

AWS Academy Learner Lab – Foundation Services

For purposes of your use of the Academy Learner Lab feature, the following restrictions apply.

Region restriction

All service access is limited to the **us-east-1** and **us-west-2** Regions. If you load a service console page in another AWS Region you will see access error messages.

Service usage and other restrictions

The following services can be used. Specific limitations apply as documented. Services restrictions are subject to change.

- Application Auto Scaling
- Amazon Aurora
- AWS Cloud9
 - Supported Instance types: nano, micro, small, and medium.
- Amazon CloudFormation
- Amazon CloudFront
- AWS CloudShell
- AWS CloudTrail
- Amazon CloudWatch
- AWS CodeCommit
- Amazon Cognito
- Amazon Comprehend
- AWS DeepRacer
- Amazon DynamoDB
- Amazon EC2 Auto Scaling
 - Supported Instance types: nano, micro, small, medium, and large.
 - Maximum of 32 vCPU used by concurrently running instances in an AWS Region. For example, t2.micro instances use 1 vCPU each, so you could run up to 32 of them. However, t3.large instances use 2 vCPUs each, so you can run up to 16 of them. You can run a mix of instance types as long as you do not exceed the 32 vCPU threshold. Note that you are also limited to launching no more than nine (9) instances (of any size) in a Region at once. *Recommendation:* size to your actual need to avoid using up your cost budget.
- AWS Elastic Beanstalk
 - Supported Instance types: nano, micro, small, medium, and large. If you attempt to launch a larger instance type, it will be terminated.
 - When you first create an environment in the console, it will use the default security settings. However after the environment has been created, you may want to edit the configuration's security settings. Change the service role to **LabRole**. Similarly, set the IAM instance profile to **LabInstanceProfile**. If the environment

is in the us-east-1 AWS Region, you could set the EC2 key pair to **vockey**. These settings will give you more permissions than the defaults.

- Amazon Elastic Block Store (EBS)
 - Maximum volume size is 100GB
 - PIOPs not supported
- Amazon Elastic Compute Cloud (EC2)
 - AMIs - Amazon provided Linux and Windows AMIs only.
 - Supported Instance types - nano, micro, small, medium, and large.
 - Instance quantity - Maximum of 32 vCPU used by concurrently running instances in an AWS Region. For example, t2.micro instances use 1 vCPU each, so you could run up to 32 of them. However, t3.large instances use 2 vCPUs each, so you can run up to 16 of them. You can run a mix of instance types as long as you do not exceed the 32 vCPU threshold. Note that you are also limited to launching no more than nine (9) instances (of any size) in a Region at once. *Recommendation:* size to your actual need to avoid using up your cost budget.
 - On-Demand instances only
 - EBS volumes - sizes up to 100 GB and type must be General Purpose SSD (gp2, gp3) cold HDD (sc1), or standard.
 - Key pairs - If you are creating an EC2 instance in any AWS Region other than us-east-1, the vockey key pair will not be available. In such cases, you should create a new key pair and download it when creating the EC2 instance. Then use the new key pair to connect to that instance.
 - A role named **LabRole** and an instance profile named **LabInstanceProfile** have been pre-created for you. You can attach the role (via the instance profile) to an EC2 instance when you want to access an EC2 instance (terminal in the browser) using AWS Systems Manager Session Manager. The role also grants permissions to any applications running on the instance to access many other AWS services from the instance.
 - **Tip:** to preserve your lab budget, stop any running EC2 instances before you are done using the account for the day (or terminate them if not longer needed). When your session ends, the lab environment *may* place any running instances into a 'stopped' state. Keep this in mind when you start a new session, that you may need to start the stopped instance(s). Also, instances that have been stopped and started again, will be assigned a new IPv4 public IP address unless you have an elastic IP address associated with the instance.
- Amazon Elastic Container Registry (ECR)
- Amazon Elastic File System (EFS)
- Amazon Elastic Inference
- Elastic Load Balancing
- Amazon EventBridge
- Amazon Forecast
- AWS Glue
- AWS Glue DataBrew
- AWS Identity and Access Management (IAM)
 - Extremely limited access. You cannot create users or groups. You cannot create roles, except that you can create service-linked roles.

- Service role creation is generally permitted. If the service needs to create a role for you, you may need to retry role creation if it fails the first time.
- A role named **LabRole** has been pre-created for you. This role is designed to be used when you want to attach a role to a resource in an AWS service. It grants many AWS services access to other AWS services and has permissions very similar to the permissions you have as a user in the console.
 - Example use: attach the LabRole via the instance profile named **LabInstanceProfile** to an EC2 instance for terminal in the browser access to an EC2 instance guest OS using AWS Systems Manager Session Manager.
 - Another example: Attach the LabRole to a Lambda function so that the Lambda function can access an S3, CloudWatch, RDS, or some other service.
 - Another example: Attach the LabRole to a SageMaker notebook instance so that the instance can access files in an S3 bucket.
- AWS Key Management Service (KMS)
- AWS Lambda
 - **Tip:** Attach the existing **LabRole** to any function that you create if that function will need permissions to interact with other AWS services.
- Amazon Lex
- Amazon Lightsail
 - If you choose vCPU and memory specs that are too high (such as as 8 vCPU and 32GB) the instance may be terminated. Smaller sizes are supported.
- Amazon Marketplace Subscriptions (Amazon ML)
 - Extremely limited read-only access.
- Amazon Polly
- Amazon Rekognition
- Amazon Relational Database Service (RDS)
 - Supported instance types: nano, micro, small, and medium.
 - Supported database engines: Amazon Aurora, MySQL, PostgreSQL and MariaDB.
 - EBS volumes - size up to 100 GB and type General Purpose SSD (gp2).
 - On-Demand DB instance class types only
 - Multi-AZ deployments are not supported (choose Dev/Test or Free tier template if prompted and do not create a standby instance).
 - Enhanced monitoring is not supported (you must *uncheck* this default setting).
 - Tip: to preserve your lab budget, stop any running RDS instances before you are done using the account for the day (or terminate them if not longer needed). Be aware that if you do stop an RDS instance and leave it stopped for seven days, AWS will start it again automatically, which will increase the cost impact.
- AWS Resource Groups & Tag Editor
- AWS RoboMaker
 - Supported Instance types for development environments: *nano*, *micro*, *small*, *medium*, *large*, and *c4.xlarge* only.
- Amazon SageMaker
 - Supported instance types: *medium*, *large*, and *xlarge* only.

- AWS Secrets Manager
- AWS Security Token Service (STS)
- AWS Service Catalog
- Amazon Simple Notification Service (SNS)
- Amazon Simple Queue Service (SQS)
- Amazon Simple Storage Service (S3)
- Amazon Simple Storage Service Glacier (S3 Glacier)
 - You cannot create a vault lock
- AWS Step Functions
- AWS Systems Manager (SSM)
 - A role named **LabRole** and an instance profile named **LabInstanceProfile** have been pre-created for you. You can attach the role (via the instance profile) to an EC2 instance when you want to access an EC2 instance (terminal in the browser) using AWS Systems Manager Session Manager.
- Amazon Textract
- Amazon Translate
- AWS Trusted Advisor
- Amazon Virtual Private Cloud (Amazon VPC)
- AWS Well-Architected Tool