databricks

Databricks AWS Automated Configuration Workspace Deployment

Walkthrough

Shared Technical Services
Last Updated: 2026-1-13

Deploying Databricks Classic Workspaces on AWS using Automated Configuration

- This guide explains the process to deploy classic Databricks workspaces on AWS using the Automated Configuration flow
- Automated Configuration uses AWS IAM Temporary delegation to automatically provision all of the required resources that are required for a fully functional Databricks workspace
- This automated flow prevents common configuration errors and provides built-in approval workflows for users who need AWS admin authorization, if they do not have the required permissions to create the necessary AWS resources
- All automated actions and activities are logged in AWS CloudTrail

High-level Process

To deploy Databricks Workspaces using Automated Configuration

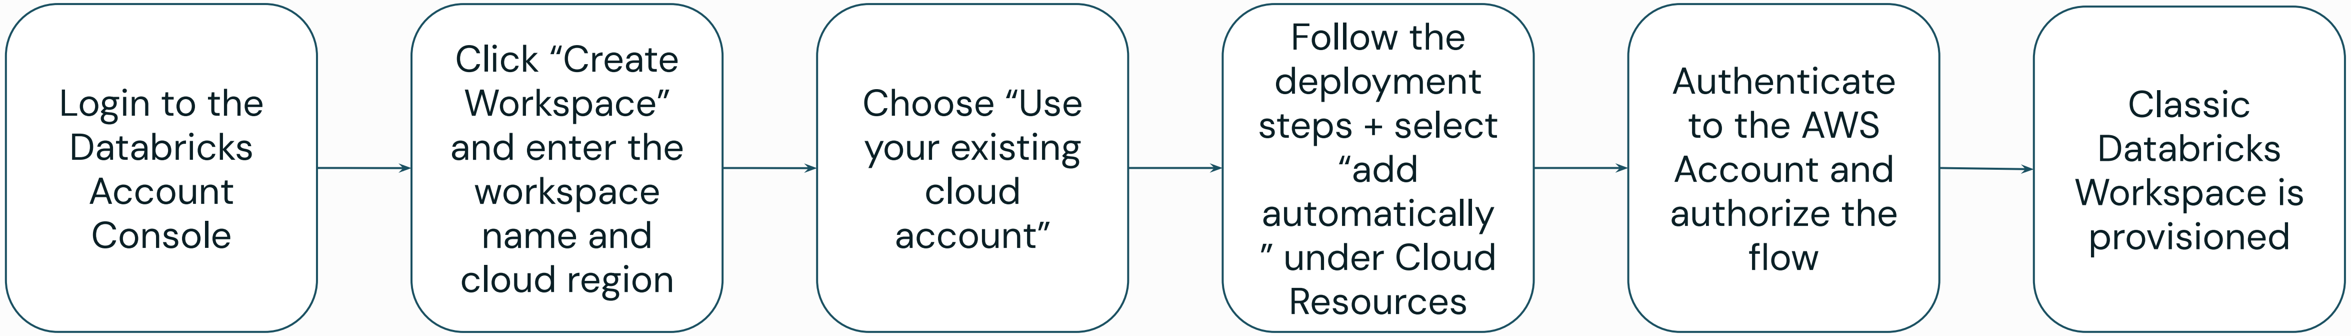

Login to the Databricks Account Console

Click “Create Workspace” and enter the workspace name and cloud region

Choose “Use your existing cloud account”

Follow the deployment steps + select “add automatically ” under Cloud Resources

Authenticate to the AWS Account and authorize the flow

Classic Databricks Workspace is provisioned

Appendix I

Detailed Deployment Steps

Login To The Databricks Account Console

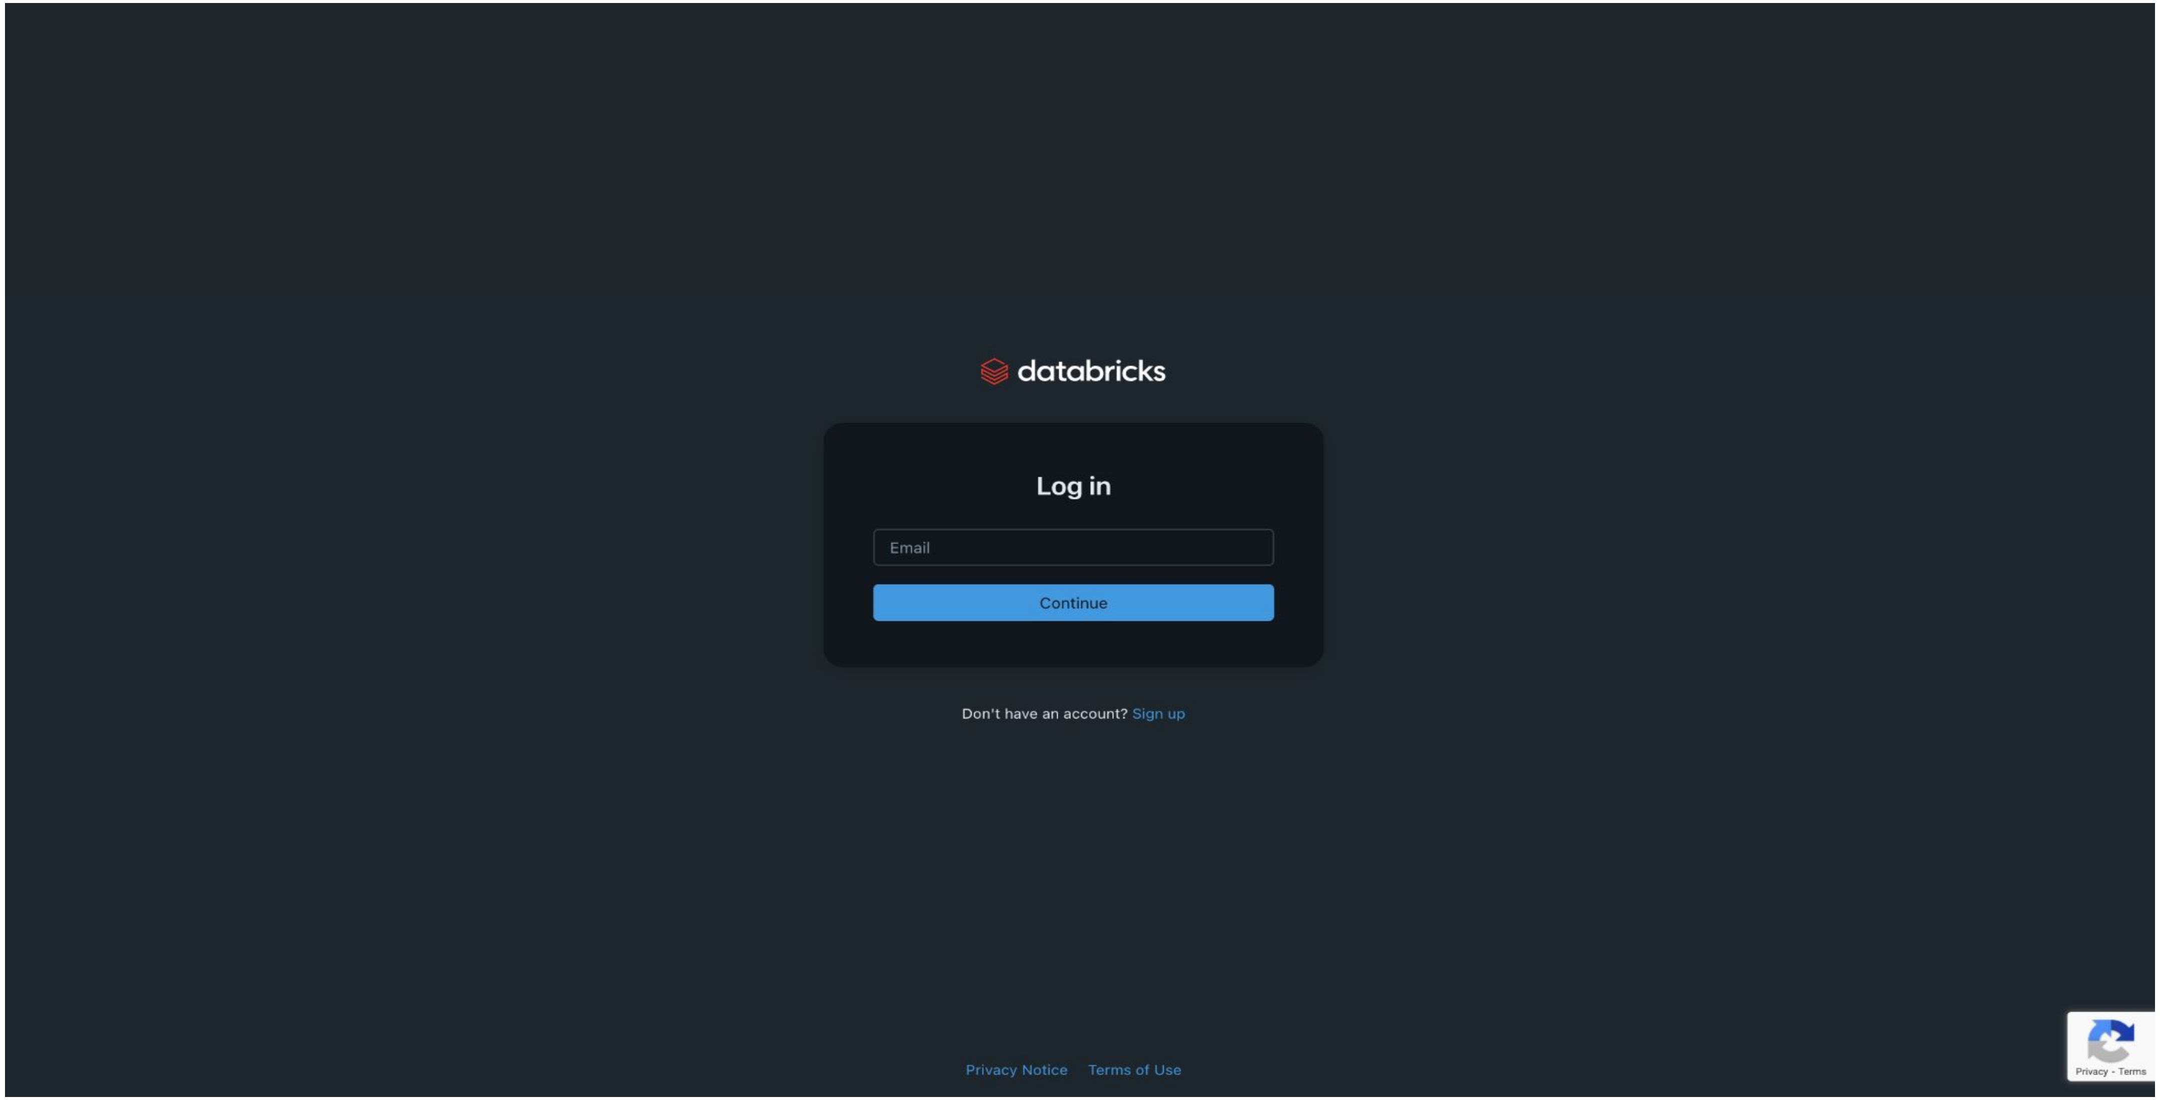

databricks

Log in

Email

Continue

Don't have an account? Sign up

Privacy Notice Terms of Use

Create Workspace

On the main page, select “Workspaces” and click on “Create Workspace” in the top right corner

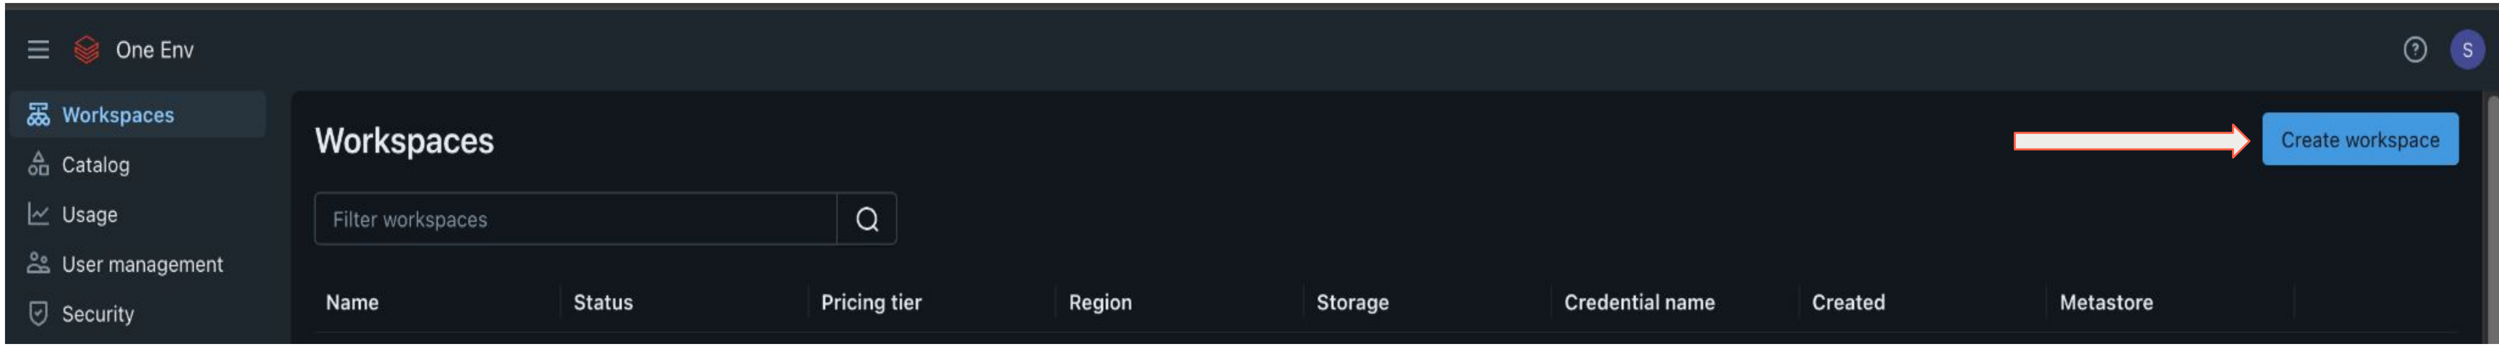

Name	Status	Pricing tier	Region	Storage	Credential name	Created	Metastore

Create Workspace

Enter a workspace name and choose the cloud region. Select “Use your existing cloud account” to provision a classic workspace.

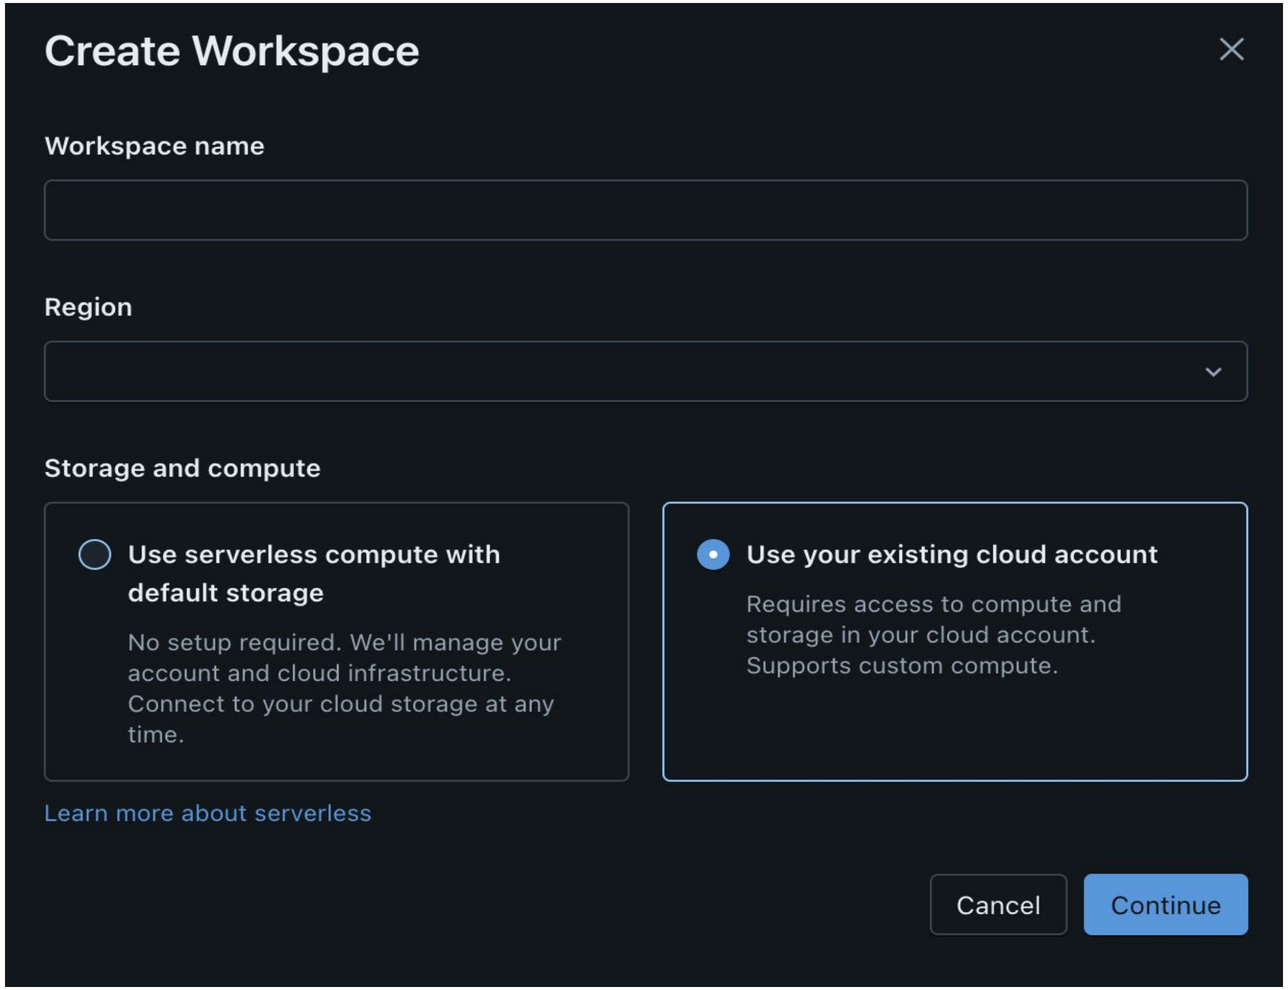

Create Workspace

Workspace name

Region

Storage and compute

○ **Use serverless compute with default storage**

No setup required. We'll manage your account and cloud infrastructure. Connect to your cloud storage at any time.

● **Use your existing cloud account**

Requires access to compute and storage in your cloud account. Supports custom compute.

Learn more about serverless

Add Cloud Credentials

Select the “Add automatically” option to allow automated configuration to provision the required AWS resources for workspace deployment

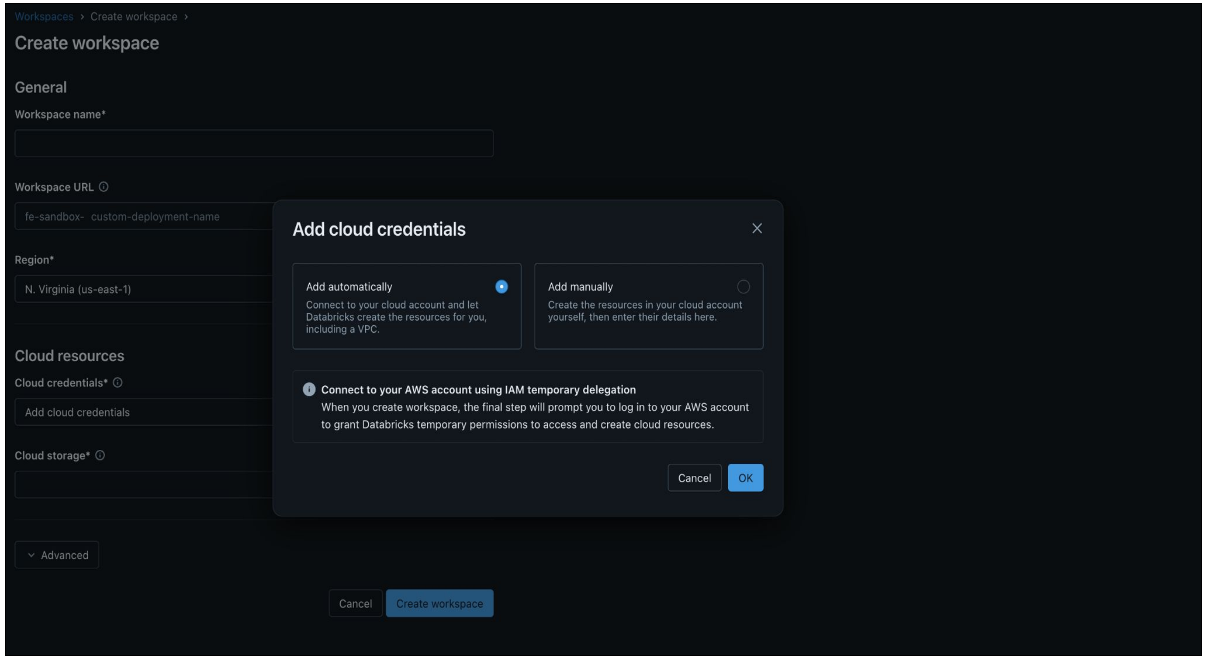

Workspaces › Create workspace ›

Create workspace

General

Workspace name*

Workspace URL

fe-sandbox- custom-deployment-name

Region*

N. Virginia (us-east-1)

Cloud resources

Cloud credentials*

Add cloud credentials

Cloud storage*

Add cloud credentials

● Add automatically
Connect to your cloud account and let Databricks create the resources for you, including a VPC.

○ Add manually
Create the resources in your cloud account yourself, then enter their details here.

Connect to your AWS account using IAM temporary delegation
When you create workspace, the final step will prompt you to log in to your AWS account to grant Databricks temporary permissions to access and create cloud resources.

1. Under the dropdown for “Cloud credentials”, choose “Add cloud credentials”.
2. Select “Add automatically” for the automated configuration flow using AWS IAM Permission Delegation.
3. Under the “Advanced” tab, make sure to select the metastore to be used by this workspace.
4. Click “Log in to AWS and create workspace”.

Review the Resources that will be Provisioned and Initiate workspace Creation

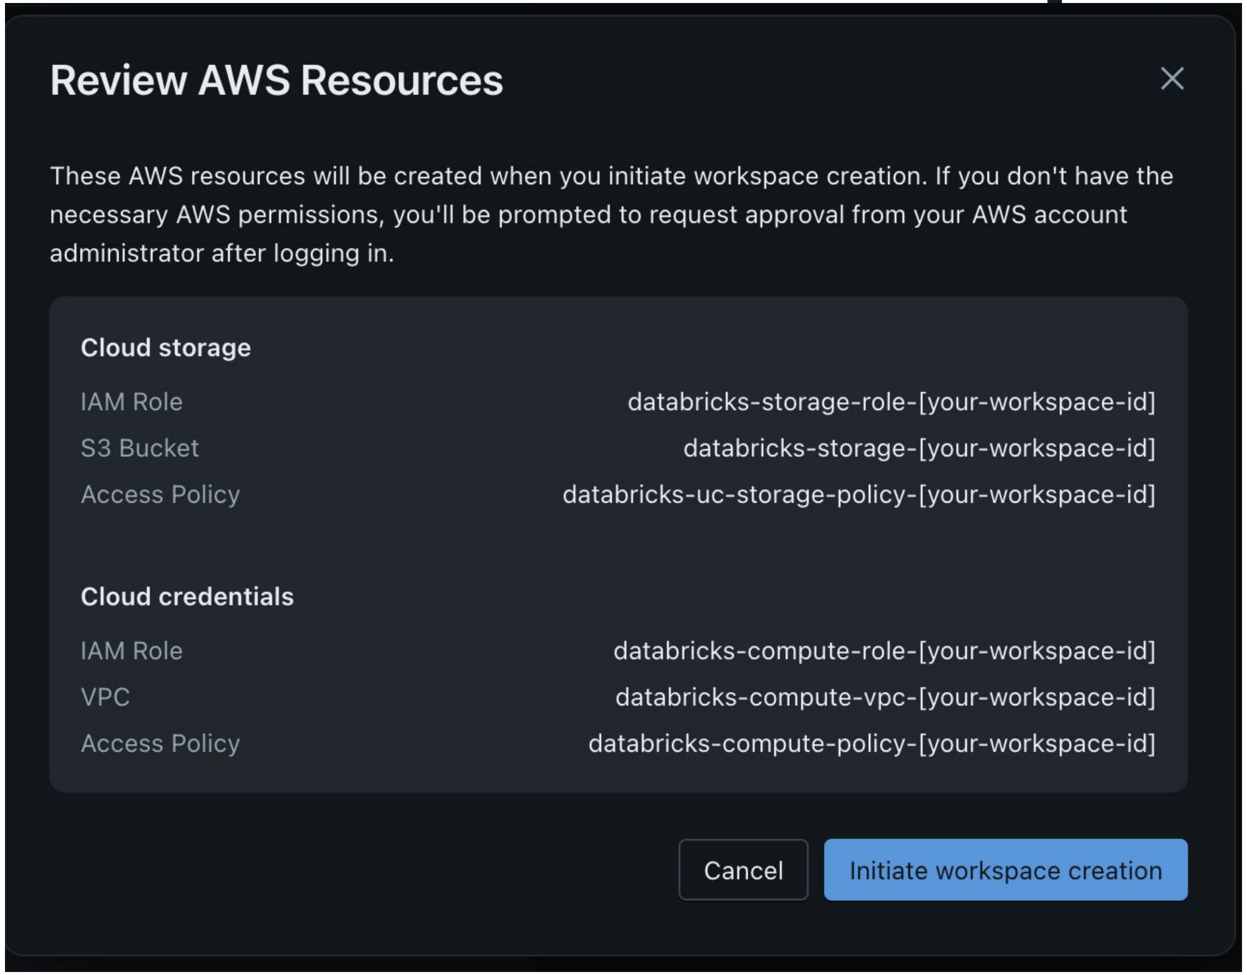

Review AWS Resources

These AWS resources will be created when you initiate workspace creation. If you don't have the necessary AWS permissions, you'll be prompted to request approval from your AWS account administrator after logging in.

Cloud storage

IAM Role	databricks-storage-role-[your-workspace-id]
S3 Bucket	databricks-storage-[your-workspace-id]
Access Policy	databricks-uc-storage-policy-[your-workspace-id]

Cloud credentials

IAM Role	databricks-compute-role-[your-workspace-id]
VPC	databricks-compute-vpc-[your-workspace-id]
Access Policy	databricks-compute-policy-[your-workspace-id]

Approve and Authorize the AWS Resource Provisioning

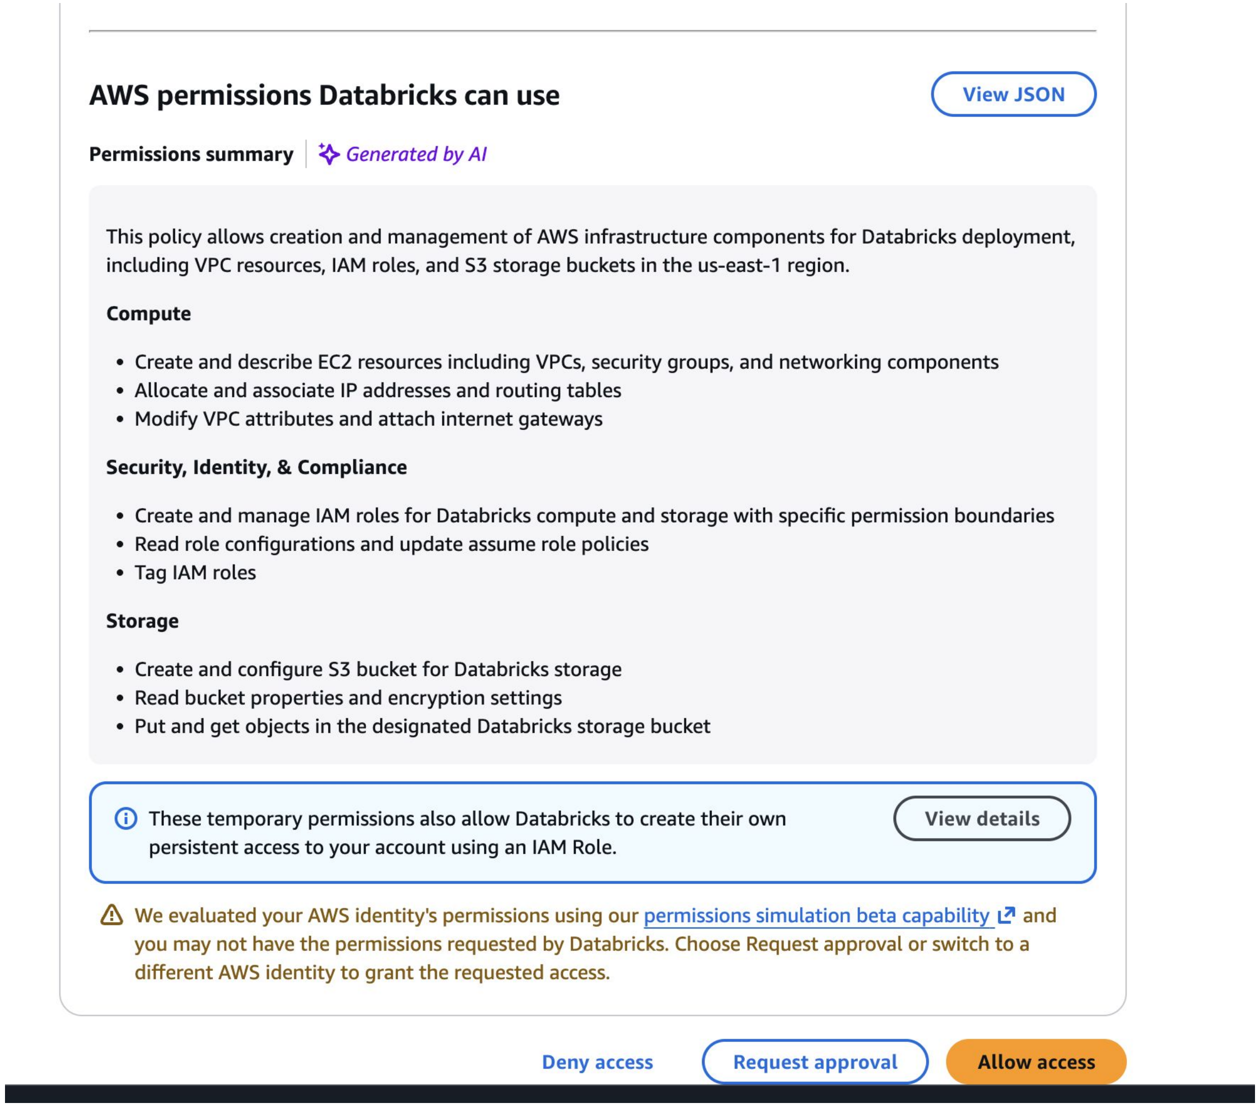

AWS permissions Databricks can use

Permissions summary | *Generated by AI*

This policy allows creation and management of AWS infrastructure components for Databricks deployment, including VPC resources, IAM roles, and S3 storage buckets in the us-east-1 region.

Compute

- Create and describe EC2 resources including VPCs, security groups, and networking components
- Allocate and associate IP addresses and routing tables
- Modify VPC attributes and attach internet gateways

Security, Identity, & Compliance

- Create and manage IAM roles for Databricks compute and storage with specific permission boundaries
- Read role configurations and update assume role policies
- Tag IAM roles

Storage

- Create and configure S3 bucket for Databricks storage
- Read bucket properties and encryption settings
- Put and get objects in the designated Databricks storage bucket

ⓘ These temporary permissions also allow Databricks to create their own persistent access to your account using an IAM Role.

⚠ We evaluated your AWS identity's permissions using our permissions simulation beta capability and you may not have the permissions requested by Databricks. Choose Request approval or switch to a different AWS identity to grant the requested access.

- In the AWS account pop-up, approve and authorize the AWS permissions that Databricks can use to provision the required resources.
- If your IAM identity/role does not have the required permissions to deploy the resources, click on “Request approval” to have the proposed permissions and changes be reviewed by an AWS admin.
- After Databricks is granted temporary access, the workspace will begin to provision.
- All delegated permissions are time-bounded and automatically expire after deployment.

Provisioned Resources

AWS/Databricks Resources that will be provisioned via Automated Configuration

AWS

- Cross-account IAM role with an access policy
- Customer-managed VPC with default subnets, security groups, and routing tables
- Root workspace S3 bucket to store workspace assets and the workspace default UC catalog
- IAM role with an access policy to allow access to the root workspace S3 bucket

Databricks

- Credential Configuration – represents the cross-account IAM role to manage lifecycle of classic compute instances
- Storage Configuration – represents the S3 bucket and IAM role used for workspace root storage
- Network Configuration – represents the customer managed VPC, subnets, and security group used by the workspace

databricks