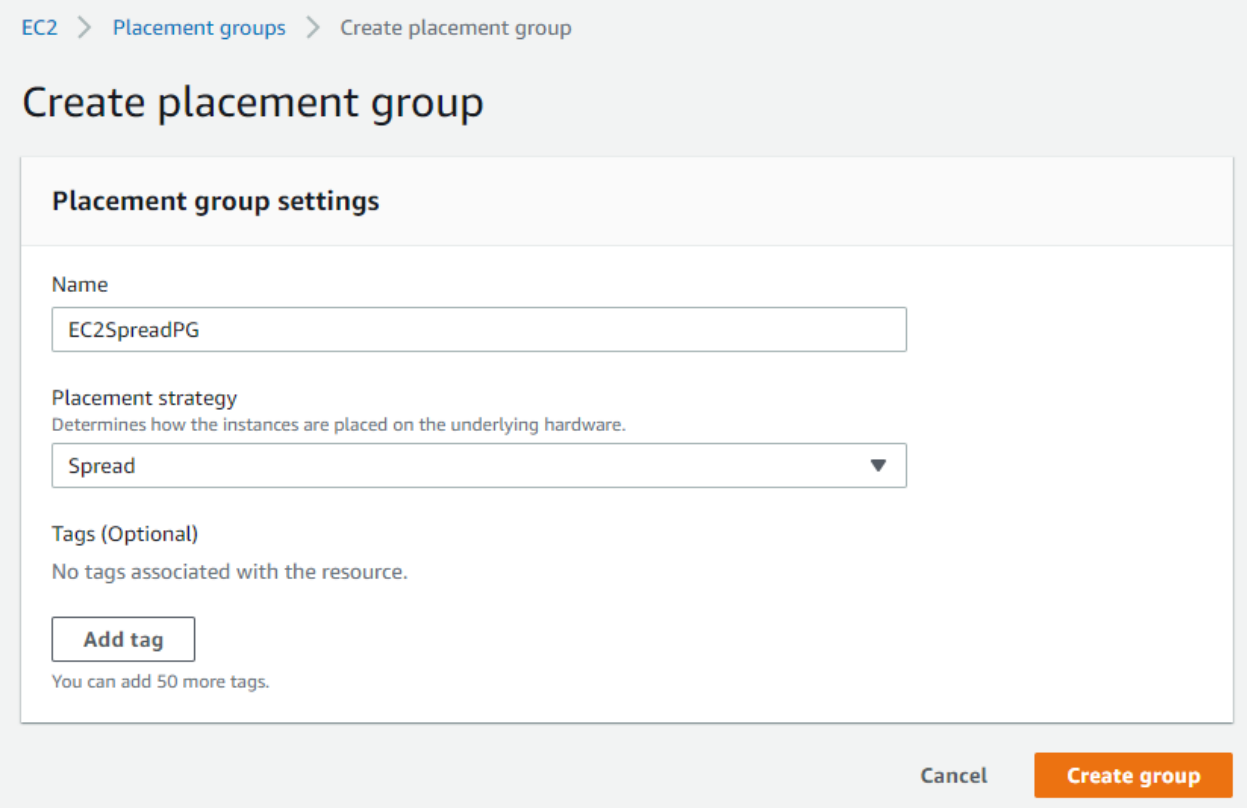

For the Partition and Spread placement group, there are times when a unique hardware is unavailable to accommodate all instance requests. When this happens, try to request again later as more hardware becomes available over time.

Reference:
https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/placement-groups.html

EC2 Image Builder

EC2 Image Builder is an AWS service that automates the process of creating, managing, and deploying machine and docker images both for your AWS environment and on-premises. You can keep your images updated through the image builder, automate image customization, validate image integrity and functionality through testing, and deploy images on different AWS regions. The image builder is pretty straightforward; it lets you create an Image Pipeline, configure an Image recipe, define the infrastructure, and set the image distribution.

Image Pipelines

To automate the creation of images, AWS allows you to create a pipeline where you can configure the necessary components of your custom images. The image creation will run based on the defined build schedule and frequency or can be manually run.

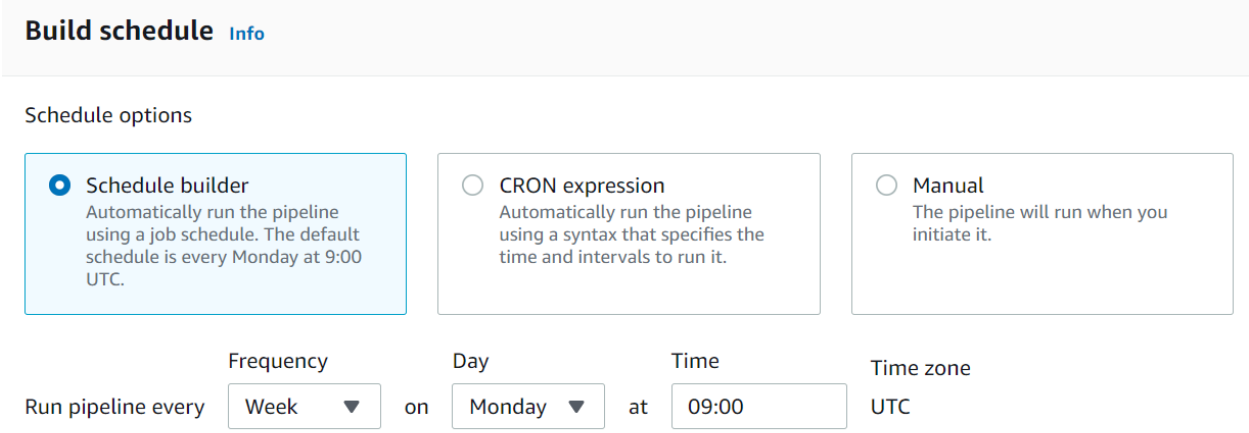

Image Recipes Configuration

Image recipes are where you define the customization and testing of your images. Image recipes are reusable and have version control. It consists of the following components.

Source Images

The source image will be the baseline of your custom image. Image builder supports the customization for Amazon Machine Image (AMI) and Docker image. For AMI, this can be AWS-managed images or a custom AMI. Likewise, for Docker images, it can be AWS-managed images, an ECR image, or a public image from Docker Hub.

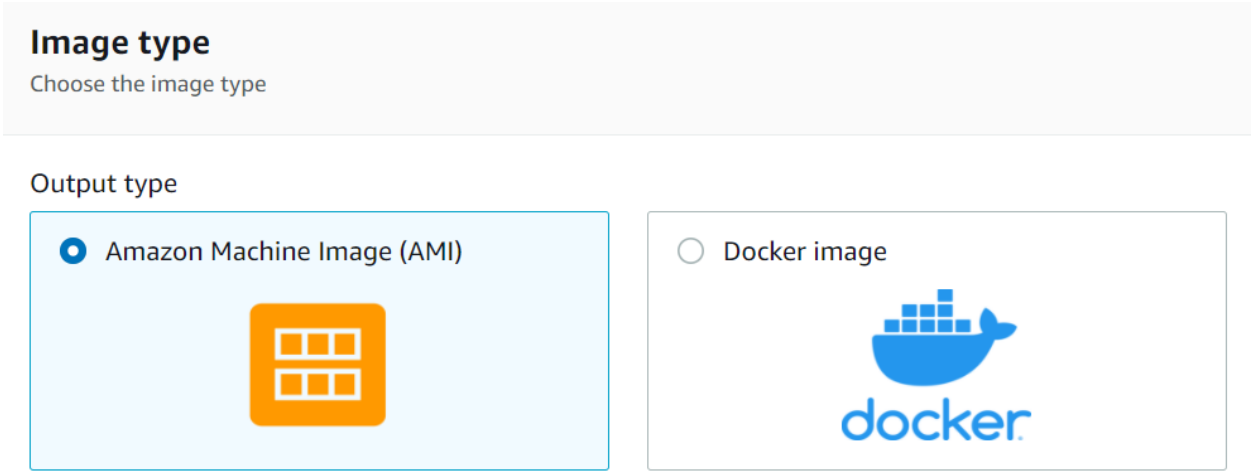

You can select from different operating systems and versions; availability depends on the image type.

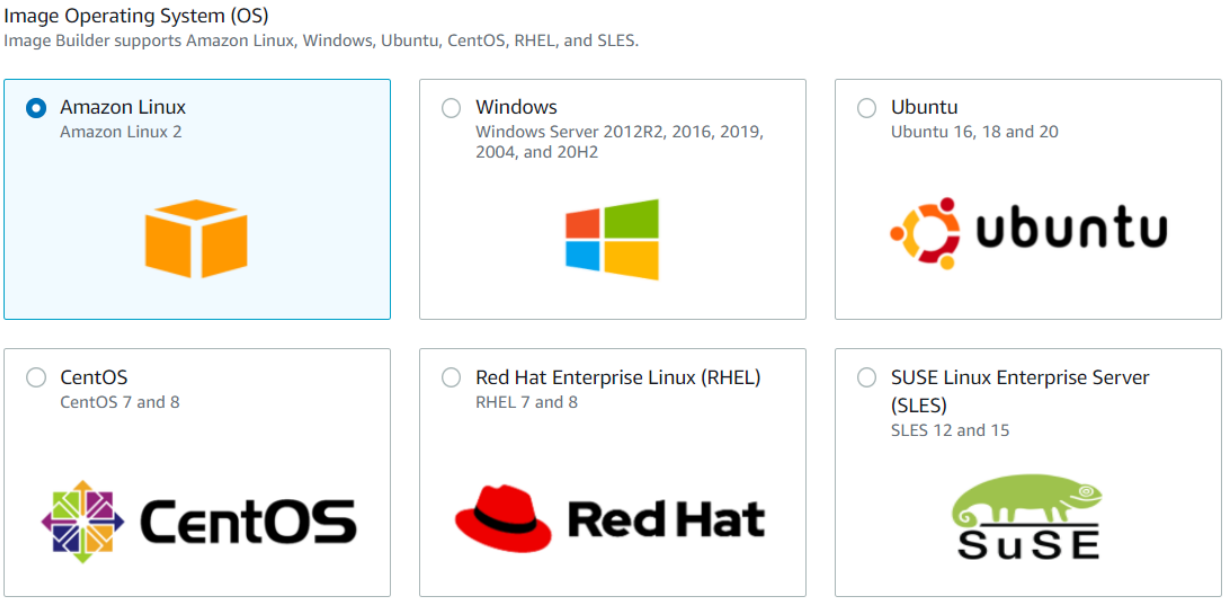

Image builder installs SSM Agent during the build process, but you can remove the agent after the pipeline execution.

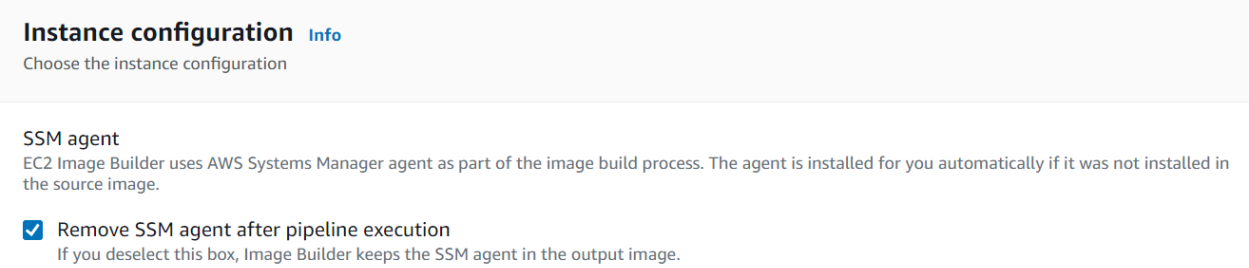

Should it be necessary to run a command on the instance launch, you can set it on the User Data. Note that defining User data requires your source image to have the SSM Agent pre-installed or that you include SSM Agent installation on the User Data.

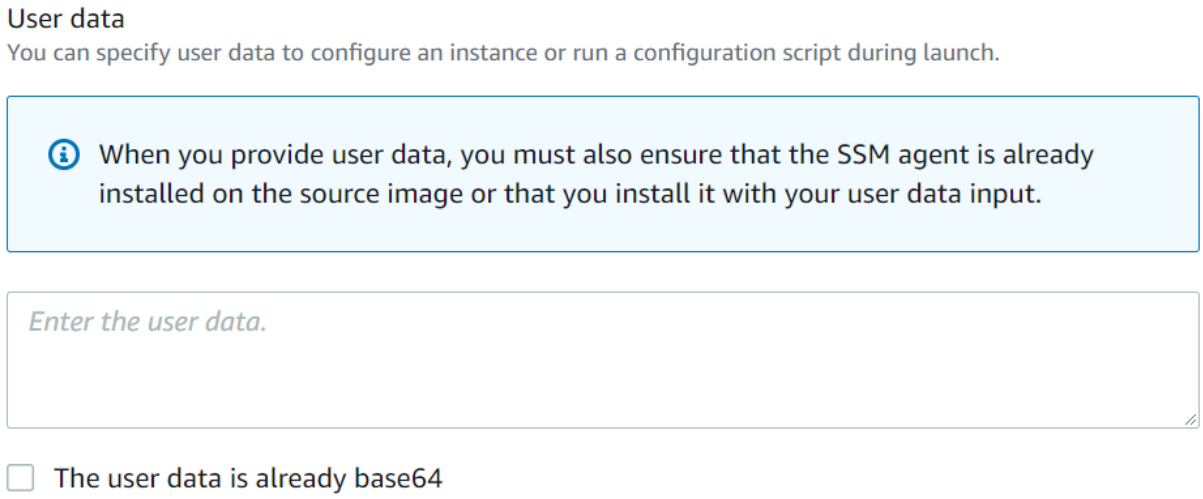

Build and Test Components

A build component installs software packages to your source image. You can select from Amazon-managed build components, share build components to your AWS account, or create a new one. See the example Amazon-managed build components below.

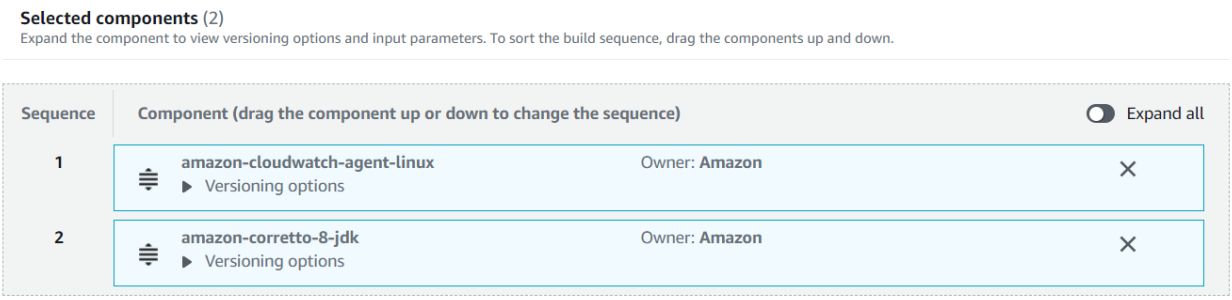

Test Components are optional, but it's a better option to configure this to validate the integrity and functionality of the output image. You can also use Amazon-managed, shared, or create a new test component. See the example Amazon-manage test components below.

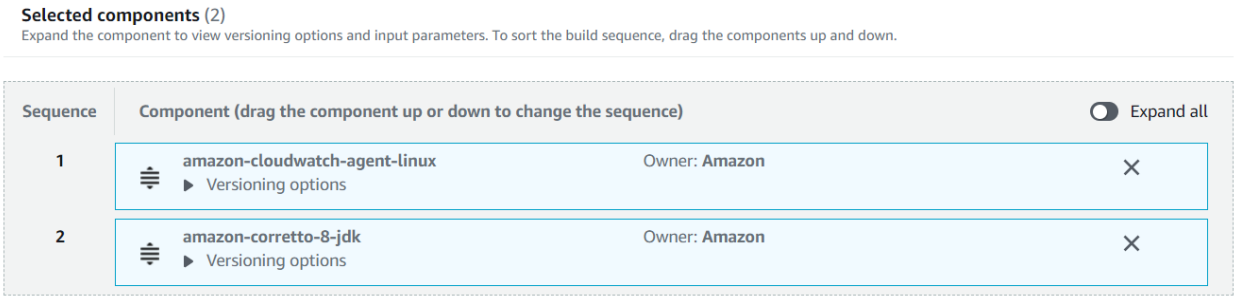

Storage

Storage configuration is optional. You can configure this during the instance launch.

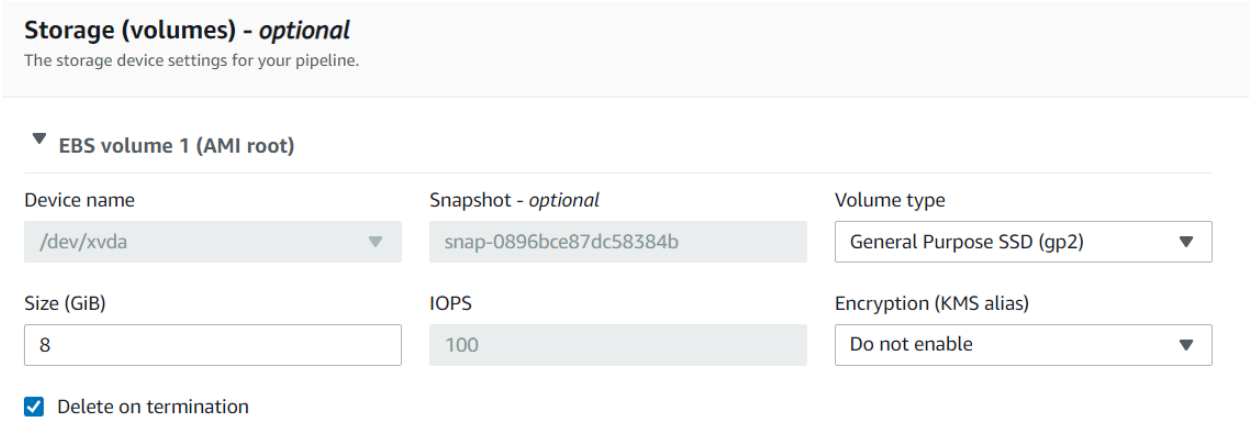

Infrastructure Configuration

The Infrastructure configuration is an optional configuration on the image pipeline. You can configure the Instance Type, VPC settings, IAM role, and Tags for the output image. A notification can also be published using SNS.

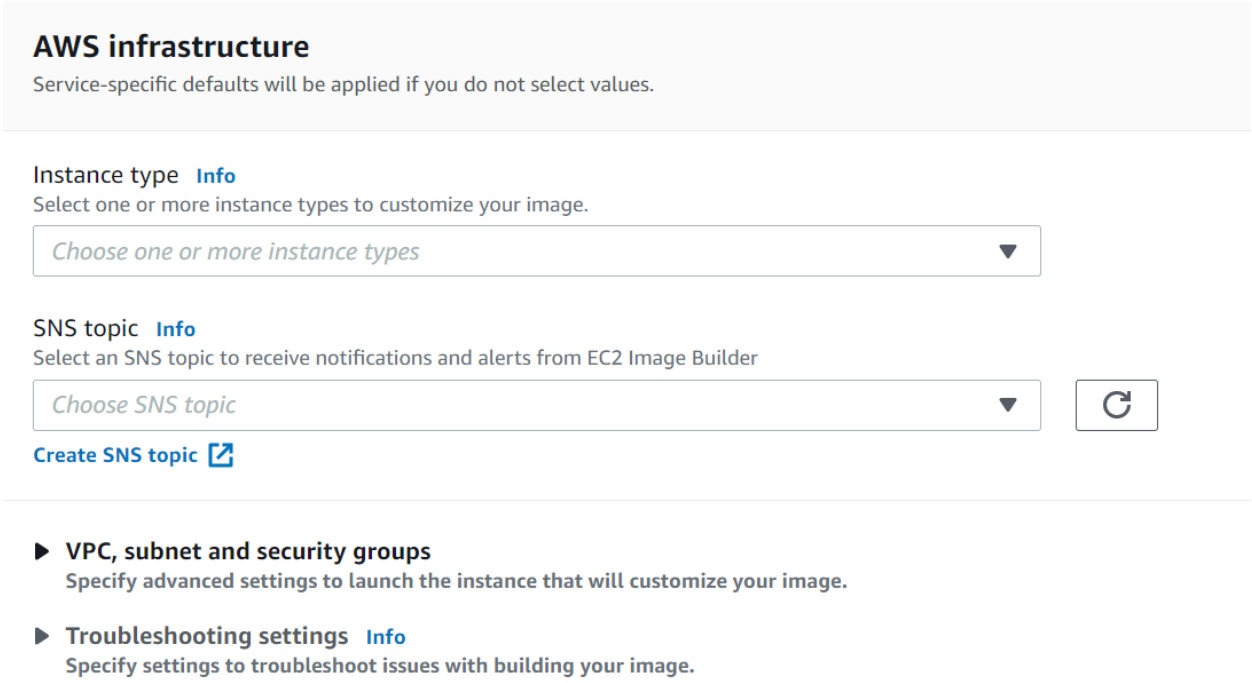

Besides the default IAM policies that the image builder uses, the configured IAM role should also have the necessary permissions to execute all the build and test components defined on the image recipe.

Default IAM Policies for Image Builder:

- *EC2InstanceProfileForImageBuilder*
- *EC2InstanceProfileForImageBuilderECRContainerBuilds*
- *AmazonSSMManagedInstanceCore*

Distribution Settings

You can configure the image deployment on the Distribution settings. You can choose multiple AWS Regions as image destinations. For Amazon Machine Images, you can configure the output image name, AMI sharing, and the license and launch template configuration. For the docker images, you need to specify the Regions and ECR repository name.

Reference:
https://docs.aws.amazon.com/imagebuilder/latest/userguide/what-is-image-builder.html

Amazon EC2Rescue

While AWS takes care of the underlying infrastructure for EC2, customers are responsible for configuring, maintaining, and troubleshooting their instances.

EC2Rescue for Windows Server

EC2Rescue for Windows Server is a downloadable tool for Windows Server instances to help you diagnose and troubleshoot issues. You can also use EC2Rescue to detect potential problems in your current instances.

Diagnose and Rescue an Offline Instance

EC2Rescue scans and diagnoses the Amazon EBS root volumes of the problematic instances. To do this, EC2Rescue requires a host instance where it will be installed. The EBS root volume should be detached from the problematic instance and attached to the EC2Rescue instance host.

Reminders:

- The EC2Rescue tool only runs on Windows Server 2008 R2 or later and requires .NET Framework 3.5 SPI or later.
- The EC2Rescue instance host should also be accessible using an RDP connection.
- The instance where EC2Rescue is installed and the instance to be diagnosed should reside on the same Availability Zone.

The following instructions will guide you on how to check an instance using EC2Rescue.

1. Connect to the EC2Rescue host and download the tool here using a browser or using the PowerShell command below.

```
Invoke-WebRequest https://s3.amazonaws.com/ec2rescue/windows/EC2Rescue_latest.zip -OutFile $env:USERPROFILE\Desktop\EC2Rescue_latest.zip
```

2. Extract the downloaded zip file. Once extracted, run the EC2Rescue application.

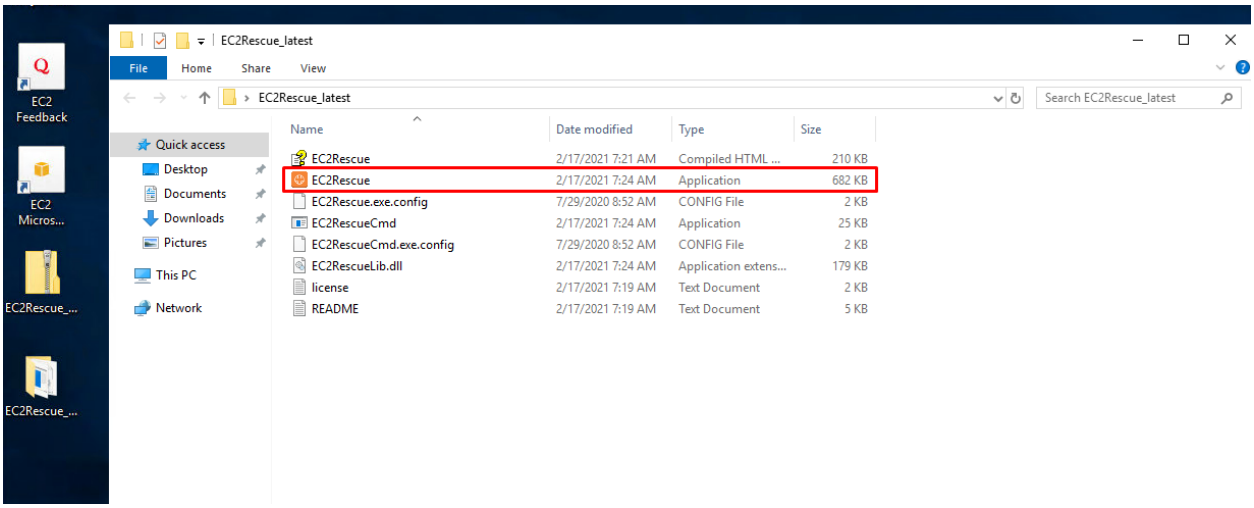

3. Click ***I agree*** on the license agreement and click ***Next*** on the Welcome screen.

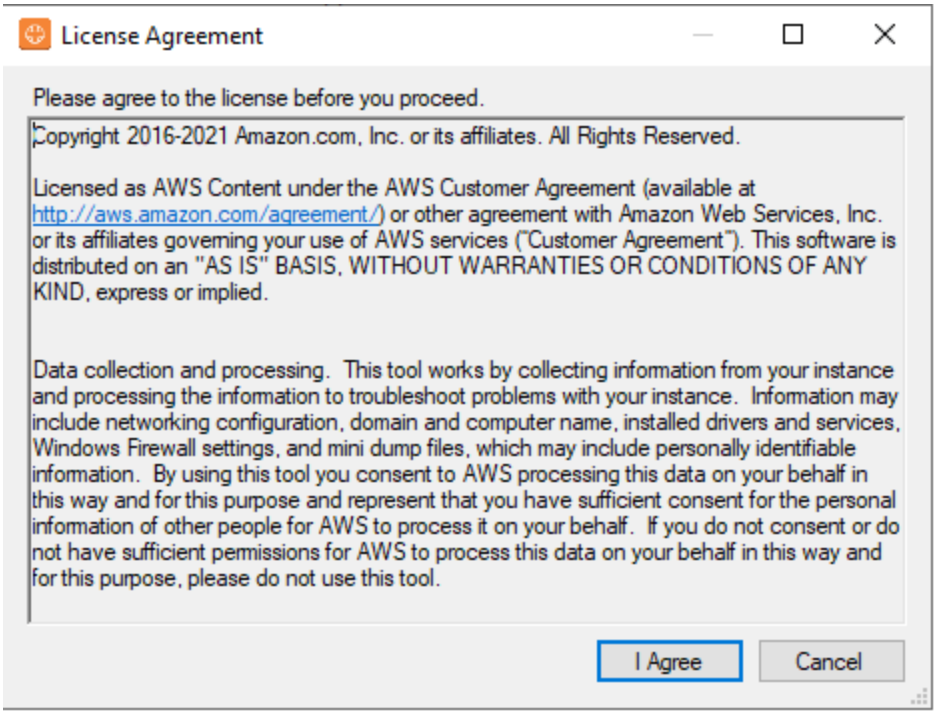

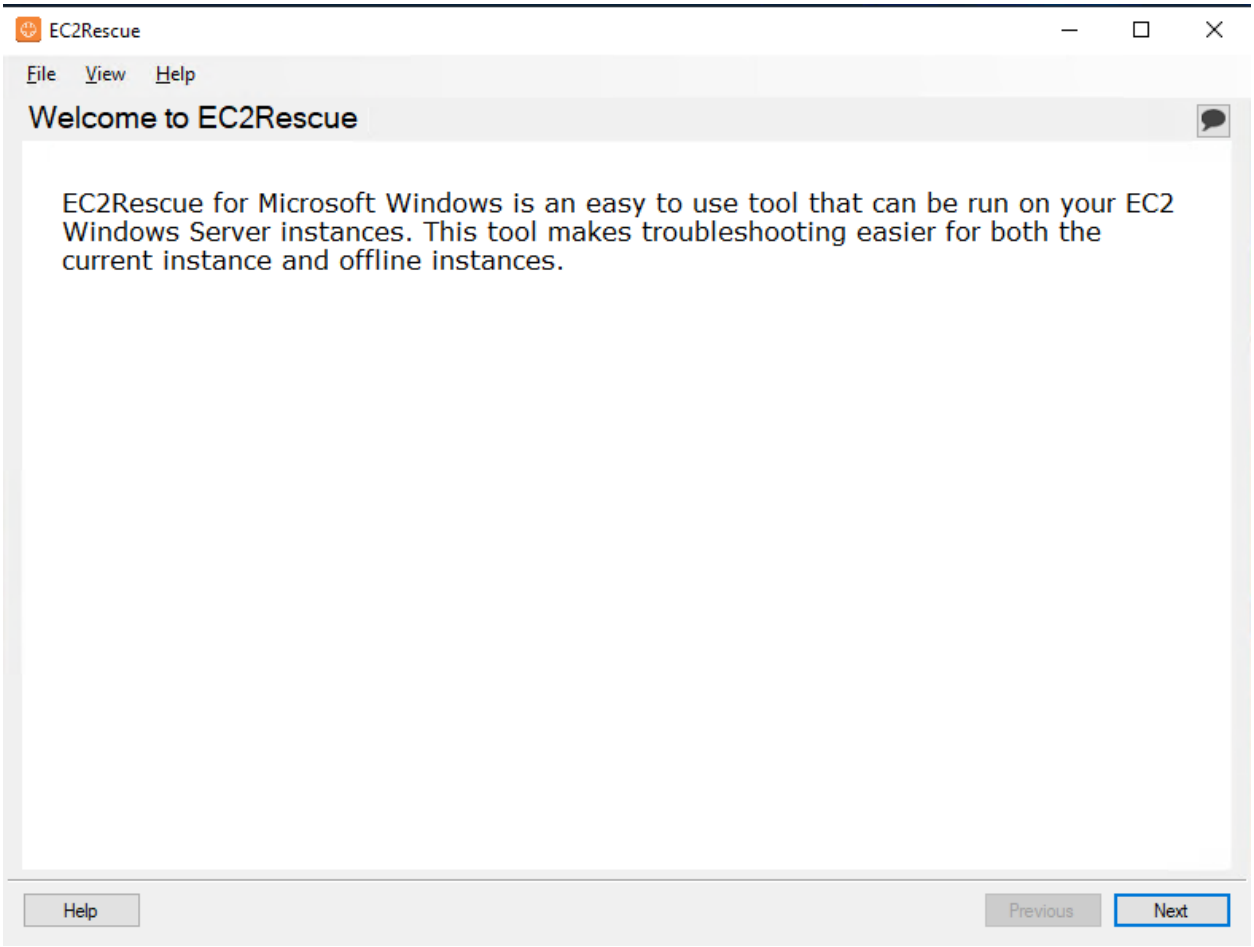

4. Select ***Offline instance*** and click ***Next***.

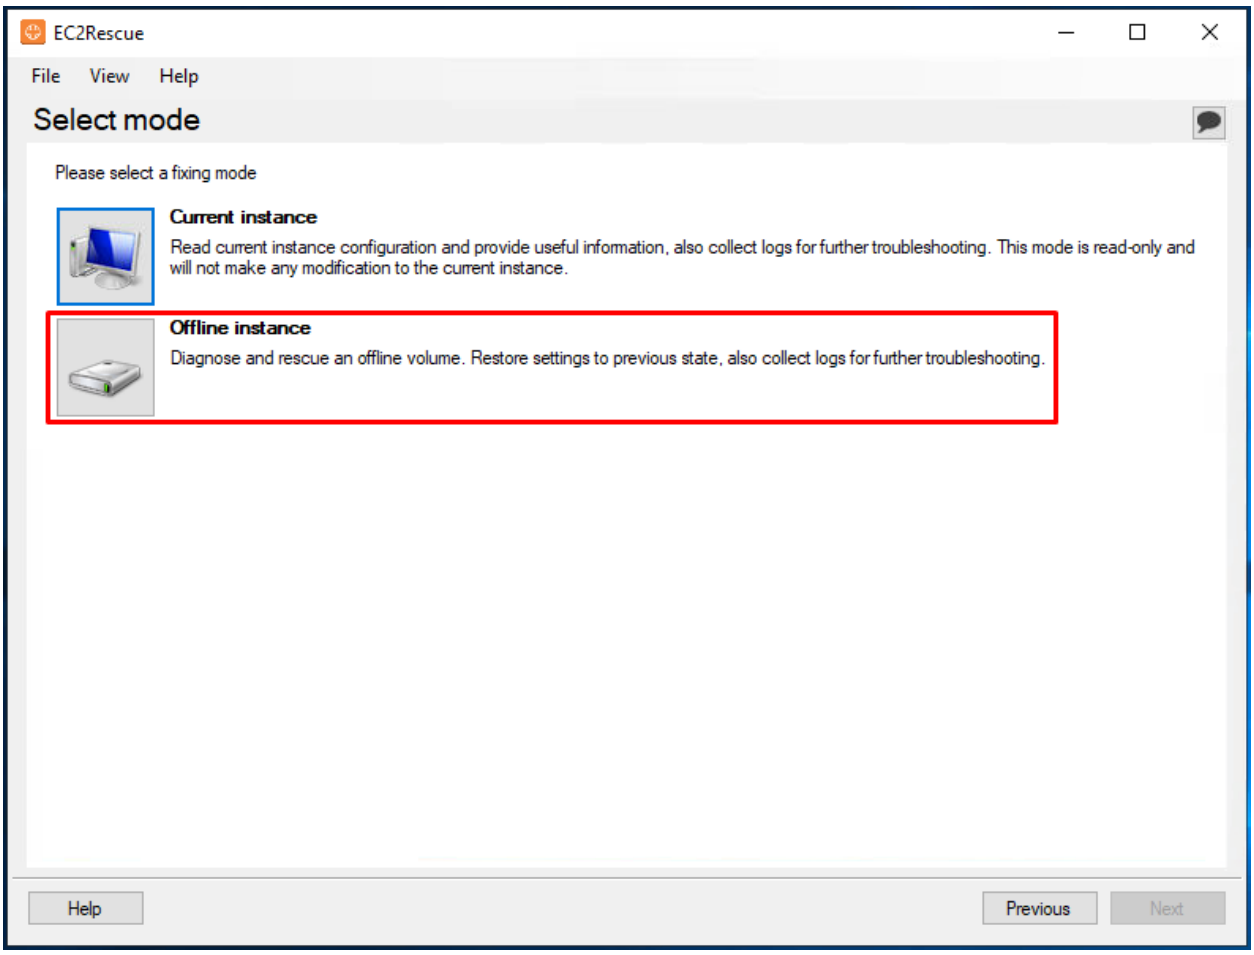

5. Select the ***EBS volume*** of the problematic instance and click ***Next***. If you are checking multiple root volumes, note the device name when attaching the volumes to the EC2Rescue host.

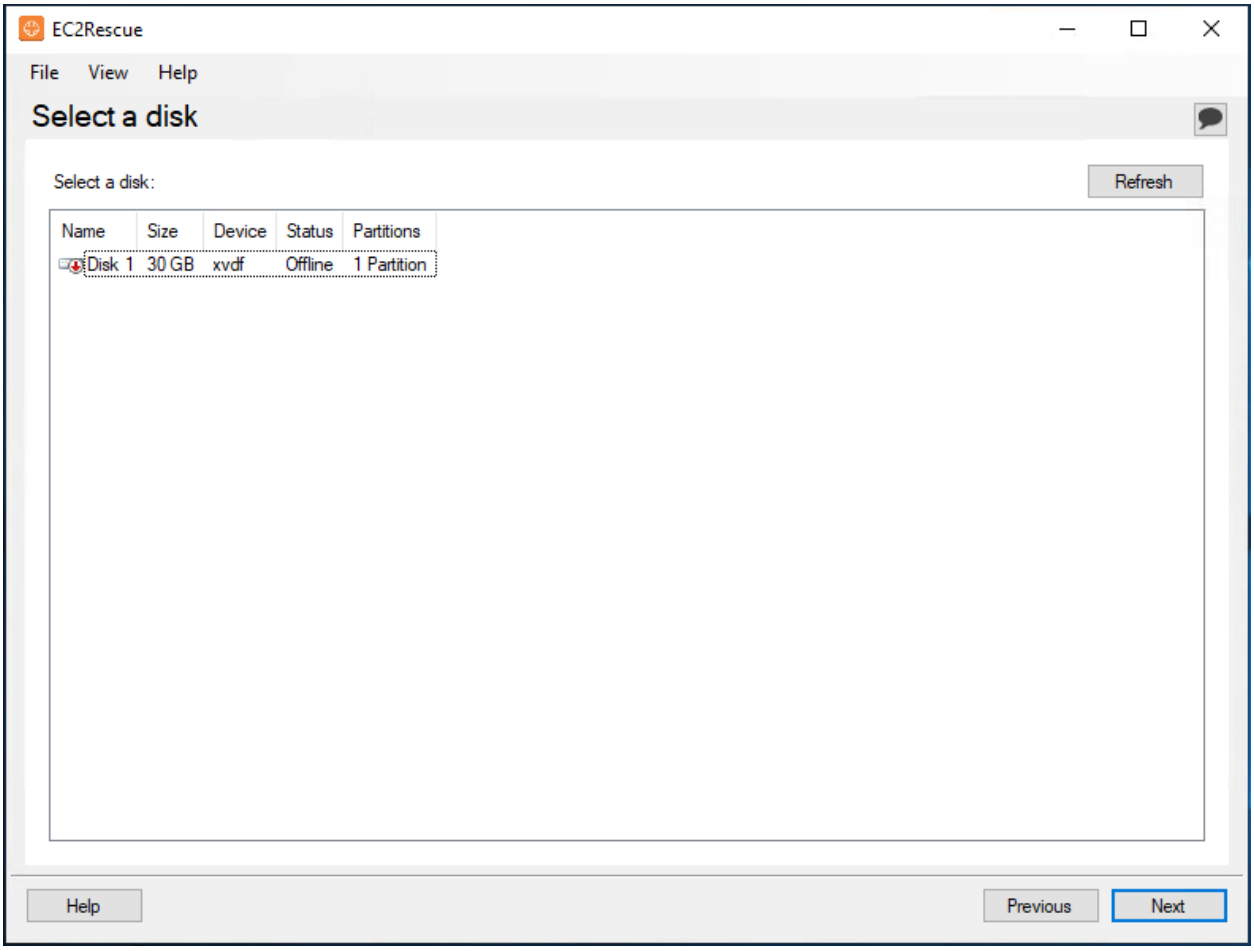

6. Click ***Yes*** to confirm. A popup window will show once the EBS volume is successfully loaded.

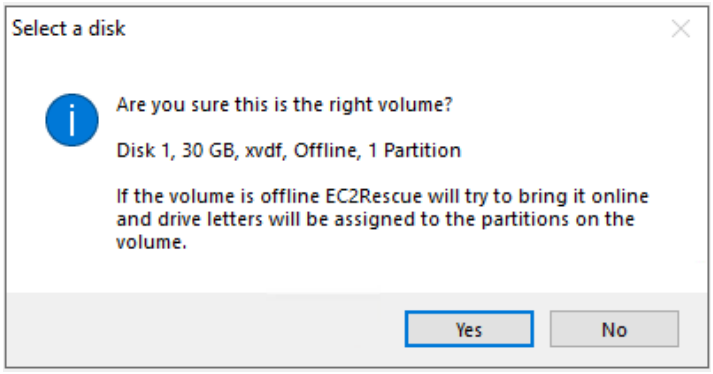

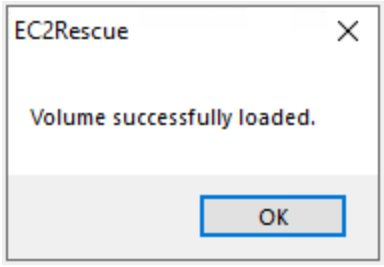

7. Once the volume is loaded, the EC2Rescue tool will display system information of the instance. You will also see different offline instance options. In this case, select ***Diagnose and Rescue*** to proceed.

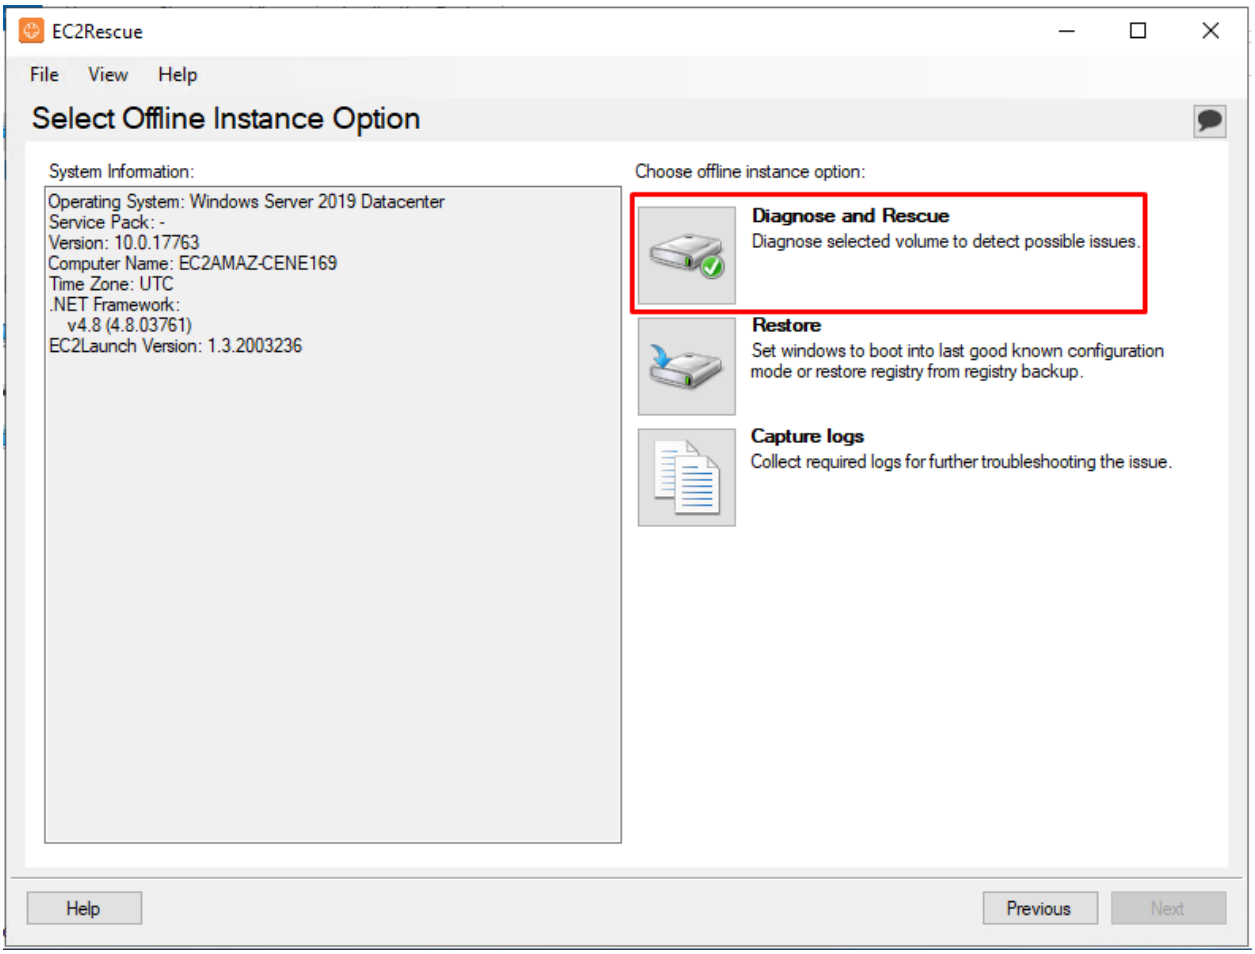

8. The EC2Rescue tool will now start scanning and diagnosing the volume. Once the diagnostic is done, it will summarize necessary configurations, including their status and description. Click ***Next*** to proceed.

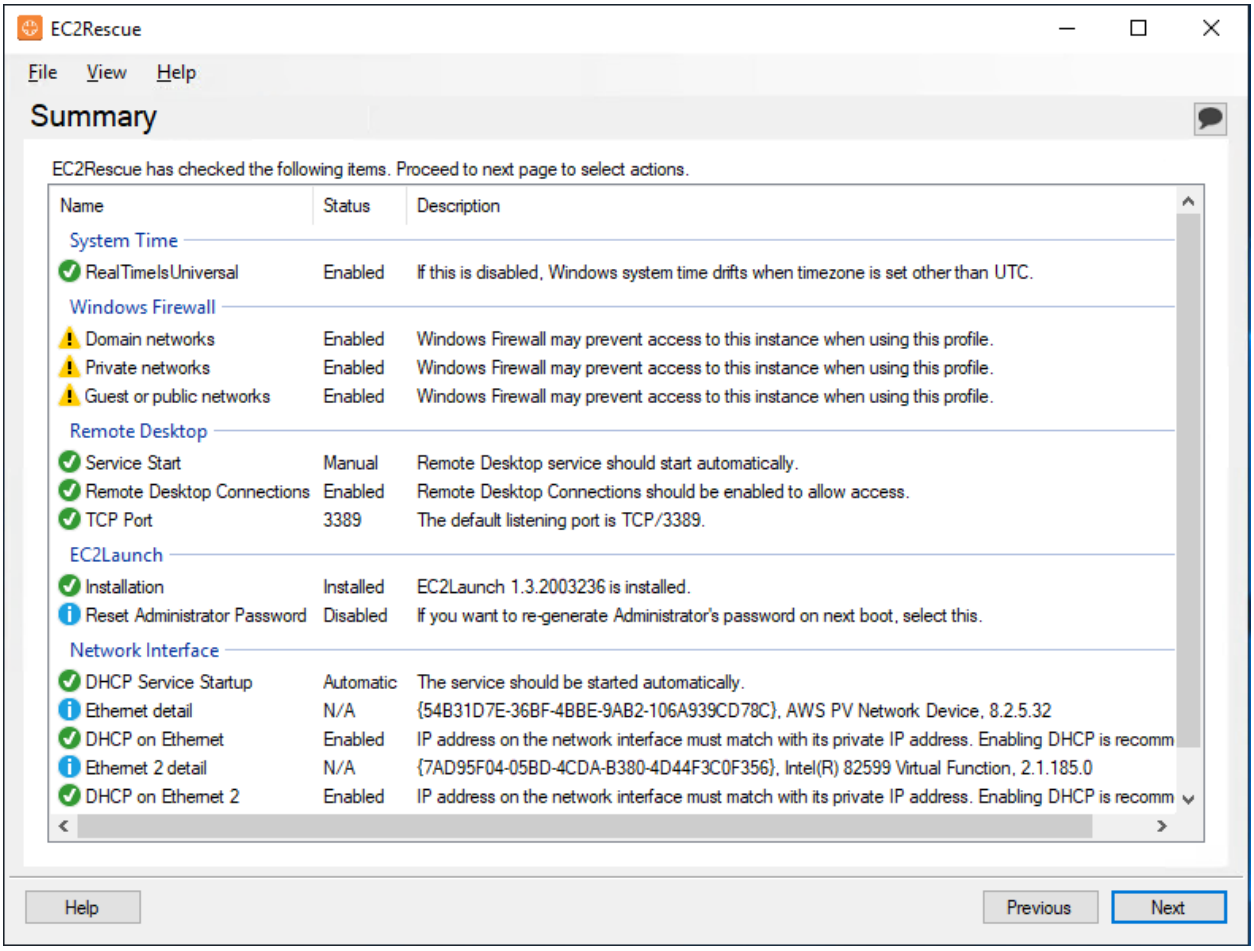

9. EC2Rescue will give you a list of potential issues of the instance. From this, you can select the fixes you find necessary for your instance. Click ***Next*** to proceed.

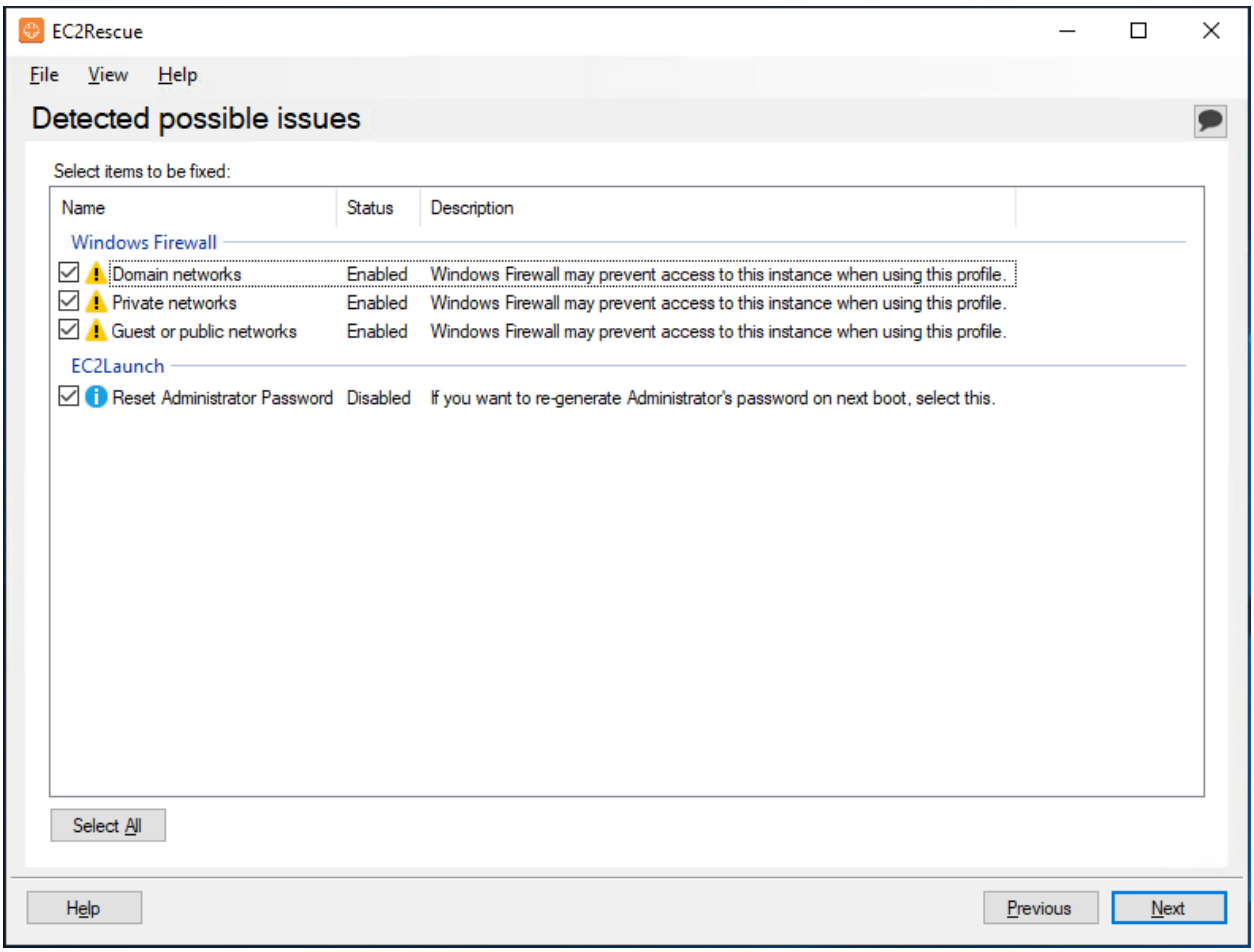

10. Once confirmed, click ***Rescue***.

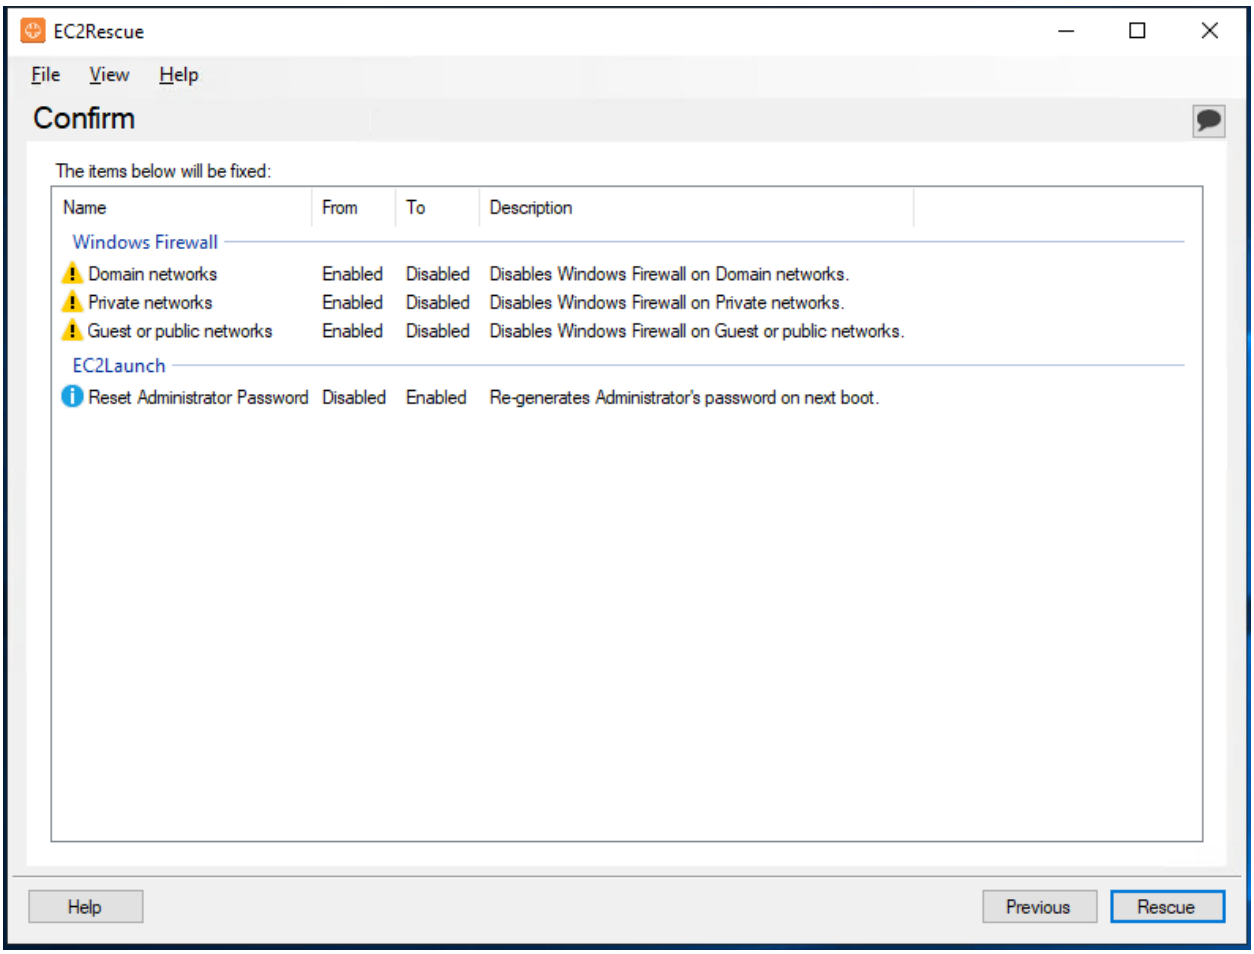

11. Click ***Next*** to continue applying the changes.

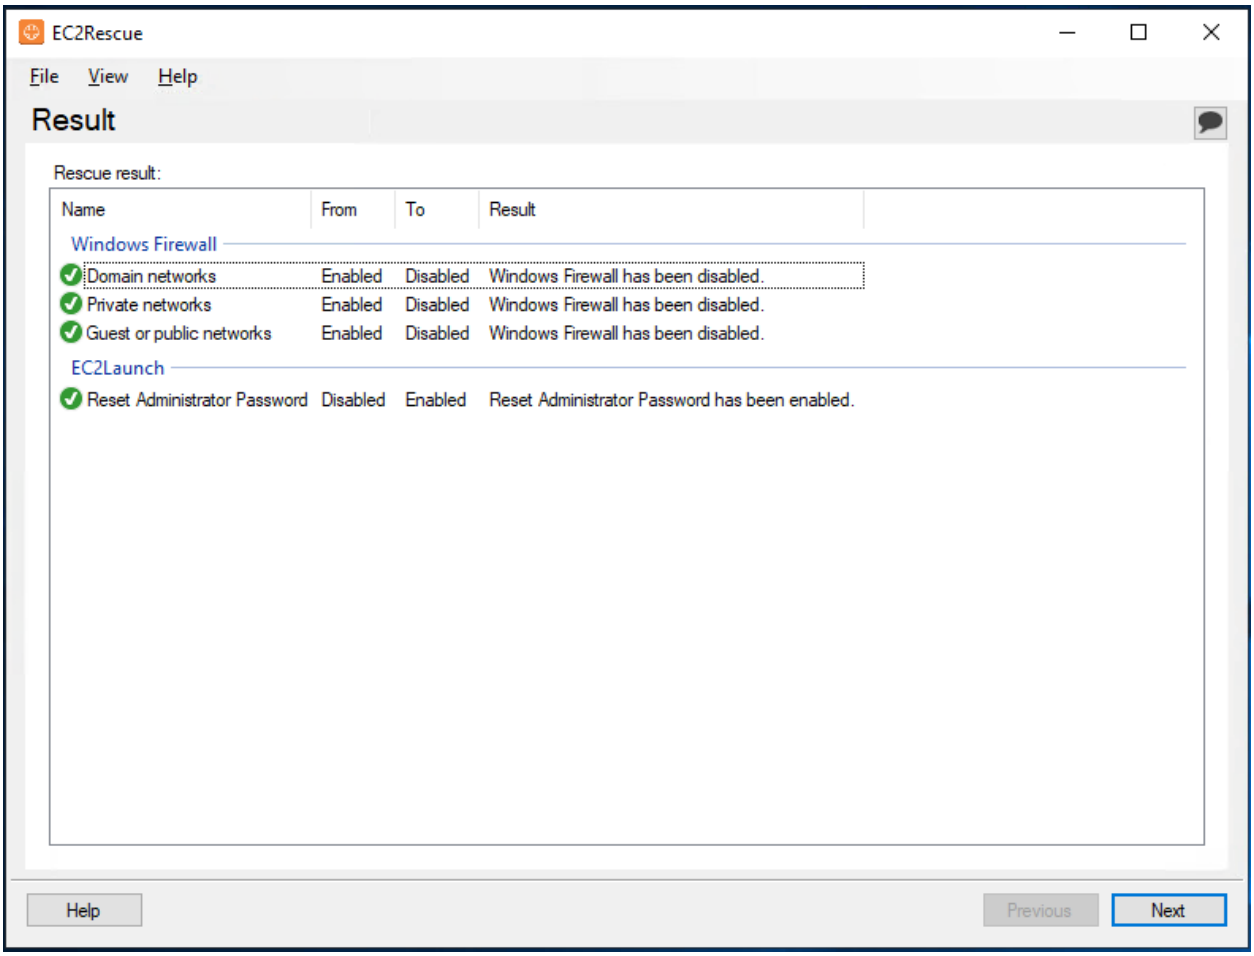

12. Click ***Finish***.