

User Guide

Research and Engineering Studio



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

Research and Engineering Studio: 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

Overview	1
Features and benefits	1
Concepts and definitions	3
Architecture overview	5
Architecture diagram	5
AWS services in this product	6
Demo environment	10
Create a one-click demo stack	10
Prerequisites	10
Create resources and input parameters	11
Post deployment steps	12
Plan your deployment	14
Cost	14
Security	14
IAM roles	15
Security groups	15
Data encryption	15
Product security considerations	15
Quotas	18
Quotas for AWS services in this product	19
AWS CloudFormation quotas	19
Planning for resilience	19
Supported AWS Regions	20
Deploy the product	22
Prerequisites	22
Create an AWS account with an administrative user	23
Create an Amazon EC2 SSH key pair	23
Increase service quotas	23
Create a public domain (optional)	24
Create domain (GovCloud only)	24
Provide external resources	25
Configure LDAPS in your environment (optional)	26
Configure a private VPC (optional)	26
Create external resources	38

Step 1: Launch the product	44
Step 2: Sign in for the first time	51
Update the product	53
Major version updates	53
Minor version updates	53
Uninstall the product	55
Using the AWS Management Console	55
Using AWS Command Line Interface	55
Deleting the shared-storage-security-group	55
Deleting the Amazon S3 buckets	56
Configuration guide	57
Managing users and groups	57
Setting up SSO with IAM Identity Center	57
Configuring your identity provider for SSO	61
Setting passwords for users	71
Creating subdomains	71
Create an ACM certificate	72
Amazon CloudWatch Logs	73
Setting custom permission boundaries	74
Configure RES-ready AMIs	79
Prepare IAM role to access RES environment	79
Create EC2 Image Builder component	81
Prepare your EC2 Image Builder recipe	85
Configure EC2 Image Builder infrastructure	87
Configure Image Builder image pipeline	88
Run Image Builder image pipeline	89
Register a new software stack in RES	89
Administrator guide	90
Secrets management	90
Cost monitoring and control	93
Session management	97
Dashboard	99
Sessions	100
Software Stacks (AMIs)	103
Debugging	107
Desktop settings	108

Er	nvironment management	109
	Environment status	110
	Environment settings	111
	Users	111
	Groups	112
	Projects	113
	Permission policy	120
	File Systems	135
	Snapshot management	138
	Amazon S3 buckets	144
Use	the product	161
S	SH access	161
Vi	irtual desktops	161
	Launch a new desktop	162
	Access your desktop	163
	Control your desktop state	165
	Modify a virtual desktop	167
	Retrieve session information	168
	Schedule virtual desktops	168
	VDI autostop	171
SI	nared desktops	173
	Share a desktop	173
	Access a shared desktop	175
Fi	le browser	
	Upload file(s)	175
	Delete file(s)	176
	Manage favorites	177
	Edit files	177
	Transfer files	178
Trou	bleshooting	180
G	eneral Debugging and Monitoring	183
	Useful log and event information sources	183
	Typical Amazon EC2 Console Appearance	188
	Windows DCV debugging	190
	Find Amazon DCV Version Information	191
Ic	sue RunRooks	191

Installation issues	193
Identity management issues	202
Storage	206
Snapshots	211
Infrastructure	212
Launching Virtual Desktops	213
Virtual Desktop Component	217
Env deletion	224
Demo environment	231
Known Issues	233
Known Issues 2024.x	233
Notices	251
Revisions	252

Overview



Important

This version of the User Guide covers release 2024.10 of Research and Engineering Studio on AWS. For the current version, see the Research and Engineering Studio on AWS User Guide.

Research and Engineering Studio (RES) is an AWS supported, open source product that enables IT administrators to provide a web portal for scientists and engineers to run technical computing workloads on AWS. RES provides a single pane of glass for users to launch secure virtual desktops to conduct scientific research, product design, engineering simulations, or data analysis workloads. Users can connect to the RES portal using their existing corporate credentials and work on individual or collaborative projects.

Administrators can create virtual collaboration spaces called projects for a specific set of users to access shared resources and collaborate. Admins can build their own application software stacks (using Amazon Machine Images or AMIs) and allow RES users to launch Windows or Linux virtual desktops, and enable access to project data through shared file-systems. Admins can assign software stacks and file-systems and restrict access to only those project users. Admins can use built-in telemetry to monitor the environment usage and troubleshoot user issues. They can also set budgets for individual projects to prevent overconsumption of resources. As the product is open source, customers can also customize the user-experience of the RES portal to suit their own needs.

RES is available at no additional charge, and you pay only for the AWS resources needed to run your applications.

This guide provides an overview of Research and Engineering Studio on AWS, its reference architecture and components, considerations for planning the deployment, and configuration steps for deploying RES to the Amazon Web Services (AWS) Cloud.

Features and benefits

Research and Engineering Studio on AWS provides the following features:

Features and benefits

Web-based user interface

RES provides a web-based portal that administrators, researchers, and engineers can use to access and manage their research and engineering workspaces. Scientists and engineers do not need to have an AWS account or cloud expertise to use RES.

Project-based configuration

Use projects to define access permissions, allocate resources, and manage budgets for a set of tasks or activities. Assign specific software stacks (operating systems and approved applications) and storage resources to a project for consistency and compliance. Monitor and manage spending on a per-project basis.

Collaboration tools

Scientists and engineers can invite other members of their project to collaborate with them, setting the permissions levels they want those colleagues to have. Those individuals can sign in to RES to connect to those desktops.

Integration with existing identity management infrastructure

Integrate with your existing identity management and directory services infrastructure to enable connection to the RES portal with a user's existing corporate identity and assign permissions to projects using existing user and group memberships.

Persistent storage and access to shared data

To provide users access to shared data across virtual desktop sessions, connect to your existing file systems within RES. Supported storage services include Amazon Elastic File System for Linux desktops and Amazon FSx for NetApp ONTAP for Windows and Linux desktops.

Monitoring and reporting

Use the analytics dashboard to monitor resource usage for instance types, software stacks, and operating system types. The dashboard also provides a breakdown of resource usage by projects for reporting.

Budget and cost management

Link AWS Budgets to your RES projects to monitor costs for each project. If you exceed your budget, you can limit the launch of VDI sessions.

Features and benefits 2

Concepts and definitions

This section describes key concepts and defines terminology specific to Research and Engineering Studio on AWS:

File browser

A file browser is a part of the RES user interface where users who are currently logged-in can view their file system.

File system

The file system acts as a container for project data (often referred to as datasets). It provides a storage solution within a project's boundaries and improves collaboration and data access control.

Global administrator

An administrative delegate with access to RES resources that are shared across a RES environment. Scope and permissions span multiple projects. They can create or modify projects and assign project owners. They can delegate or assign permissions to project owners and project members. Sometimes the same person acts as the RES administrator depending on the size of the organization.

Project

A project is a logical partition within the application that serves as a distinct boundary for data and compute resources; this ensures governance over data flow and prevents sharing data and VDI hosts across projects.

Project-based permissions

Project-based permissions describes a logical partition of both data and VDI hosts in a system where multiple projects can exist. A user's access to data and VDI hosts within a project is determined by their associated role(s). A user must be assigned access (or project membership) for each project to which they require access. Otherwise, a user is unable to access project data and VDIs when they have not been granted membership.

Project member

An end user of RES resources (VDI, storage, etc). Scope and permissions are restricted to the projects they are assigned to. They cannot delegate or assign any permissions.

Concepts and definitions

Project owner

An administrative delegate with access to, and ownership over, a specific project. Scope and permissions are restricted to the project(s) they own. They can assign permissions to project members in the projects they own.

Software stack

Software stacks are <u>Amazon Machine Images (AMI)</u> with RES-specific metadata based on any operating system a user has selected to provision for their VDI host.

VDI hosts

Virtual desktop instance (VDI) hosts allow project members to access project-specific data and compute environments, ensuring secure and isolated workspaces.

For a general reference of AWS terms, see the AWS glossary in the AWS General Reference.

Concepts and definitions

Architecture overview

This section provides an architecture diagram for the components deployed with this product.

Architecture diagram

Deploying this product with the default parameters deploys the following components in your AWS account.

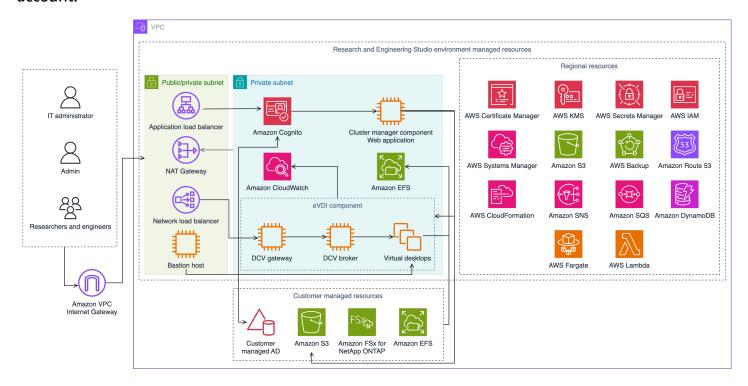


Figure 1: Research and Engineering Studio on AWS architecture



Note

AWS CloudFormation resources are created from AWS Cloud Development Kit (AWS CDK) constructs.

The high-level process flow for the product components deployed with the AWS CloudFormation template is as follows:

- 1. RES installs components for the web portal as well as:
 - a. Engineering Virtual Desktop (eVDI) component for interactive workloads

Architecture diagram

b. Metrics component

Amazon CloudWatch receives metrics from the eVDI components.

c. Bastion Host component

Administrators may use SSH to connect to the bastion host component to manage the underlying infrastructure.

- 2. RES installs components in private subnets behind a NAT gateway. Administrators access the private subnets via the Application Load Balancer (ALB) or the Bastion Host component.
- 3. Amazon DynamoDB stores the environment configuration.
- 4. AWS Certificate Manager (ACM) generates and stores a public certificate for the Application Load Balancer (ALB).



Note

We recommend using AWS Certificate Manager to generate a trusted certificate for your domain.

- 5. Amazon Elastic File System (EFS) hosts the default /home file system mounted on all applicable infrastructure hosts and eVDI Linux sessions.
- 6. RES uses Amazon Cognito to create an initial bootstrap user called 'clusteradmin' within and sends temporary credentials to the email address provided during installation. The 'clusteradmin' must change the password the first time they login.
- 7. Amazon Cognito integrates with your organization's Active Directory and user identities for permissions management.
- 8. Security zones allow administrators to restrict access to specific components within the product based on permissions.

AWS services in this product

AWS service	Туре	Description
Amazon Elastic Compute Cloud	Core	Provides the underlying compute services to create virtual desktops with their

AWS service	Туре	Description
		chosen operating system and software stack.
Elastic Load Balancing	Core	Bastion, cluster-manager, and VDI hosts are created in Auto Scaling groups behind the load balancer. ELB balances traffic from the web portal across RES hosts.
Amazon Virtual Private Cloud	Core	All core product components are created within your VPC.
Amazon Cognito	Core	Manages user identities and authentication. Active Directory users are mapped to Amazon Cognito users and groups to authenticate access levels.
Amazon Elastic File System	Core	Provides the /home file system for the file browser and VDI hosts, as well as shared external file systems.
Amazon DynamoDB	Core	Stores configuration data such as users, groups, projects, file systems, and component settings.
AWS Systems Manager	Core	Stores documents for performing commands for VDI session management.

AWS service	Туре	Description
AWS Lambda	Core	Supports product functiona lities such as updating settings within the DynamoDB table, starting Active Directory sync workflows, and updating the prefix list.
Amazon CloudWatch	Supporting	Provides metrics and activity logs for all Amazon EC2 hosts and Lambda functions.
Amazon Simple Storage Service	Supporting	Stores application binaries for host bootstrapping and configuration.
AWS Key Management Service	Supporting	Used for encryption at rest with Amazon SQS queues, DynamoDB tables, and Amazon SNS topics.
AWS Secrets Manager	Supporting	Stores service account credentials in Active Directory and self-signed certificates for VDIs.
AWS CloudFormation	Supporting	Provides a deployment mechanism for the product.
AWS Identity and Access Management	Supporting	Restricts the access level for hosts.
Amazon Route 53	Supporting	Creates private hosted zone for resolving the internal load balancer and the bastion host domain name.

AWS service	Туре	Description
Amazon Simple Queue Service	Supporting	Creates task queues to support asynchronous executions.
Amazon Simple Notification Service	Supporting	Supports the publication- subscriber model between VDI components such as the controller and hosts.
AWS Fargate	Supporting	Installs, updates, and deletes environments using Fargate tasks.
Amazon FSx File Gateway	Optional	Provides external shared file system.
Amazon FSx for NetApp ONTAP	Optional	Provides external shared file system.
AWS Certificate Manager	Optional	Generates a trusted certificate for your custom domain.
AWS Backup	Optional	Offers backup capabilities for Amazon EC2 hosts, file systems, and DynamoDB.

Create a demo environment

Follow the steps in this section to try out Research and Engineering Studio on AWS. This demo deploys a non-production environment with a minimal set of parameters using the Research and Engineering Studio on AWS demo environment stack template. It uses a Keycloak server for SSO.

Note that after you deploy the stack, you must follow the <u>Post deployment steps</u> below to set up users in the environment before you login.

Create a one-click demo stack

This AWS CloudFormation stack creates all the components required by Research and Engineering Studio.

Time to deploy: ~90 minutes

Prerequisites

Topics

- Create an AWS account with an administrative user
- Create an Amazon EC2 SSH key pair
- Increase service quotas

Create an AWS account with an administrative user

You must have an AWS account with an administrative user:

- 1. Open https://portal.aws.amazon.com/billing/signup.
- 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.

Create a one-click demo stack

Create an Amazon EC2 SSH key pair

If you do not have Amazon EC2 SSH key pair, you will need to create one. For more information, see Create a key pair using Amazon EC2 in the Amazon EC2 User Guide.

Increase service quotas

We recommend increasing the service quotas for:

- Amazon VPC
 - Increase the Elastic IP address quota per NAT gateway from five to eight
 - Increase the NAT gateways per Availability Zone from five to ten
- Amazon EC2
 - Increase the EC2-VPC Elastic IPs from five to ten

Your AWS account has default quotas, formerly referred to as limits, for each AWS service. Unless otherwise noted, each quota is Region-specific. You can request increases for some quotas, and other quotas cannot be increased. For more information, see the section called "Quotas for AWS services in this product".

Create resources and input parameters

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



Note

Make sure you are in your administrator account.

- 2. Launch the template in the console.
- Under **Parameters**, review the parameters for this product template and modify them as necessary.

Parameter	Default	Description
EnvironmentName	<res-demo></res-demo>	A unique name given to your RES environment

Parameter	Default	Description
		starting with res-, no longer than 11 characters, and no capital letters.
AdministratorEmail		The email address for the user completing setup of the product. This user additionally functions as a break-glass user if there is an Active Directory single sign on integration failure.
KeyPair		The key pair used to connect to infrastructure hosts.
ClientIPCidr	<0.0.0/0>	IP address filter which limits connection to the system. You can update the ClientIpCidr after deployment.
InboundPrefixList		(Optional) Provide a managed prefix list for IPs allowed to directly access the web UI and SSH into the bastion host.

4. Choose Create stack.

Post deployment steps

- 1. Reset user passwords in AWS Directory Service—The demo stack creates four users with usernames which you can use: admin1, user1, admin2, and user2.
 - a. Go to the Directory Service console.

Post deployment steps 12

- b. Select the Directory Id for your environment. You can get the Directory Id from the output of <StackName>*DirectoryService* stack.
- c. From the top right **Action** dropdown menu, select **Reset user password**.
- d. For all the users you want to use, put the username and type in the password you want to have and choose **Reset Password**.
- 2. Once you have reset the user passwords, you will need to wait for Research and Engineering Studio to sync the users in the environment. Research and Engineering Studio syncs the users every hour at xx.00. You can either wait for that to happen or follow the steps listed in User added in Active Directory, but missing from RES to sync the users immediately.

Your deployment is now ready. Use the EnvironmentUrl you received in your email to access the UI, or you can also get the same URL from the output of the deployed stack. You may now login to the Research and Engineering Studio environment with the user and password that you reset the password for in Active Directory.

Post deployment steps 13

Plan your deployment

This section contains information on cost, security, supported regions, and quotas that can help you plan your deployment of Research and Engineering Studio on AWS.

Cost

Research and Engineering Studio on AWS is available at no additional charge, and you pay only for the AWS resources needed to run your applications. For more information, see AWS services in this product.



Note

You are responsible for the cost of the AWS services used while running this product. We recommend creating a budget through AWS Cost Explorer to help manage costs. Prices are subject to change. For full details, see the pricing webpage for each AWS service used in this product.

Security

Cloud security at AWS is the highest priority. As an AWS customer, 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. Thirdparty 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 Research and Engineering Studio on AWS, 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.

Cost

To understand how to apply the shared responsibility model with the AWS services used by Research and Engineering Studio, see <u>Security considerations for services in this product</u>. For more information about AWS security, visit AWS Cloud Security.

IAM roles

AWS Identity and Access Management (IAM) roles allow customers to assign granular access policies and permissions to services and users on the AWS Cloud. This product creates IAM roles that grant the product's AWS Lambda functions and Amazon EC2 instances access to create Regional resources.

RES supports identity-based policies within IAM. When deployed, RES creates policies to define the administrator permission and access. The administrator who implements the product creates and manages end users and project leaders within the existing customer Active Directory integrated with RES. For more information, see Creating IAM policies in the AWS Identity and Access Management User Guide.

Your organization's administrator can manage user access with an active directory. When end users access the RES user interface, RES authenticates with Amazon Cognito.

Security groups

The security groups created in this product are designed to control and isolate network traffic between the Lambda functions, EC2 instances, file systems CSR instances, and remote VPN endpoints. We recommend that you review the security groups and further restrict access as needed once the product is deployed.

Data encryption

By default, Research and Engineering Studio on AWS (RES) encrypts customer data at rest and in transit using an RES owned key. When you deploy RES, you may specify an AWS KMS key. RES uses your credentials to grant key access. If you supply a customer owned and managed AWS KMS key, customer data at rest will be encrypted using that key.

RES encrypts customer data in transit using SSL/TLS. We require TLS 1.2, but recommend TLS 1.3.

Security considerations for services in this product

For more detailed information regarding security considerations for the services used by Research and Engineering Studio, follow the links in this table:

IAM roles 15

AWS service security info	Service type	How the service is used in RES
Amazon Elastic Compute Cloud	Core	Provides the underlying compute services to create virtual desktops with their chosen operating system and software stack.
Elastic Load Balancing	Core	Bastion, cluster-manager, and VDI hosts are created in Auto Scaling groups behind the load balancer. ELB balances traffic from the web portal across RES hosts.
Amazon Virtual Private Cloud	Core	All core product components are created within your VPC.
Amazon Cognito	Core	Manages user identities and authentication. Active Directory users are mapped to Amazon Cognito users and groups to authenticate access levels.
Amazon Elastic File System	Core	Provides the /home file system for the file browser and VDI hosts, as well as shared external file systems.
Amazon DynamoDB	Core	Stores configuration data such as users, groups, projects, file systems, and component settings.

AWS service security info	Service type	How the service is used in RES
AWS Systems Manager	Core	Stores documents for performing commands for VDI session management.
AWS Lambda	Core	Supports product functiona lities such as updating settings within the DynamoDB table, starting Active Directory sync workflows, and updating the prefix list.
Amazon CloudWatch	Supporting	Provides metrics and activity logs for all Amazon EC2 hosts and Lambda functions.
Amazon Simple Storage Service	Supporting	Stores application binaries for host bootstrapping and configuration.
AWS Key Management Service	Supporting	Used for encryption at rest with Amazon SQS queues, DynamoDB tables, and Amazon SNS topics.
AWS Secrets Manager	Supporting	Stores service account credentials in Active Directory and self-signed certificates for VDIs.
AWS CloudFormation	Supporting	Provides a deployment mechanism for the product.
AWS Identity and Access Management	Supporting	Restricts the access level for hosts.

AWS service security info	Service type	How the service is used in RES
Amazon Route 53	Supporting	Creates private hosted zone for resolving the internal load balancer and the bastion host domain name.
Amazon Simple Queue Service	Supporting	Creates task queues to support asynchronous executions.
Amazon Simple Notification Service	Supporting	Supports the publication- subscriber model between VDI components such as the controller and hosts.
AWS Fargate	Supporting	Installs, updates, and deletes environments using Fargate tasks.
Amazon FSx File Gateway	Optional	Provides external shared file system.
Amazon FSx for NetApp ONTAP	Optional	Provides external shared file system.
AWS Certificate Manager	Optional	Generates a trusted certificate for your custom domain.
AWS Backup	Optional	Offers backup capabilities for Amazon EC2 hosts, file systems, and DynamoDB.

Quotas

Service quotas, also referred to as limits, are the maximum number of service resources or operations for your AWS account.

Quotas 18

Quotas for AWS services in this product

Make sure you have sufficient quota for each of the <u>services implemented in this product</u>. For more information, see AWS service quotas.

For this product, we recommend raising quotas for the following services:

- Amazon Virtual Private Cloud
- Amazon EC2

To request a quota increase, see <u>Requesting a Quota Increase</u> in the <u>Service Quotas User Guide</u>. If the quota is not yet available in Service Quotas, use the <u>limit increase</u> form.

AWS CloudFormation quotas

Your AWS account has AWS CloudFormation quotas that you should be aware of when <u>launching</u> <u>the stack</u> in this product. By understanding these quotas, you can avoid limitation errors that would prevent you from deploying this product successfully. For more information, see <u>AWS</u> <u>CloudFormation quotas</u> in the in the *AWS CloudFormation User's Guide*.

Planning for resilience

The product deploys a default infrastructure with the minimum number and size of Amazon EC2 instances to operate the system. To improve resilience in large-scale production environments, we recommend increasing the default minimum capacity settings within the infrastructure's Auto Scaling groups (ASG). Increasing the value from one instance to two instances provides the benefit of multiple Availability Zones (AZ) and reduces the time to restore system functionality in the event of unexpected data loss.

ASG settings can be customized within the Amazon EC2 console at https://console.aws.amazon.com/ec2/. The product creates four ASGs by default with each name ending with -asg. You can change the minimum and desired values to an amount appropriate for your production environment. Select the group you want to modify, and then choose **Actions** and select **Edit**. For more information on ASGs, see Scaling group in the Amazon EC2 Auto Scaling User Guide.

Supported AWS Regions

This product uses services which are not currently available in all AWS Regions. You must launch this product in an AWS Region where all services are available. For the most current availability of AWS services by Region, see the AWS Regional Services List.

Research and Engineering Studio on AWS is supported in the following AWS Regions:

Region name	Region	Previous versions	Latest version (2024.10)
US East (N. Virginia)	us-east-1	yes	yes
US East (Ohio)	us-east-2	yes	yes
US West (N. Californi a)	us-west-1	yes	yes
US West (Oregon)	us-west-2	yes	yes
Asia Pacific (Tokyo)	ap-northeast-1	yes	yes
Asia Pacific (Seoul)	ap-northeast-2	yes	yes
Asia Pacific (Mumbai)	ap-south-1	yes	yes
Asia Pacific (Singapor e)	ap-southeast-1	yes	yes
Asia Pacific (Sydney)	ap-southeast-2	yes	yes
Canada (Central)	ca-central-1	yes	yes
Europe (Frankfurt)	eu-central-1	yes	yes
Europe (Milan)	eu-south-1	yes	yes
Europe (Ireland)	eu-west-1	yes	yes
Europe (London)	eu-west-2	yes	yes

Supported AWS Regions 20

Region name	Region	Previous versions	Latest version (2024.10)
Europe (Paris)	eu-west-3	yes	yes
Europe (Stockholm)	eu-north-1	no	yes
Israel (Tel Aviv)	il-central-1	yes	yes
AWS GovCloud (US- West)	us-gov-west-1	yes	yes

Supported AWS Regions 21

Deploy the product



Note

This product uses AWS CloudFormation templates and stacks to automate its deployment. The CloudFormation templates describe the AWS resources included in this product and their properties. The CloudFormation stack provisions the resources that are described in the templates.

Before you launch the product, review the cost, architecture, network security, and other considerations discussed earlier in this guide.

Topics

- **Prerequisites**
- Create external resources
- Step 1: Launch the product
- Step 2: Sign in for the first time

Prerequisites

Topics

- Create an AWS account with an administrative user
- Create an Amazon EC2 SSH key pair
- Increase service quotas
- Create a public domain (optional)
- Create domain (GovCloud only)
- Provide external resources
- Configure LDAPS in your environment (optional)
- Configure a private VPC (optional)

Prerequisites 22

Create an AWS account with an administrative user

You must have an AWS account with an administrative user:

- 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.

Create an Amazon EC2 SSH key pair

If you do not have Amazon EC2 SSH key pair, you will need to create one. For more information, see Create a key pair using Amazon EC2 in the Amazon EC2 User Guide.

Increase service quotas

We recommend increasing the service quotas for:

- Amazon VPC
 - Increase the Elastic IP address quota per NAT gateway from five to eight.
 - Increase the NAT gateways per Availability Zone from five to ten.
- Amazon EC2
 - Increase the EC2-VPC Elastic IPs from five to ten

Your AWS account has default quotas, formerly referred to as limits, for each AWS service. Unless otherwise noted, each quota is Region-specific. You can request increases for some quotas, and other quotas cannot be increased. For more information, see Quotas for AWS services in this product.

Create a public domain (optional)

We recommend using a custom domain for the product in order to have a user-friendly URL. You will need to register a domain using Amazon Route 53 or another provider and import a certificate for the domain using AWS Certificate Manager. If you already have a public domain and certificate, you may skip this step.

- 1. Follow the directions to <u>register a domain</u> with Route53. You should receive a confirmation email.
- 2. Retrieve the hosted zone for your domain. This is created automatically by Route53.
 - a. Open the Route53 console.
 - b. Choose **Hosted zones** from the left navigation.
 - c. Open the hosted zone created for your domain name and copy the **Hosted zone ID**.
- 3. Open AWS Certificate Manager and follow these steps to <u>request a domain certificate</u>. Ensure you are in the Region where you plan to deploy the solution.
- 4. Choose **List certificates** from the navigation, and find your certificate request. The request should be pending.
- 5. Choose your **Certificate ID** to open the request.
- 6. From the **Domains** section, choose **Create records in Route53**. It will take approximately ten minutes for the request to process.
- 7. Once the certificate is issued, copy the **ARN** from the **Certificate status** section.

Create domain (GovCloud only)

If you are deploying in the AWS GovCloud (US-West) Region and you are using a custom domain for Research and Engineering Studio, you will need to complete these prerequisite steps.

- Deploy the <u>Certificate AWS CloudFormation stack</u> in the commercial-partition AWS Account where the public hosted domain was created.
- From the Certificate CloudFormation Outputs, find and note the CertificateARN and PrivateKeySecretARN.
- 3. In the GovCloud partition account, create a secret with the value of the CertificateARN output. Note the new secret ARN and add two tags to the secret so vdc-gateway can access the secret value:

- a. res:ModuleName = virtual-desktop-controller
- b. res:EnvironmentName = [environment name] (This could be res-demo.)
- 4. In the GovCloud partition account, create a secret with the value of the PrivateKeySecretArn output. Note the new secret ARN and add two tags to the secret so vdc-gateway can access the secret value:
 - a. res:ModuleName = virtual-desktop-controller
 - b. res:EnvironmentName = [environment name] (This could be res-demo.)

Provide external resources

Research and Engineering Studio on AWS expects the following external resources to exist when it is deployed.

Networking (VPC, Public Subnets, and Private Subnets)

This is where you will run the EC2 instances used to host the RES environment, the Active Directory (AD), and shared storage.

• Storage (Amazon EFS)

The storage volumes contain files and data needed for the virtual desktop infrastructure (VDI).

• Directory service (AWS Directory Service for Microsoft Active Directory)

The directory service authenticates users to the RES environment.

A secret that contains the service account password

Research and Engineering Studio accesses <u>secrets</u> that you provide, including the service account password, using AWS Secrets Manager.



If you are deploying a demo environment and do not have these external resources available, you can use AWS High Performance Compute recipes to generate the external resources. See the following section, Create external resources, to deploy resources in your account.

Provide external resources 25

For demo deployments in the AWS GovCloud (US-West) Region, you will need to complete the prerequisite steps in Create domain (GovCloud only).

Configure LDAPS in your environment (optional)

If you plan to use LDAPS communication in your environment, you must complete these steps to create and attach certificates to the AWS Managed Microsoft AD (AD) domain controller to provide communication between AD and RES.

- Follow the steps provided in <u>How to enable server-side LDAPS for your AWS Managed</u> <u>Microsoft AD</u>. You can skip this step if you have already enabled LDAPS.
- 2. After confirming that LDAPS is configured on the AD, export the AD certificate:
 - a. Go to your Active Directory server.
 - b. Open PowerShell as an administrator.
 - c. Run certmgr.msc to open the certificate list.
 - d. Open the certificate list by first opening the Trusted Root Certification Authorities and then Certificates.
 - e. Select and hold (or right-click) the certificate with the same name as your AD server and choose **All tasks** and then **Export**.
 - f. Select Base-64 encoded X.509 (.CER) and choose Next.
 - g. Select a directory and then choose **Next**.
- 3. Create a secret in AWS Secrets Manager:

When creating your Secret in the Secrets Manager, choose **Other type of secrets** under **secret type** and paste your PEM encoded certificate in the **Plaintext** field.

4. Note the ARN created and input it as the DomainTLSCertificateSecretARN parameter in Step 1: Launch the product.

Configure a private VPC (optional)

Deploying Research and Engineering Studio in an isolated VPC offers enhanced security to meet your organization's compliance and governance requirements. However, the standard RES deployment relies on internet access for installing dependencies. To install RES in a private VPC, you will need to satisfy the following prerequisites:

Topics

- Prepare Amazon Machine Images (AMIs)
- Set up VPC endpoints
- Connect to services without VPC endpoints
- Set private VPC deployment parameters

Prepare Amazon Machine Images (AMIs)

- Download <u>dependencies</u>. To deploy in an isolated VPC, the RES infrastructure requires the availability of dependencies without having public internet access.
- 2. Create an IAM role with Amazon S3 read-only access and trusted identity as Amazon EC2.
 - a. Open the IAM console at https://console.aws.amazon.com/iam/.
 - b. From Roles, choose Create role.
 - c. On the **Select trusted entity** page:
 - Under Trusted entity type, choose AWS service.
 - For Use case under Service or use case, choose EC2 and choose Next.
 - d. On Add permissions, select the following permission policies and then choose Next:
 - AmazonS3ReadOnlyAccess
 - AmazonSSMManagedInstanceCore
 - EC2InstanceProfileForImageBuilder
 - e. Add a **Role name** and **Description**, and then choose **Create role**.
- 3. Create the EC2 image builder component:
 - a. Open the EC2 Image Builder console at https://console.aws.amazon.com/imagebuilder.
 - b. Under **Saved resources**, choose **Components** and choose **Create component**.
 - c. On the **Create component** page, enter the following details:
 - For Component type, choose Build.
 - For Component details choose:

Parameter

User entry

Linux

Compatible OS Versions

Amazon Linux 2

Component name

Enter a name such as: <researchand-engineering-studio-inf
rastructure>

Component version

We recommend starting with 1.0.0.

Description

Optional user entry.

- d. On the **Create component** page, choose **Define document content**.
 - i. Before entering the definition document content, you will need a file URI for the tar.gz file. Upload the tar.gz file provided by RES to an Amazon S3 bucket and copy the file's URI from the bucket properties.
 - ii. Enter the following:

Note

AddEnvironmentVariables is optional, and you may remove it if you do not require custom environment variables in your infrastructure hosts. If you are setting up http_proxy and https_proxy environment variables, the no_proxy parameters are required to prevent the instance from using proxy to query localhost, instance metadata IP addresses, and the services that support VPC endpoints.

```
# Copyright Amazon.com, Inc. or its affiliates. All Rights Reserved.
#
# Licensed under the Apache License, Version 2.0 (the "License"). You may
not use this file except in compliance
# with the License. A copy of the License is located at
#
http://www.apache.org/licenses/LICENSE-2.0
```

```
# or in the 'license' file accompanying this file. This file is
distributed on an 'AS IS' BASIS, WITHOUT WARRANTIES
# OR CONDITIONS OF ANY KIND, express or implied. See the License for the
specific language governing permissions
# and limitations under the License.
name: research-and-engineering-studio-infrastructure
description: An RES EC2 Image Builder component to install required RES
 software dependencies for infrastructure hosts.
schemaVersion: 1.0
parameters:
  - AWSAccountID:
      type: string
      description: RES Environment AWS Account ID
  - AWSRegion:
      type: string
      description: RES Environment AWS Region
phases:
  - name: build
    steps:
       - name: DownloadRESInstallScripts
         action: S3Download
         onFailure: Abort
         maxAttempts: 3
         inputs:
            - source: '<s3 tar.gz file uri>'
              destination: '/root/bootstrap/res_dependencies/
res_dependencies.tar.gz'
              expectedBucketOwner: '{{ AWSAccountID }}'
       - name: RunInstallScript
         action: ExecuteBash
         onFailure: Abort
         maxAttempts: 3
         inputs:
            commands:
                - 'cd /root/bootstrap/res_dependencies'
                - 'tar -xf res_dependencies.tar.gz'
                'cd all_dependencies'
                - '/bin/bash install.sh'
       - name: AddEnvironmentVariables
         action: ExecuteBash
         onFailure: Abort
         maxAttempts: 3
```

```
inputs:
            commands:
                  echo -e "
                  http_proxy=http://<ip>:<port>
                  https_proxy=http://<ip>:<port>
no_proxy=127.0.0.1,169.254.169.254,169.254.170.2,localhost,
{{ AWSRegion }}.res,{{ AWSRegion }}.vpce.amazonaws.com,
{{ AWSRegion }}.elb.amazonaws.com,s3.
{{ AWSRegion }}.amazonaws.com,s3.dualstack.
{{ AWSRegion }}.amazonaws.com,ec2.{{ AWSRegion }}.amazonaws.com,ec2.
{{ AWSRegion }}.api.aws,ec2messages.{{ AWSRegion }}.amazonaws.com,ssm.
{{ AWSRegion }}.amazonaws.com,ssmmessages.
{{ AWSRegion }}.amazonaws.com,kms.
{{ AWSRegion }}.amazonaws.com, secretsmanager.
{{ AWSRegion }}.amazonaws.com,sqs.
{{ AWSRegion }}.amazonaws.com,elasticloadbalancing.
{{ AWSRegion }}.amazonaws.com,sns.{{ AWSRegion }}.amazonaws.com,logs.
{{ AWSRegion }}.amazonaws.com,logs.
{{ AWSRegion }}.api.aws,elasticfilesystem.
{{ AWSRegion }}.amazonaws.com,fsx.{{ AWSRegion }}.amazonaws.com,dynamodb.
{{ AWSRegion }}.amazonaws.com,api.ecr.
{{ AWSRegion }}.amazonaws.com,.dkr.ecr.
{{ AWSRegion }}.amazonaws.com,kinesis.{{ AWSRegion }}.amazonaws.com,.data-
kinesis.{{ AWSRegion }}.amazonaws.com,.control-
kinesis.{{ AWSRegion }}.amazonaws.com,events.
{{ AWSRegion }}.amazonaws.com,cloudformation.
{{ AWSRegion }}.amazonaws.com,sts.
{{ AWSRegion }}.amazonaws.com,application-autoscaling.
{{ AWSRegion }}.amazonaws.com,monitoring.{{ AWSRegion }}.amazonaws.com
                  " > /etc/environment
```

- e. Choose Create component.
- 4. Create an Image Builder image recipe.
 - a. On the **Create recipe** page, enter the following:

Section	Parameter	User entry
Recipe details	Name	Enter an appropriate name such as res-recipe-linux-x 86.
	Version	Enter a version, typically starting with 1.0.0.
	Description	Add an optional descripti on.
Base image	Select image	Select managed images.
	os	Amazon Linux
	Image origin	Quick start (Amazon-m anaged)
	Image name	Amazon Linux 2 x86
	Auto-versioning options	Use latest available OS version.
Instance configuration	_	Keep everything in the default settings, and make sure Remove SSM agent after pipeline execution is not selected.
Working directory	Working directory path	/root/bootstrap/re s_dependencies

Section	Parameter	User entry
Components	Build components	Search for and select the following:
		 Amazon-managed: aws- cli-version-2-linux
		 Amazon-managed: amazon-cloudwatch- agent-linux
		 Owned by you: Amazon EC2 component created previously. Put your AWS account ID and current AWS Region in the fields.
	Test components	Search for and select:
		 Amazon-managed: simple-boot-test-linux

- b. Choose **Create recipe**.
- 5. Create Image Builder infrastructure configuration.
 - a. Under **Saved resources**, choose **Infrastructure configurations**.
 - b. Choose **Create infrastructure configuration**.
 - c. On the **Create infrastructure configuration** page, enter the following:

Section	Parameter	User entry
General	Name	Enter an appropriate name such as res-infra-linux-x86.
	Description	Add an optional descripti
		on.

Section	Parameter	User entry
	IAM role	Select the IAM role created previously.
AWS infrastructure	Instance type	Choose t3.medium.
	VPC, subnet, and security groups	Select an option that permits internet access and access to the Amazon S3 bucket. If you need to create a security group, you can create one from the Amazon EC2 console with the following inputs: • VPC: Select the same VPC being used for the infrastructure configuration. This VPC must have internet access. • Inbound rule: • Type: SSH • Source: Custom • CIDR block: 0.0.0.0/0

- d. Choose **Create infrastructure configuration**.
- 6. Create a new EC2 Image Builder pipeline:
 - a. Go to Image pipelines, and choose Create image pipeline.
 - b. On the **Specify pipeline details** page, enter the following and choose **Next**:
 - Pipeline name and optional description
 - For **Build schedule**, set a schedule or choose **Manual** if you want to start the AMI baking process manually.
 - c. On the Choose recipe page, choose Use existing recipe and enter the Recipe name created previously. Choose Next.

- d. On the **Define image process** page, select the default workflows and choose **Next**.
- e. On the **Define infrastructure configuration** page, choose **Use existing infrastructure configuration** and enter the name of the previously created infrastructure configuration. Choose **Next**.
- f. On the **Define distribution settings** page, consider the following for your selections:
 - The output image must reside in the same region as the deployed RES environment, so that RES can properly launch infrastructure host instances from it. Using service defaults, the output image will be created in the region where the EC2 Image Builder service is being used.
 - If you want to deploy RES in multiple regions, you can choose **Create a new distribution settings** and add more regions there.
- g. Review your selections and choose **Create pipeline**.
- 7. Run the EC2 Image Builder pipeline:
 - a. From **Image pipelines**, find and select the pipeline you created.
 - b. Choose **Actions**, and select **Run pipeline**.

The pipeline may take approximately 45 minutes to an hour to create an AMI image.

8. Note the AMI ID for the generated AMI and use it as the input for the InfrastructureHostAMI parameter in the section called "Step 1: Launch the product".

Set up VPC endpoints

To deploy RES and launch virtual desktops, AWS services require access to your private subnet. You must set up VPC endpoints to provide the required access, and you will need to repeat these steps for each endpoint.

- 1. If endpoints have not previously been configured, follow the instructions provided in <u>Access</u> an AWS service using an interface VPC endpoint.
- 2. Select one private subnet in each of the two availability zones.

AWS service	Service name
Application Auto Scaling	com.amazonaws. <i>region</i> .application-autoscaling

AWS service	Service name
AWS CloudFormation	com.amazonaws. <i>region</i> .cloudformation
Amazon CloudWatch	com.amazonaws. <i>region</i> .monitoring
Amazon CloudWatch Logs	com.amazonaws. <i>region</i> .logs
Amazon DynamoDB	com.amazonaws. <i>region</i> .dynamodb (Requires gateway endpoint)
Amazon EC2	com.amazonaws. <i>region</i> .ec2
Amazon ECR	com.amazonaws. <i>region</i> .ecr.api
	com.amazonaws. <i>region</i> .ecr.dkr
Amazon Elastic File System	com.amazonaws. <i>region</i> .elasticfilesystem
Elastic Load Balancing	com.amazonaws. <i>region</i> .elasticloadbalancing
Amazon EventBridge	com.amazonaws. <i>region</i> .events
Amazon FSx	com.amazonaws. <i>region</i> .fsx
AWS Key Management Service	com.amazonaws. <i>region</i> .kms
Amazon Kinesis Data Streams	com.amazonaws. <i>region</i> .kinesis-streams
AWS Lambda	com.amazonaws. <i>region</i> .lambda
Amazon S3	com.amazonaws. <i>region</i> .s3 (Requires a gateway endpoint that is created by default in RES.)
	Additional Amazon S3 interface endpoints are required for cross-mounting buckets in an isolated environment. See Accessing Amazon Simple Storage Service interface endpoints .
AWS Secrets Manager	com.amazonaws. <i>region</i> .secretsmanager

AWS service	Service name
Amazon SES	com.amazonaws. <i>region</i> .email-smtp (Not supported in the following Availability Zones: use-1-az2, use1-az3, use1-az5, usw1-az2, usw2-az4, apne2-az4, cac1-az3, and cac1-az4.)
AWS Security Token Service	com.amazonaws. <i>region</i> .sts
Amazon SNS	com.amazonaws. <i>region</i> .sns
Amazon SQS	com.amazonaws. <i>region</i> .sqs
AWS Systems Manager	com.amazonaws. <i>region</i> .ec2messages
	com.amazonaws. <i>region</i> .ssm
	com.amazonaws. <i>region</i> .ssmmessages

Connect to services without VPC endpoints

To integrate with services that do not support VPC endpoints, you can set up a proxy server in a public subnet of your VPC. Follow these steps to create a proxy server with the minimum necessary access for a Research and Engineering Studio deployment using AWS Identity Center as your identity provider.

- 1. Launch a Linux instance in the public subnet of the VPC you will use for your RES deployment.
 - Linux family Amazon Linux 2 or Amazon Linux 3
 - Architecture x86
 - Instance type t2.micro or higher
 - Security group TCP on port 3128 from 0.0.0.0/0
- 2. Connect to the instance to set up a proxy server.
 - a. Open the http connection.
 - b. Allow connection to the following domains from all relevant subnets:
 - .amazonaws.com (for generic AWS services)

- .amazoncognito.com (for Amazon Cognito)
- .awsapps.com (for Identity Center)
- .signin.aws (for Identity Center)
- .amazonaws-us-gov.com (for Gov Cloud)
- c. Deny all other connections.
- d. Activate and start the proxy server.
- e. Note the PORT on which the proxy server listens.
- 3. Configure your route table to allow access to the proxy server.
 - Go to your VPC console and identify the route tables for the subnets you will be using for Infrastructure Hosts and VDI hosts.
 - Edit route table to allow all incoming connections to go to the proxy server instance created in the previous steps.
 - c. Do this for route tables for all the subnets (without internet access) which you are going to use for Infrastructure/VDIs.
- 4. Modify the security group of the proxy server EC2 instance and make sure it allows inbound TCP connections on the PORT on which the proxy server is listening.

Set private VPC deployment parameters

In <u>the section called "Step 1: Launch the product"</u>, you are expected to input certain parameters in the AWS CloudFormation template. Be sure to set the following parameters as noted to successfully deploy into the private VPC you just configured.

Parameter	Input
InfrastructureHostAMI	Use the infrastructure AMI ID created in the section called "Prepare Amazon Machine Images (AMIs)".
IsLoadBalancerInternetFacing	Set to false.
LoadBalancerSubnets	Choose private subnets without internet access.

Parameter	Input
InfrastructureHostSubnets	Choose private subnets without internet access.
VdiSubnets	Choose private subnets without internet access.
ClientIP	You can choose your VPC CIDR to allow access for all VPC IP addresses.

This CloudFormation stack creates networking, storage, active directory, and domain certificates (if a PortalDomainName is provided). You must have these external resources available to deploy the product.

You may download the recipes template before deployment.

Time to deploy: Approximately 40-90 minutes

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



Note

Make sure you are in your administrator account.

2. Launch the template in the console.

> If you are deploying in the AWS GovCloud (US-West) Region, launch the template in the GovCloud partition account.

Enter the template parameters:

Parameter	Default	Description
DomainName	corp.res.com	Domain used for the active directory. The default value

Parameter	Default	Description
		is supplied in the LDIF file which sets up bootstrap users. If you would like to use the default users, leave the value as default. To change the value, update and provide a separate LDIF file. This does not need to match the domain used for active directory.
SubDomain (GovCloud only)		This parameter is optional for commercial regions, but required for GovCloud regions.
		If you provide a SubDomain, the parameter will be prefixed to the DomainNam e provided. The provided Active Directory domain name will become a subdomain.

Parameter	Default	Description
AdminPassword		The password for the active directory administrator (username Admin). This user is created in the active directory for the initial bootstrapping phase and is not used after. Important: the format of this field can either be (1) a plain text password or (2) the ARN of an AWS Secret formatted as a key/value pair {"password"} Note: The password for this user must meet the password complexity requirements for active directory
		<u>directory</u> .

Parameter	Default	Description
ServiceAccountPassword		Password used to create a service account (ReadOnlyU ser). This account is used for synchronization. Important: the format of this field can either be (1) a plain text password or (2) the ARN of an AWS Secret formatted as a key/ value pair {"password"} . Note: The password for this user must meet the password complexity requirements for active directory.
Keypair		Connects the administr ative instances using an SSH client. Note: AWS Systems Manager Session Manager can also be used to connect to instances.

Parameter	Default	Description
LDIFS3Path	<pre>aws-hpc-recipes/ma in/recipes/res/res _demo_env/assets/r es.ldif</pre>	The Amazon S3 path to an LDIF file imported during the bootstrapping phase of active directory setup. For more information, see LDIF Support. The parameter prepopulates with a file that creates a number of users in the active directory. To view the file, see the res.ldif file available in GitHub.
ClientIpCidr		The IP address from which you will access the site. For example, you can select your IP address and use [IPADDRESS]/32 to only allow access from your host. You can update this postdeployment.
ClientPrefixList		Enter a prefix list to provide access to the active directory management nodes. For information on creating a managed prefix list, see Work with customer-managed prefix lists.

Parameter	Default	Description
EnvironmentName	res-[environment name]	If the PortalDomainName is provided, this parameter is used to add tags to the secrets generated so that they can be used within the environment. This will need to match the Environme ntName parameter used when creating the RES stack. If you are deploying multiple environments in your account, this will need to be unique.
PortalDomainName		For GovCloud deploymen ts, do not enter this parameter. The certifica tes and secrets were manually created during the prerequisites. The domain name in Amazon Route 53 for the account. If this is provided, then a public certificate and key file will be generated and uploaded to AWS Secrets Manager. If you have your own domain and certificates, this parameter and EnvironmentName can be left blank.

4. Acknowledge all checkboxes in **Capabilities**, and choose **Create stack**.

Follow the step-by-step instructions in this section to configure and deploy the product into your account.

Time to deploy: Approximately 60 minutes

You can download the CloudFormation template for this product before deploying it.

If you are deploying in AWS GovCloud (US-West), use this template.

res-stack - Use this template to launch the product and all associated components. The default configuration deploys the RES main stack and authentication, frontend, and backend resources.



Note

AWS CloudFormation resources are created from AWS Cloud Development Kit (AWS CDK) (AWS CDK) constructs.

The AWS CloudFormation template deploys Research and Engineering Studio on AWS in the AWS Cloud. You must meet the prerequisites before launching the stack.

- Sign in to the AWS Management Console and open the AWS CloudFormation console at 1. https://console.aws.amazon.com/cloudformation.
- 2. Launch the template.

To deploy in AWS GovCloud (US-West), launch this template.

3. The template launches in the US East (N. Virginia) Region by default. To launch the solution in a different AWS Region, use the Region selector in the console navigation bar.



Note

This product uses the Amazon Cognito service, which is not currently available in all AWS Regions. You must launch this product in an AWS Region where Amazon Cognito is available. For the most current availability by Region, see the AWS Regional Services List.

4. Under **Parameters**, review the parameters for this product template and modify them as necessary. If you deployed the automated external resources, you can find these parameters in the **Outputs** tab of the external resources stack.

Parameter	Default	Description
EnvironmentName	<res-demo></res-demo>	A unique name given to your RES environment starting with res-, no longer than 11 characters, and no capital letters.
AdministratorEmail		The email address for the user completing setup of the product. This user additionally functions as a break-glass user if there is an active directory single sign on integration failure.
InfrastructureHostAMI	ami-[numbers or letters only]	(Optional) You may provide a custom AMI id to use for all the infrastructure hosts. The current supported base OS is Amazon Linux 2. For more information, see Configure RES-ready AMIs.
SSHKeyPair		The key pair used to connect to infrastructure hosts.
ClientIP	x.x.x.0/24 or x.x.x.0/32	IP address filter which limits connection to the system. You can update the ClientIpCidr after deployment.

Parameter	Default	Description
ClientPrefixList		(Optional) Provide a managed prefix list for IPs allowed to directly access the web UI and SSH into the bastion host.
IAMPermissionBoundary		(Optional) You may provide a managed policy ARN that will be attached as a permission boundary to all roles created in RES. For more information, see Setting custom permission boundaries.
Vpcld		ID for the VPC where instances will launch.
IsLoadBalancerInternetFacin g		Select true to deploy internet facing load balancer (Requires public subnets for load balancer). For deployments that need restricted internet access, select false.

Parameter	Default	Description
LoadBalancerSubnets		Select at least two subnets in different Availability Zones where load balancers will launch. For deployments that need restricted internet access, select private subnets. For deployments that need internet access, select public subnets. If more than two were created by the external networkin g stack, select all that were created.
InfrastructureHostSubnets		Select at least two private subnets in different Availabil ity Zones where infrastru cture hosts will launch. If more than two were created by the external networkin g stack, select all that were created.
VdiSubnets		Select at least two private subnets in different Availabil ity Zones where VDI instances will launch. If more than two were created by the external networkin g stack, select all that were created.

Parameter	Default	Description
ActiveDirectoryName	corp.res.com	Domain for the active directory. It does not need to match the portal domain name.
ADShortName	corp	The short name for the active directory. This is also called the NetBIOS name.
LDAP Base	DC=corp,DC=res,DC= com	An LDAP path to the base within the LDAP hierarchy.
LDAPConnectionURI		A single ldap:// path that can be reached by the active directory's host server. If you deployed the automated external resources with the default AD domain, you can use ldap://corp.res.com.
ServiceAccountCred entialsSecretArn		Provide a Secret ARN which contains the username and password for the Active Directory ServiceAc count user, formatted as a username:password key/value pair.
UsersOU		Organizational unit within AD for users that will sync.
GroupsOU		Organizational unit within AD for groups that will sync.

Parameter	Default	Description
SudoersGroupName	RESAdministrators	Group name that contains all users with sudoer access on instances at install and administrator access on RES.
ComputersOU		Organizational unit within AD that instances will join.
DomainTLSCertifica teSecretARN		(Optional) Provide a domain TLS certificate secret ARN to enable TLS communication to AD.
EnableLdapIDMapping		Determines if UID and GID numbers are generated by SSSD or if the numbers provided by the AD are used. Set to True to use SSSD generated UID and GID, or False to use UID and GID provided by the AD. For most cases this parameter should be set to True.
DisableADJoin	False	To prevent Linux hosts from joining the directory domain, change to True. Otherwise, leave in the default setting of False.
ServiceAccountUserDN		Provide the distinguished name (DN) of the service account user in Directory.

Parameter	Default	Description
SharedHomeFilesystemID		An EFS ID to use for the shared home filesystem for Linux VDI hosts.
CustomDomainNamefo rWebApp		(Optional) Subdomain used by the web portal to provide links for the web portion of the system.
CustomDomainNameforVDI		(Optional) Subdomain used by the web portal to provide links for the VDI portion of the system.
ACMCertificateARNf orWebApp		(Optional) When using the default configuration, the product hosts the web application under the domain amazonaws .com. You may host the product services under your domain. If you deployed the automated external resources, this was generated for you and the information can be found in the Outputs of the resbi stack. If you need to generate a certificate for your web application, see Configuration guide.

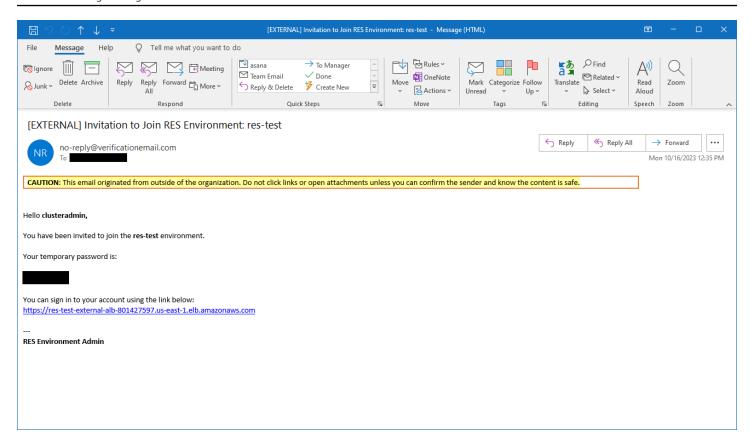
Parameter	Default	Description
CertificateSecretARNforVDI		(Optional) This ARN secret stores the public certifica te for your web portal's public certificate. If you set a portal domain name for your automated external resources, you can find this value under the Outputs tab of the res-bi stack.
PrivateKeySecretARNforVDI		(Optional) This ARN secret stores the private key for your web portal's certificate. If you set a portal domain name for your automated external resources, you can find this value under the Outputs tab of the res-bi stack.

5. Choose **Create stack** to deploy the stack.

You can view the status of the stack in the AWS CloudFormation console in the **Status** column. You should receive a CREATE_COMPLETE status in approximately 60 minutes.

Step 2: Sign in for the first time

Once the product stack has deployed in your account, you will receive an email with your credentials. Use the URL to sign in to your account and configure the workspace for other users.



Once you have signed in for the first time, you can configure settings in the web portal to connect to the SSO provider. For post-deployment configuration information, see the <u>Configuration guide</u>. Note that clusteradmin is a break-glass account— you can use it to create projects and assign user or group membership to those projects; it cannot assign software stacks or deploy a desktop for itself.

Update the product

Research and Engineering Studio (RES) has two methods of updating the product which depend on if the version update is major or minor.

RES uses a date-based versioning scheme. A major release uses the year and month, and a minor release adds a sequence number when necessary. For example, version 2024.01 was released in January 2024 as a major release; version 2024.01.01 was a minor release update of that version.

Topics

- Major version updates
- Minor version updates

Major version updates

Research and Engineering Studio uses snapshots to support migration from a previous RES environment to the latest without losing your environment settings. You can also use this process to test and verify updates to your environment before onboarding users.

To update your environment with the latest version of RES:

- 1. Create a snapshot of your current environment. See the section called "Create a snapshot".
- 2. Redeploy RES with the new version. See the section called "Step 1: Launch the product".
- 3. Apply the snapshot to your updated environment. See the section called "Apply a snapshot".
- 4. Verify all data migrated successfully to the new environment.

Minor version updates

For minor version updates to RES, a new install is not required. You can update the existing RES stack by updating its AWS CloudFormation template. Check the version of your current RES environment in AWS CloudFormation before deploying the update. You can find the version number at the beginning of the template.

For example: "Description": "RES_2024.1"

Major version updates 53

To make a minor version update:

- 1. Download the latest AWS CloudFormation template in the section called "Step 1: Launch the product".
- 2. Open the AWS CloudFormation console at https://console.aws.amazon.com/cloudformation.
- 3. From **Stacks**, find and select the primary stack. It should appear as <stack-name>.
- 4. Choose **Update**.
- 5. Choose Replace current template.
- 6. For **Template source**, choose **Upload a template file**.
- 7. Choose **Choose file** and upload the template you downloaded.
- 8. On **Specify stack details**, choose **Next**. You do not need to update the parameters.
- 9. On **Configure stack options**, choose **Next**.
- 10. On Review <stack-name>, choose Submit.

Minor version updates 54

Uninstall the product

You can uninstall the Research and Engineering Studio on AWS product from the AWS Management Console or by using the AWS Command Line Interface. You must manually delete the Amazon Simple Storage Service (Amazon S3) buckets created by this product. This product does not automatically delete <EnvironmentName>-shared-storage-security-group in case you have stored data to retain.

Using the AWS Management Console

- Sign in to the AWS CloudFormation console.
- 2. On the **Stacks** page, select this product's installation stack.
- 3. Choose **Delete**.

Using AWS Command Line Interface

Determine whether the AWS Command Line Interface (AWS CLI) is available in your environment. For installation instructions, see What Is the AWS Command Line Interface in the AWS CLI User Guide. After confirming that the AWS CLI is available and configured to the administrator account in the Region where the product was deployed, run the following command.

\$ aws cloudformation delete-stack --stack-name <RES-stack-name>

Deleting the shared-storage-security-group



Marning

The product retains this file system by default to protect against unintentional data loss. If you choose to delete the security group and associated file systems, any data retained within those systems will be permanently deleted. We recommend backing up data or reassigning the data to a new security group.

Sign in to the AWS Management Console and open the Amazon EFS console at https:// console.aws.amazon.com/efs/.

- Delete all file systems associated with <<u>RES-stack-name</u>>-shared-storage-securitygroup. Alternatively, you may reassign these file systems to another security group to maintain the data.
- 3. Sign in to the AWS Management Console and open the Amazon EC2 console at https://console.aws.amazon.com/ec2/.
- 4. Delete the <RES-stack-name>-shared-storage-security-group.

Deleting the Amazon S3 buckets

This product is configured to retain the product-created Amazon S3 bucket (for deploying in an opt-in Region) if you decide to delete the AWS CloudFormation stack to prevent accidental data loss. After uninstalling the product, you can manually delete this S3 bucket if you do not need to retain the data. Follow these steps to delete the Amazon S3 bucket.

- 1. Sign in to the AWS Management Console and open the Amazon S3 console at https://console.aws.amazon.com/s3/.
- 2. Choose **Buckets** from the navigation pane.
- 3. Locate the stack-name S3 buckets.
- 4. Select each Amazon S3 bucket, then choose **Empty**. You must empty each bucket.
- 5. Select the S3 bucket and choose **Delete**.

To delete S3 buckets using AWS CLI, run the following command:

\$ aws s3 rb s3://<bucket-name> --force



The --force command empties the bucket of its contents.

Configuration guide

This configuration guide provides post-deployment instructions for a technical audience on how to further customize and integrate with the Research and Engineering Studio on AWS product.

Topics

- Managing users and groups
- Creating subdomains
- Create an ACM certificate
- Amazon CloudWatch Logs
- Setting custom permission boundaries
- Configure RES-ready AMIs

Managing users and groups

Research and Engineering Studio can use any SAML 2.0 compliant identity provider. If you deployed RES using the external resources or plan to use IAM Identity center, see <u>Setting up single sign-on (SSO) with IAM Identity Center</u>. If you have your own SAML 2.0 compliant identity provider, see <u>Configuring your identity provider for single sign-on (SSO)</u>.

Topics

- Setting up single sign-on (SSO) with IAM Identity Center
- Configuring your identity provider for single sign-on (SSO)
- Setting passwords for users

Setting up single sign-on (SSO) with IAM Identity Center

If you do not already have an identity center connected to the managed Active Directory, start with <u>Step 1: Set up an identity center</u>. If you already have an identity center connected with the managed Active Directory, start with <u>Step 2: Connect to an identity center</u>.

Managing users and groups 57



Note

If you are deploying to the AWS GovCloud (US-West) Region, set up SSO in the AWS GovCloud (US) partition account where you deployed Research and Engineering Studio.

Step 1: Set up an identity center

Enabling IAM Identity Center

- Sign in to the AWS Identity and Access Management console.
- 2. Open the **Identity Center**.
- Choose **Enable**.
- 4. Choose **Enable with AWS Organizations**.
- 5. Choose **Continue**.



Note

Make sure you are in the same Region where you have your managed Active Directory.

Connecting IAM Identity Center to a managed Active Directory

After you enable IAM Identity Center, complete these recommended set up steps:

- In the navigation pane, choose **Settings**. 1.
- Under Identity source, choose Actions and choose Change identity source. 2.
- 3. Under **Existing directories**, select your directory.
- Choose Next. 4.
- Review your changes and enter ACCEPT in the confirmation box. 5.
- Choose Change identity source. 6.

Syncing users and groups to identity center

Once the changes made in Connecting IAM Identity Center to a managed Active Directory are complete, a green confirmation banner appears.

- 1. In the confirmation banner, choose **Start guided setup**.
- 2. From Configure attribute mappings, choose Next.
- 3. Under the **User** section, enter the users you want to sync.
- 4. Choose **Add**.
- Choose Next.
- 6. Review your changes, then choose **Save configuration**.
- 7. The sync process may take a few minutes. If you receive a warning message about users not syncing, choose **Resume sync**.

Enabling users

- 1. From the menu, choose **Users**.
- 2. Select the user(s) for whom you want to enable access.
- Choose Enable user access.

Step 2: Connect to an identity center

Setting up the application in IAM Identity Center

- 1. Open the IAM Identity Center console.
- 2. Choose **Applications**.
- 3. Choose Add application.
- 4. Under **Setup preference**, choose **I have an application I want to set up**.
- 5. Under Application type, choose SAML 2.0.
- 6. Choose **Next**.
- 7. Enter the display name and description you would like to use.
- 8. Under IAM Identity Center metadata, copy the link for the IAM Identity Center SAML metadata file. You will need this when configuring IAM Identity Center with the RES portal.
- Under Application properties, enter your Application start URL. For example, <yourportal-domain>/sso.
- 10. Under **Application ACS URL**, enter the redirect URL from the RES portal. To find this:
 - a. Under Environment management, choose General settings.

- b. Select the **Identity provider** tab.
- c. Under Single Sign-On, you will find the SAML Redirect URL.
- 11. Under Application SAML audience, enter the Amazon Cognito URN.

To create the urn:

- a. From the RES portal, open General Settings.
- b. Under the **Identity provider** tab, locate the **User Pool ID**.
- c. Add the **User Pool ID** to this string:

```
urn:amazon:cognito:sp:<user_pool_id>
```

12. After you enter the Amazon Cognito URN, choose **Submit**.

Configuring attribute mappings for the application

- 1. From the **Identity Center**, open the details for your created application.
- 2. Choose **Actions**, then choose **Edit attribute mappings**.
- Under Subject, enter \${user:email}.
- 4. Under Format, choose emailAddress.
- 5. Choose **Add new attribute mapping**.
- 6. Under **User attribute in the application**, enter 'email'.
- 7. Under Maps to this string value or user attribute in IAM Identity Center, enter \${user:email}.
- 8. Under **Format**, enter 'unspecified'.
- 9. Choose **Save changes**.

Adding users to the application in IAM Identity Center

- From the Identity Center, open Assigned users for your created application and choose Assign users.
- 2. Select the users you want to assign application access.
- 3. Choose Assign users.

Setting up IAM Identity Center within the RES environment

- From the Research and Engineering Studio environment, under **Environment management**, open **General settings**.
- Open the **Identity provider** tab. 2.
- Under Single Sign-On, choose Edit (next to Status). 3.
- 4. Complete the form with the following information:
 - Choose SAML. a.
 - b. Under **Provider name**, enter a user friendly name.
 - Choose Enter metadata document endpoint URL. c.
 - Enter the URL you copied during Setting up the application in IAM Identity Center. d.
 - e. Under Provider email attribute, enter 'email'.
 - f. Choose Submit.
- Refresh the page and check that the **Status** displays as enabled.

Configuring your identity provider for single sign-on (SSO)

Research and Engineering Studio integrates with any SAML 2.0 identity provider to authenticate user access to the RES portal. These steps provide directions to integrate with your chosen SAML 2.0 identity provider. If you intend to use IAM Identity Center, please see Setting up single sign-on (SSO) with IAM Identity Center.



Note

The user's email must match in the IDP SAML assertion and Active Directory. You will need to connect your identity provider with your Active Directory and periodically sync users.

Topics

- Configure your identity provider
- Configure RES to use your identity provider
- Configuring your identity provider in a non-production environment
- Debugging SAML IdP issues

Configure your identity provider

This section provides the steps to configure your identity provider with information from the RES Amazon Cognito user pool.

- 1. RES assumes that you have an AD (AWS Managed AD or a self-provisioned AD) with the user identities allowed to access the RES portal and projects. Connect your AD to your identity service provider and sync the user identities. Check your identity provider's documentation to learn how to connect your AD and sync user identities. For example, see Using Active Directory as an identity source in the AWS IAM Identity Center User Guide.
- 2. Configure a SAML 2.0 application for RES in your identity provider (IdP). This configuration requires the following parameters:
 - **SAML Redirect URL** The URL that your IdP uses to send the SAML 2.0 response to the service provider.



Depending on the IdP, the SAML Redirect URL might have a different name:

- Application URL
- Assertion Consumer Service (ACS) URL
- ACS POST Binding URL

To get the URL

- 1. Sign in to RES as an **admin** or **clusteradmin**.
- 2. Navigate to **Environment Management** ⇒ **General Settings** ⇒ **Identity Provider**.
- 3. Choose **SAML Redirect URL**.
- **SAML Audience URI** The unique ID of the SAML audience entity on the service provider side.



Note

Depending on the IdP, the SAML Audience URI might have a different name:

- ClientID
- Application SAML Audience
- SP entity ID

Provide the input in the following format.

```
urn:amazon:cognito:sp:user-pool-id
```

To find your SAML Audience URI

- 1. Sign in to RES as an **admin** or **clusteradmin**.
- 2. Navigate to **Environment Management** ⇒ **General Settings** ⇒ **Identity Provider**.
- 3. Choose User Pool Id.
- 3. The SAML assertion posted to RES must have the following fields/claims set to the user's email address:
 - SAML Subject or NameID
 - SAML email
- 4. Your IdP adds fields/claims to the SAML assertion, based on the configuration. RES requires these fields. Most providers automatically fill these fields by default. Refer to the following field inputs and values if you have to configure them.
 - AudienceRestriction Set to urn: amazon: cognito:sp:user-pool-id. Replace user-pool-id with the ID of your Amazon Cognito user pool.

```
<saml:AudienceRestriction>
     <saml:Audience> urn:amazon:cognito:sp:user-pool-id
</saml:AudienceRestriction>
```

Response — Set InResponseTo to https://user-pool-domain/saml2/idpresponse.
 Replace user-pool-domain with the domain name of your Amazon Cognito user pool.

```
<saml2p:Response
Destination="http://user-pool-domain/saml2/idpresponse"
ID="id123"
InResponseTo="_dd0a3436-bc64-4679-a0c2-cb4454f04184"
IssueInstant="Date-time stamp"</pre>
```

```
Version="2.0"
xmlns:saml2p="urn:oasis:names:tc:SAML:2.0:protocol"
xmlns:xs="http://www.w3.org/2001/XMLSchema">
```

• **SubjectConfirmationData** — Set Recipient to your user pool sam12/idpresponse endpoint and InResponseTo to the original SAML request ID.

```
<saml2:SubjectConfirmationData
InResponseTo="_dd0a3436-bc64-4679-a0c2-cb4454f04184"
NotOnOrAfter="Date-time stamp"
Recipient="https://user-pool-domain/saml2/idpresponse"/>
```

AuthnStatement — Configure as the following:

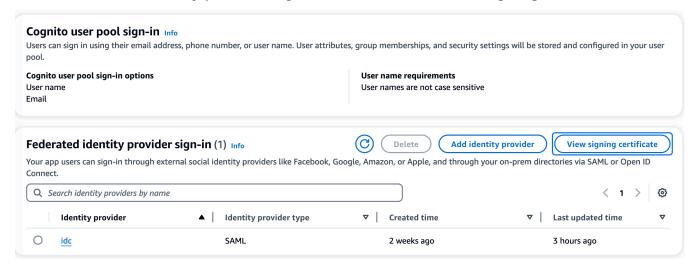
5. If your SAML application has a logout URL field, set it to: <domain-url>/saml2/logout.

To get the domain URL

- 1. Sign in to RES as an admin or clusteradmin.
- 2. Navigate to Environment Management ⇒ General Settings ⇒ Identity Provider.
- 3. Choose **Domain URL**.
- 6. If your IdP accepts a signing certificate to establish trust with Amazon Cognito, download the Amazon Cognito signing certificate and upload it in your IdP.

To get the signing certificate

- Open the Amazon Cognito console in the <u>Getting Started with the AWS Management</u> Console
- 2. Select your user pool. Your user pool should be res-<environment name>-user-pool.
- 3. Select the **Sign-in experience** tab.
- 4. In the Federated identity provider sign-in section, choose View signing certificate.



You can use this certificate to set up Active Directory IDP, add a relying party trust, and enable SAML support on this relying party.

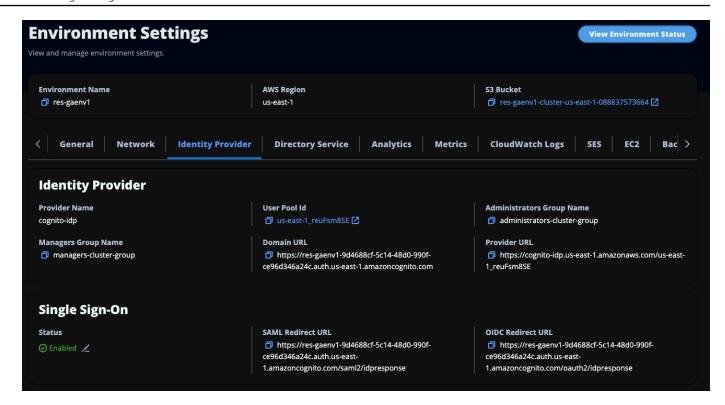


5. After the application setup is complete, download the SAML 2.0 application metadata XML or URL. You use it in the next section.

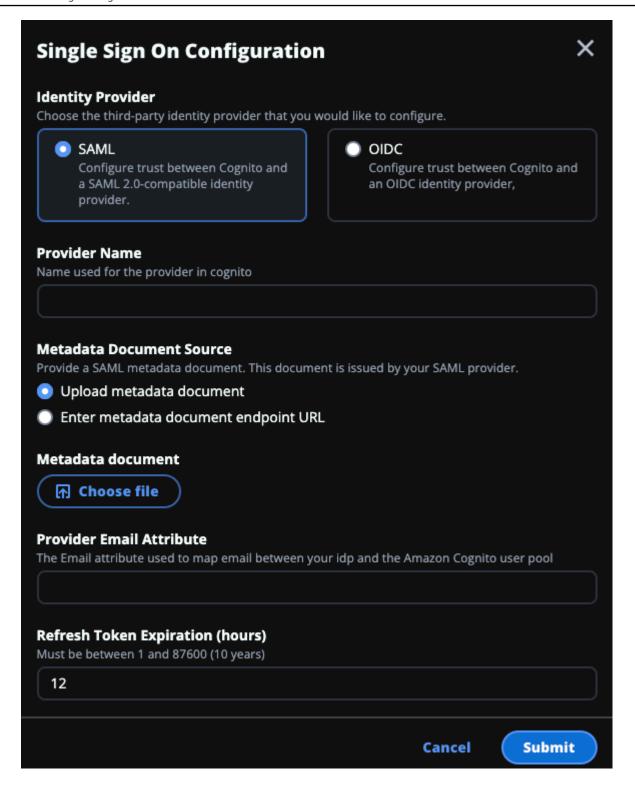
Configure RES to use your identity provider

To complete the single sign-on setup for RES

- 1. Sign in to RES as an admin or clusteradmin.
- 2. Navigate to Environment Management ⇒ General Settings ⇒ Identity Provider.



3. Under **Single Sign-On**, choose the edit icon next to the status indicator to open the **Single Sign On Configuration** page.



- a. For **Identity Provider**, choose **SAML**.
- b. For **Provider Name**, enter a unique name for your identity provider.



Note

The following names are not allowed:

- Cognito
- IdentityCenter
- Under Metadata Document Source, choose the appropriate option and upload the metadata XML document or provide the URL from the identity provider.
- For **Provider Email Attribute**, enter the text value email.
- Choose **Submit**. e.
- Reload the **Environment Settings** page. Single sign-on is enabled if the configuration was correct.

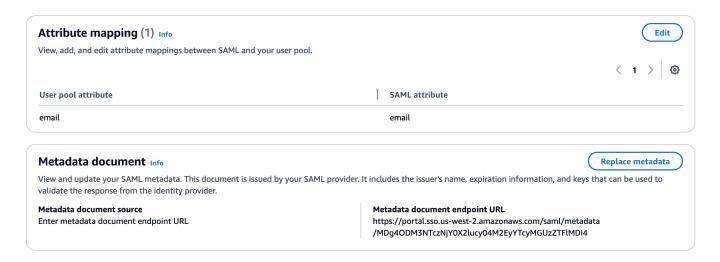
Configuring your identity provider in a non-production environment

If you used the provided external resources to create a non-production RES environment and configured IAM Identity Center as your identity provider, you may want to configure a different identity provider such as Okta. The RES SSO enablement form asks for three configuration parameters:

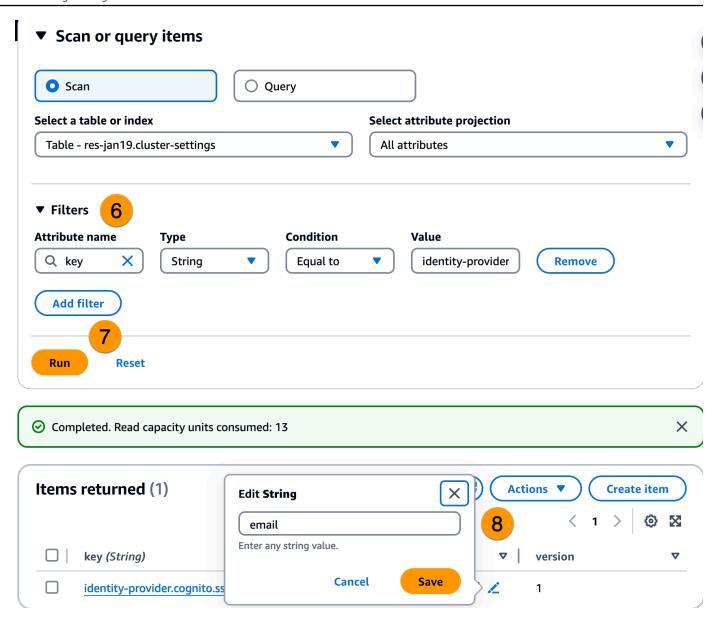
- 1. Provider name Cannot be modified
- 2. Metadata document or URL Can be modified
- 3. Provider email attribute Can be modified

To modify the metadata document and provider email attribute, do the following:

- 1. Go to the Amazon Cognito console.
- 2. From the navigation, choose **User pools**.
- 3. Select your user pool to view the **User pool overview**.
- From the **Sign-in experience** tab, go to **Federated identity provider sign-in** and open your configured identity provider.
- Generally, you will only be required to change the metadata and leave the attribute mapping unchanged. To update Attribute mapping, choose Edit. To update the Metadata document, choose Replace metadata.



- 6. If you edited the attribute mapping, you will need to update the <environment name>.cluster-settings table in DynamoDB.
 - a. Open the DynamoDB console and choose **Tables** from the navigation.
 - b. Find and select the <environment name>.cluster-settings table, and from the **Actions** menu select **Explore items**.
 - c. Under **Scan or query items**, go to **Filters** and enter the following parameters:
 - Attribute name key
 - Value identity-provider.cognito.sso_idp_provider_email_attribute
 - d. Choose Run.
- 7. Under **Items returned**, find the identity-provider.cognito.sso_idp_provider_email_attribute string and choose **Edit** to modify the string to match your changes in Amazon Cognito.



Debugging SAML IdP issues

SAML-tracer — You can use this extension for the Chrome browser to track SAML requests and check the SAML assertion values. For more information, see <u>SAML-tracer</u> at the Chrome web store.

SAML developer tools — OneLogin provides tools that you can use to decode the SAML encoded value and check the required fields in the SAML assertion. For more information, see Base 64
Decode + Inflate
at the OneLogin web site.

Amazon CloudWatch Logs — You can check your RES logs in CloudWatch Logs for errors or warnings. Your logs are in a log group with the name format res-environment-name/clustermanager.

Amazon Cognito documentation — For more information about SAML integration with Amazon Cognito, see Adding SAML identity providers to a user pool in the Amazon Cognito Developer Guide.

Setting passwords for users

- 1. From the AWS Directory Service console, select the directory for the created stack.
- 2. Under the **Actions** menu, select **Reset user password**.
- 3. Select the user and enter a new password.
- 4. Choose **Reset password**.

Creating subdomains

If you are using a custom domain, you will need to set up subdomains to support the web and VDI portions of your portal.

Note

If you are deploying to the AWS GovCloud (US-West) Region, set up the web application and VDI subdomains in the commercial partition account hosting the domain public hosted zone.

- 1. Open the Route 53 console.
- 2. Find the domain you created and choose **Create record**.
- Enter 'web' as the **Record name**. 3.
- 4. Select **CNAME** as the **Record type**.
- 5. For **Value**, enter the link you received in the initial email.
- Choose Create records. 6.
- To create a record for the VDC, retrieve the NLB address. 7.
 - a. Open the AWS CloudFormation console.
 - Choose <environment-name>-vdc.

Setting passwords for users 71

- Choose **Resources** and open <environmentname>-vdc-external-nlb. c.
- Copy the DNS name from the NLB. d.
- Open the Route 53 console.
- Find your domain and choose **Create record**.
- 10. Under **Record name**, enter vdc.
- 11. Under **Record type**, select **CNAME**.
- 12. For the NLB, enter the DNS.
- 13. Choose **Create record**.

Create an ACM certificate

By default, RES hosts the web portal under an application load balancer using the domain amazonaws.com. To use your own domain, you will need to configure a public SSL/TLS certificate provided by you or requested from AWS Certificate Manager (ACM). If you use ACM, you will receive an AWS resource name you will need to provide as a parameter to encrypt the SSL/TLS channel between the client and web services host.



🚺 Tip

If you are deploying the external resources demo package, you will need to enter your chosen domain in PortalDomainName when deploying the external resources stack in Create external resources.

To create a certificate for custom domains:

- From the console, open AWS Certificate Manager to request a public certificate. If you are deploying in AWS GovCloud (US-West), create the certificate in your GovCloud partition account.
- 2. Choose **Request a public certificate**, and choose **Next**.
- Under **Domain names**, request a certificate for both *.PortalDomainName and PortalDomainName.
- Under Validation method, choose DNS validation. 4.
- Choose Request.

Create an ACM certificate 72 From the **Certificates** list, open your requested certificates. Each certificate will have **Pending** validation as the status.



Note

If you do not see your certificates, refresh the list.

- Do one of the following:
 - **Commercial deployment:**

From the **Certificate details** for each requested certificate, choose **Create records in Route 53**. The status of the certificate should change to **Issued**.

GovCloud deployment:

If you are deploying in AWS GovCloud (US-West), copy the CNAME key and value. From the commercial partition account, use the values to create a new record in the Public Hosted Zone. The status of the certificate should change to **Issued**.

Copy the new certificate ARN to input as the parameter for ACMCertificateARNforWebApp. 8.

Amazon CloudWatch Logs

Research and Engineering Studio creates the following log groups in CloudWatch during installation. See the following table for default retentions:

CloudWatch Log groups	Retention
<pre>/aws/lambda/ <installation-stack- name="">-cluster-endpoints</installation-stack-></pre>	Never expire
/aws/lambda/ <installation-stack- name>-cluster-manager-scheduled- ad-sync</installation-stack- 	Never expire
/aws/lambda/ <installation-stack- name>-cluster-settings</installation-stack- 	Never expire

Amazon CloudWatch Logs 73

CloudWatch Log groups	Retention
/aws/lambda/ <installation-stack- name>-oauth-credentials</installation-stack- 	Never expire
<pre>/aws/lambda/ <installation-stack- name="">-self-signed-certificate</installation-stack-></pre>	Never expire
<pre>/aws/lambda/ <installation-stack- name="">-update-cluster-prefix-list</installation-stack-></pre>	Never expire
<pre>/aws/lambda/ <installation-stack- name="">-vdc-scheduled-event-transf ormer</installation-stack-></pre>	Never expire
/aws/lambda/ <installation-stack- name="">-vdc-update-cluster-manager -client-scope</installation-stack->	Never expire
<pre>/<installation-stack-name> / cluster-manager</installation-stack-name></pre>	3 months
<pre>/<installation-stack-name> /vdc/ controller</installation-stack-name></pre>	3 months
<pre>/<installation-stack-name> /vdc/ dcv-broker</installation-stack-name></pre>	3 months
<pre>/<installation-stack-name> /vdc/ dcv-connection-gateway</installation-stack-name></pre>	3 months

If you would like to change the default retention for a log group, you can go to the <u>CloudWatch</u> <u>console</u> and follow the directions to <u>Change log data retention in CloudWatch Logs</u>.

Setting custom permission boundaries

As of 2024.04, you can optionally modify roles created by RES by attaching custom permission boundaries. A custom permission boundary may be defined as part of the RES

AWS CloudFormation installation by supplying the permission boundary's ARN as part of the IAMPermissionBoundary parameter. No permission boundary is set on any RES roles if this parameter is left empty. Below is the list of actions that RES roles require to operate. Make sure that any permission boundary that you plan to use explicitly allows for the following actions:

```
{
        "Effect": "Allow",
        "Resource": "*",
        "Sid": "ResRequiredActions",
        "Action": [
            "access-analyzer:*",
            "account:GetAccountInformation",
            "account:ListRegions",
            "acm:*",
            "airflow:*",
            "amplify:*",
            "amplifybackend:*",
            "amplifyuibuilder:*",
            "aoss:*",
            "apigateway: *",
            "appflow: *",
            "application-autoscaling:*",
            "appmesh: *",
            "apprunner: *",
            "aps:*",
            "athena: *",
            "auditmanager:*",
            "autoscaling-plans:*",
            "autoscaling:*",
            "backup-gateway: *",
            "backup-storage:*",
            "backup:*",
            "batch:*",
            "bedrock: *",
            "budgets: *",
            "ce:*",
            "cloud9:*",
            "cloudformation:*",
            "cloudfront:*",
            "cloudtrail-data:*",
            "cloudtrail:*",
             "cloudwatch: *",
```

```
"codeartifact:*",
"codebuild: *",
"codeguru-profiler:*",
"codeguru-reviewer:*",
"codepipeline:*",
"codestar-connections:*",
"codestar-notifications:*",
"codestar:*",
"cognito-identity:*",
"cognito-idp:*",
"cognito-sync:*",
"comprehend: *",
"compute-optimizer:*",
"cur:*",
"databrew:*",
"datapipeline:*",
"datasync:*",
"dax:*",
"detective: *",
"devops-guru:*",
"dlm:*",
"dms:*",
"drs:*",
"dynamodb:*",
"ebs:*",
"ec2-instance-connect:*",
"ec2:*",
"ec2messages:*",
"ecr:*",
"ecs:*",
"eks:*",
"elastic-inference:*",
"elasticache:*",
"elasticbeanstalk:*",
"elasticfilesystem:*",
"elasticloadbalancing:*",
"elasticmapreduce:*",
"elastictranscoder:*",
"es:*",
"events:*",
"firehose: *",
"fis:*",
"fms:*",
"forecast: *",
```

```
"fsx:*",
"geo:*",
"glacier:*",
"glue:*",
"grafana:*",
"guardduty:*",
"health:*",
"iam:*",
"identitystore:*",
"imagebuilder:*",
"inspector2:*",
"inspector:*",
"internetmonitor:*",
"iot:*",
"iotanalytics:*",
"kafka:*",
"kafkaconnect:*",
"kinesis:*",
"kinesisanalytics:*",
"kms:*",
"lambda:*",
"lightsail:*",
"logs:*",
"memorydb:*",
"mgh: *",
"mobiletargeting:*",
"mq:*",
"neptune-db:*",
"organizations:DescribeOrganization",
"osis:*",
"personalize:*",
"pi:*",
"pipes:*",
"polly:*",
"qldb:*",
"quicksight:*",
"rds-data:*",
"rds:*",
"redshift-data:*",
"redshift-serverless:*",
"redshift:*",
"rekognition:*",
"resiliencehub:*",
"resource-groups:*",
```

```
"route53:*",
             "route53domains:*",
             "route53resolver:*",
             "rum:*",
             "s3:*",
             "sagemaker:*",
             "scheduler:*",
             "schemas:*",
             "sdb:*",
             "secretsmanager:*",
             "securityhub:*",
            "serverlessrepo:*",
            "servicecatalog:*",
             "servicequotas:*",
             "ses:*",
             "signer:*",
             "sns:*",
             "sqs:*",
             "ssm:*",
             "ssmmessages:*",
             "states:*",
             "storagegateway: *",
             "sts:*",
             "support:*",
             "tag:GetResources",
             "tag:GetTagKeys",
             "tag:GetTagValues",
             "textract:*",
             "timestream: *",
             "transcribe:*",
             "transfer: *",
             "translate: *",
             "vpc-lattice:*",
             "waf-regional:*",
             "waf:*",
             "wafv2:*",
             "wellarchitected:*",
             "wisdom: *",
             "xray:*"
        ]
    }
]
```

Configure RES-ready AMIs

With RES-ready Amazon Machine Images (AMIs), you can pre-install RES dependencies for virtual desktop instances (VDIs) on your custom AMIs. Using RES-ready AMIs improve boot times for VDI instances using the pre-baked images. Using EC2 Image Builder, you can build and register your AMIs as new software stacks. For more information on Image Builder, see the Image Builder User Guide.

Before you begin, you must deploy the latest version of RES.

Topics

- Prepare IAM role to access RES environment
- Create EC2 Image Builder component
- Prepare your EC2 Image Builder recipe
- Configure EC2 Image Builder infrastructure
- Configure Image Builder image pipeline
- Run Image Builder image pipeline
- Register a new software stack in RES

Prepare IAM role to access RES environment

To access the RES environment service from EC2 Image Builder, you must create or modify an IAM role called RES-EC2InstanceProfileForImageBuilder. For information on configuring an IAM role for use in Image Builder, see ACCESS Management (IAM) in the Image Builder User Guide.

Your role requires:

- Trusted relationships include the Amazon EC2 service.
- AmazonSSMManagedInstanceCore and EC2InstanceProfileForImageBuilder policies.
- Custom RES policy with limited DynamoDB and Amazon S3 access to the deployed RES environment.

(This policy can be either a customer managed or customer inline policy document.)

Trusted relationship entity:

Configure RES-ready AMIs 79

JSON

RES policy:

JSON

```
}
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "RESDynamoDBAccess",
            "Effect": "Allow",
            "Action": "dynamodb:GetItem",
            "Resource": "arn:aws:dynamodb:us-east-1:{AWS-Account-ID}:table/{RES-
EnvironmentName \}.cluster-settings",
            "Condition": {
                "ForAllValues:StringLike": {
                    "dynamodb:LeadingKeys": [
                        "global-settings.gpu_settings.*",
                        "global-settings.package_config.*",
                        "cluster-manager.host_modules.*",
                        "identity-provider.cognito.enable_native_user_login"
                    ]
                }
            }
        },
            "Sid": "RESS3Access",
```

Create EC2 Image Builder component

Follow the directions to <u>Create a component using the Image Builder console</u> in the *Image Builder User Guide*.

Enter your component details:

- 1. For **Type**, choose **Build**.
- 2. For **Image operating system (OS)**, choose either Linux or Windows.
- For Component name, enter a meaningful name such as research-and-engineeringstudio-vdi-<operating-system>.
- 4. Enter your component's version number and optionally add a description.
- 5. For the **Definition document**, enter the following definition file. If you encounter any errors, the YAML file is space sensitive and is the most likely cause.

Linux

```
# Copyright Amazon.com, Inc. or its affiliates. All Rights Reserved.
#
# Licensed under the Apache License, Version 2.0 (the "License"). You may not use this file except in compliance
# with the License. A copy of the License is located at
# http://www.apache.org/licenses/LICENSE-2.0
#
# or in the 'license' file accompanying this file. This file is distributed on an 'AS IS' BASIS, WITHOUT WARRANTIES
```

```
# OR CONDITIONS OF ANY KIND, express or implied. See the License for the
specific language governing permissions
  and limitations under the License.
name: research-and-engineering-studio-vdi-linux
description: An RES EC2 Image Builder component to install required RES software
dependencies for Linux VDI.
schemaVersion: 1.0
parameters:
  - AWSAccountID:
      type: string
      description: RES Environment AWS Account ID
  - RESEnvName:
      type: string
      description: RES Environment Name
  - RESEnvRegion:
      type: string
      description: RES Environment Region
  - RESEnvReleaseVersion:
      type: string
      description: RES Release Version
phases:
  - name: build
    steps:
       - name: PrepareRESBootstrap
         action: ExecuteBash
         onFailure: Abort
         maxAttempts: 3
         inputs:
            commands:
                - 'mkdir -p /root/bootstrap/logs'
                - 'mkdir -p /root/bootstrap/latest'
       - name: DownloadRESLinuxInstallPackage
         action: S3Download
         onFailure: Abort
         maxAttempts: 3
         inputs:
            - source: 's3://{{ RESEnvName }}-cluster-{{ RESEnvRegion }}-
{{ AWSAccountID }}/idea/vdc/res-ready-install-script-packages/linux/
res_linux_install_{{    RESEnvReleaseVersion }}.tar.gz'
              destination: '/root/bootstrap/
res_linux_install_{{    RESEnvReleaseVersion }}.tar.gz'
              expectedBucketOwner: '{{ AWSAccountID }}'
       - name: RunInstallScript
```

```
action: ExecuteBash
         onFailure: Abort
         maxAttempts: 3
         inputs:
            commands:
                - 'tar -xvf
{{ build.DownloadRESLinuxInstallPackage.inputs[0].destination }} -C /root/
bootstrap/latest'
                - '/bin/bash /root/bootstrap/latest/virtual-desktop-host-linux/
install.sh -r {{ RESEnvRegion }} -n {{ RESEnvName }} -g NONE'
       - name: FirstReboot
         action: Reboot
         onFailure: Abort
         maxAttempts: 3
         inputs:
            delaySeconds: 0
       - name: RunInstallPostRebootScript
         action: ExecuteBash
         onFailure: Abort
         maxAttempts: 3
         inputs:
            commands:
                - '/bin/bash /root/bootstrap/latest/virtual-desktop-host-linux/
install_post_reboot.sh'
       - name: SecondReboot
         action: Reboot
         onFailure: Abort
         maxAttempts: 3
         inputs:
            delaySeconds: 0
```

Windows

```
# Copyright Amazon.com, Inc. or its affiliates. All Rights Reserved.
#
# Licensed under the Apache License, Version 2.0 (the "License"). You may not use this file except in compliance
# with the License. A copy of the License is located at
# http://www.apache.org/licenses/LICENSE-2.0
#
# or in the 'license' file accompanying this file. This file is distributed on an 'AS IS' BASIS, WITHOUT WARRANTIES
```

```
# OR CONDITIONS OF ANY KIND, express or implied. See the License for the
specific language governing permissions
  and limitations under the License.
name: research-and-engineering-studio-vdi-windows
description: An RES EC2 Image Builder component to install required RES software
dependencies for Windows VDI.
schemaVersion: 1.0
parameters:
  - AWSAccountID:
      type: string
      description: RES Environment AWS Account ID
  - RESEnvName:
      type: string
      description: RES Environment Name
  - RESEnvRegion:
      type: string
      description: RES Environment Region
  - RESEnvReleaseVersion:
      type: string
      description: RES Release Version
phases:
  - name: build
    steps:
       - name: CreateRESBootstrapFolder
         action: CreateFolder
         onFailure: Abort
         maxAttempts: 3
         inputs:
            - path: 'C:\Users\Administrator\RES\Bootstrap'
              overwrite: true

    name: DownloadRESWindowsInstallPackage

         action: S3Download
         onFailure: Abort
         maxAttempts: 3
         inputs:
            - source: 's3://{{ RESEnvName }}-cluster-{{ RESEnvRegion }}-
{{ AWSAccountID }}/idea/vdc/res-ready-install-script-packages/windows/
res_windows_install_{{    RESEnvReleaseVersion }}.tar.gz'
              destination:
 '{{ build.CreateRESBootstrapFolder.inputs[0].path }}\res_windows_install_{{ RESEnvRelea
              expectedBucketOwner: '{{ AWSAccountID }}'
       - name: RunInstallScript
         action: ExecutePowerShell
```

```
onFailure: Abort
maxAttempts: 3
inputs:
    commands:
        'cd {{ build.CreateRESBootstrapFolder.inputs[0].path }}'
        - 'Tar -xf
res_windows_install_{{ RESEnvReleaseVersion }}.tar.gz'
        - 'Import-Module .\virtual-desktop-host-windows\Install.ps1'
        - 'Install-WindowsEC2Instance'
        - name: Reboot
        action: Reboot
        onFailure: Abort
        maxAttempts: 3
        inputs:
        delaySeconds: 0
```

6. Create any optional tags and choose **Create component**.

Prepare your EC2 Image Builder recipe

An EC2 Image Builder recipe defines the base image to use as your starting point to create a new image, along with the set of components that you add to customize your image and verify that everything works as expected. You must either create or modify a recipe to construct the target AMI with the necessary RES software dependencies. For more information on recipes, see Managerecipes.

RES supports the following image operating systems:

- Amazon Linux 2 (x86 and ARM64)
- Ubuntu 22.04.3 (x86)
- RHEL 8 (x86), and 9 (x86)
- Windows 2019, 2022 (x86)

Create a new recipe

- 1. Open the EC2 Image Builder console at https://console.aws.amazon.com/imagebuilder.
- 2. Under **Saved resources**, choose **Image recipes**.
- Choose Create image recipe.
- 4. Enter a unique name and a version number.

- 5. Select a base image supported by RES.
- 6. Under **Instance configuration**, install an SSM agent if one does not come pre-installed. Enter the information in **User data** and any other needed user data.

Note

For information on how to install an SSM agent, see:

- Manually installing SSM Agent on EC2 instances for Linux.
- Manually installing and uninstalling SSM Agent on EC2 instances for Windows Server.
- 7. For Linux based recipes, add the Amazon-managed aws-cli-version-2-linux build component to the recipe. RES installation scripts use the AWS CLI to provide VDI access to configuration values for the DynamoDB cluster-settings. Windows does not require this component.
- 8. Add the EC2 Image Builder component created for your Linux or Windows environment and enter any required parameter values. The following parameters are required inputs: AWSAccountID, RESEnvName, RESEnvRegion, and RESEnvReleaseVersion.

▲ Important

For Linux environments, you must add these components in order with the aws-cli-version-2-linux build component added first.

- (Recommended) Add the Amazon-managed simple-boot-test-linux-or-windows> test component to verify that the AMI can be launched. This is a minimum recommendation. You may select other test components that meet your requirements.
- 10. Complete any optional sections if needed, add any other desired components, and choose **Create recipe**.

Modify a recipe

If you have an existing EC2 Image Builder recipe, you can use it by adding the following components:

- For Linux based recipes, add the Amazon-managed aws-cli-version-2-linux build component to the recipe. RES installation scripts use the AWS CLI to provide VDI access to configuration values for the DynamoDB cluster-settings. Windows does not require this component.
- 2. Add the EC2 Image Builder component created for your Linux or Windows environment and enter any required parameter values. The following parameters are required inputs: AWSAccountID, RESEnvName, RESEnvRegion, and RESEnvReleaseVersion.

For Linux environments, you must add these components in order with the aws-cli-version-2-linux build component added first.

3. Complete any optional sections if needed, add any other desired components, and choose **Create recipe**.

Configure EC2 Image Builder infrastructure

You can use infrastructure configurations to specify the Amazon EC2 infrastructure that Image Builder uses to build and test your Image Builder image. For use with RES, you can choose to create a new infrastructure configuration, or use an existing one.

- To create a new infrastructure configuration, see Create an infrastructure configuration.
- To use an existing infrastructure configuration, Update an infrastructure configuration.

To configure your Image Builder infrastructure:

- 1. For **IAM role**, enter the role you previously configured in <u>Prepare IAM role to access RES</u> environment.
- 2. For **Instance type**, choose a type with at least 4 GB of memory and supports your chosen base AMI architecture. See Amazon EC2 Instance types.
- 3. For **VPC**, **subnet**, **and security groups**, you must permit internet access to download software packages. Access must also be allowed to the cluster-settings DynamoDB table and Amazon S3 cluster bucket of the RES environment.

Configure Image Builder image pipeline

The Image Builder image pipeline assembles the base image, components for building and testing, infrastructure configuration, and distribution settings. To configure an image pipeline for RES-ready AMIs, you can choose to create a new pipeline, or use an existing one. For more information, see Create and update AMI image pipelines in the Image Builder User Guide.

Create a new Image Builder pipeline

- 1. Open the Image Builder console at https://console.aws.amazon.com/imagebuilder.
- 2. From the navigation pane, choose **Image pipelines**.
- 3. Choose **Create image pipeline**.
- 4. Specify your pipeline details by entering a unique name, optional description, schedule, and frequency.
- 5. For **Choose recipe**, choose **Use existing recipe** and select the recipe created in <u>Prepare your EC2 Image Builder recipe</u>. Verify that your recipe details are correct.
- 6. For **Define image creation process**, choose either the default or custom workflow depending on the use case. In most cases, the default workflows are sufficient. For more information, see Configure image workflows for your EC2 Image Builder pipeline.
- 7. For **Define infrastructure configuration**, choose **Choose existing infrastructure configuration** and select the infrastructure configuration created in <u>Configure EC2 Image</u>
 Builder infrastructure. Verify that your infrastructure details are correct.
- 8. For **Define distribution settings**, choose **Create distribution settings using service defaults**. The output image must reside in the same AWS Region as your RES environment.

 Using service defaults, the image will be created in the Region where Image Builder is used.
- 9. Review the pipeline details and choose **Create pipeline**.

Modify an existing Image Builder pipeline

- 1. To use an existing pipeline, modify the details to use the recipe created in Prepare your EC2 Image Builder recipe.
- 2. Choose Save changes.

Run Image Builder image pipeline

To produce the output image configured, you must initiate the image pipeline. The building process can potentially take up to an hour depending on the number of components in the image recipe.

To run the image pipeline:

- 1. From Image pipelines, select the pipeline created in Configure Image Builder image pipeline.
- 2. From **Actions**, choose **Run pipeline**.

Register a new software stack in RES

- Follow the directions in <u>the section called "Software Stacks (AMIs)"</u> to register a software stack.
- 2. For AMI ID, enter the AMI ID of the output image built in Run Image Builder image pipeline.

Administrator guide

This administrator guide provides additional instructions for a technical audience on how to further customize and integrate with the Research and Engineering Studio on AWS product.

Topics

- Secrets management
- Cost monitoring and control
- Session management
- Environment management

Secrets management

Research and Engineering Studio maintains the following secrets using AWS Secrets Manager. RES creates secrets automatically during environment creation. Secrets entered by the administrator during environment creation are entered as parameters.

Secret name	Description	RES generated	Admin entered
<pre><envname> -sso- client-secret</envname></pre>	Single Sign-On OAuth2 Client Secret for environment	✓	
<pre><envname> -vdc- client-secret</envname></pre>	vdc ClientSecret	✓	
<pre><envname> -vdc- client-id</envname></pre>	vdc ClientId	✓	
<pre><envname> - vdc-gateway- certificate-pr ivate-key</envname></pre>	Self-Signed certifica te private key for domain	✓	
<pre><envname> - vdc-gateway-</envname></pre>	Self-Signed certifica te for domain	✓	

Secrets management 90

Secret name	Description	RES generated	Admin entered
certificate-ce rtificate			
<pre><envname> -cluster- manager-c lient-secret</envname></pre>	cluster-manager ClientSecret	✓	
<pre><envname> -cluster- manager-c lient-id</envname></pre>	cluster-manager ClientId	✓	
<pre><envname> - external- private-key</envname></pre>	Self-Signed certifica te private key for domain	√	
<pre><envname> - external- certificate</envname></pre>	Self-Signed certifica te for domain	✓	
<pre><envname> - internal- private-key</envname></pre>	Self-Signed certifica te private key for domain	√	
<pre><envname> - internal- certificate</envname></pre>	Self-Signed certifica te for domain	✓	
<pre><envname> -director yservice- ServiceAc countUserDN</envname></pre>	The Distinguished Name (DN) attribute of the ServiceAccount user.	✓	

Secrets management 91

The following secret ARN values are contained in the <envname>-cluster-settings table in
DynamoDB:

Key	Source
<pre>identity-provider.cognito.sso_client_secret</pre>	
<pre>vdc.dcv_connection_gateway.certifica te.certificate_secret_arn</pre>	stack
<pre>vdc.dcv_connection_gateway.certifica te.private_key_secret_arn</pre>	stack
<pre>cluster.load_balancers.internal_alb. certificates.private_key_secret_arn</pre>	stack
directoryservice.root_username_secret_arn	
vdc.client_secret	stack
<pre>cluster.load_balancers.external_alb. certificates.certificate_secret_arn</pre>	stack
<pre>cluster.load_balancers.internal_alb. certificates.certificate_secret_arn</pre>	stack
directoryservice.root_password_secret_arn	
cluster.secretsmanager.kms_key_id	
<pre>cluster.load_balancers.external_alb. certificates.private_key_secret_arn</pre>	stack
cluster-manager.client_secret	

Secrets management 92

Cost monitoring and control



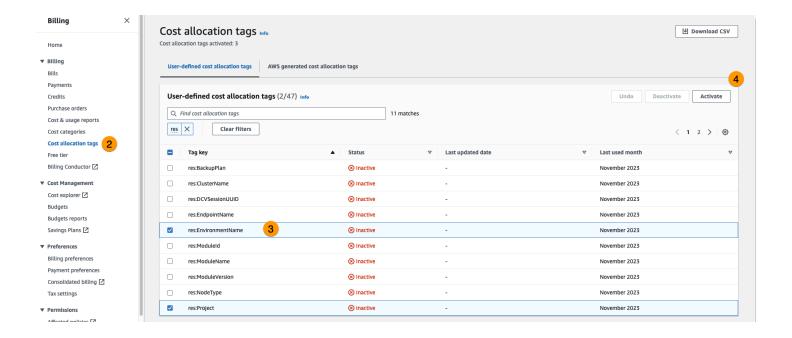
Note

Associating Research and Engineering Studio projects to AWS Budgets is not supported in AWS GovCloud (US).

We recommend creating a budget through AWS Cost Explorer to help manage costs. Prices are subject to change. For full details, see the pricing webpage for each of the the section called "AWS services in this product".

To assist with cost tracking, you can associate RES projects to budgets created within AWS Budgets. You will first need to activate the environment tags within the billing cost allocation tags.

- Sign in to the AWS Management Console and open the AWS Billing and Cost Management 1. console at https://console.aws.amazon.com/costmanagement/.
- Choose Cost allocation tags. 2.
- 3. Search for and select the res:Project and res:EnvironmentName tags.
- Choose Activate. 4.



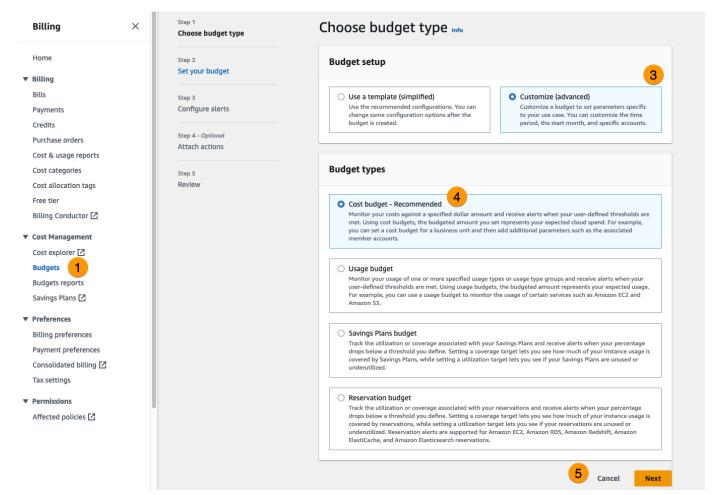


Note

It may take up to a day for RES tags to appear following deployment.

To create a budget for RES resources:

- 1. From the Billing console, choose **Budgets**.
- 2. Choose Create a budget.
- 3. Under Budget setup, choose Customize (advanced).
- 4. Under **Budget types**, choose **Cost budget - Recommended**.
- 5. Choose Next.



Under Details, enter a meaningful Budget name for your budget to distinguish it from other budgets in your account. For example, < EnvironmentName > - < ProjectName > -<BudgetName>.

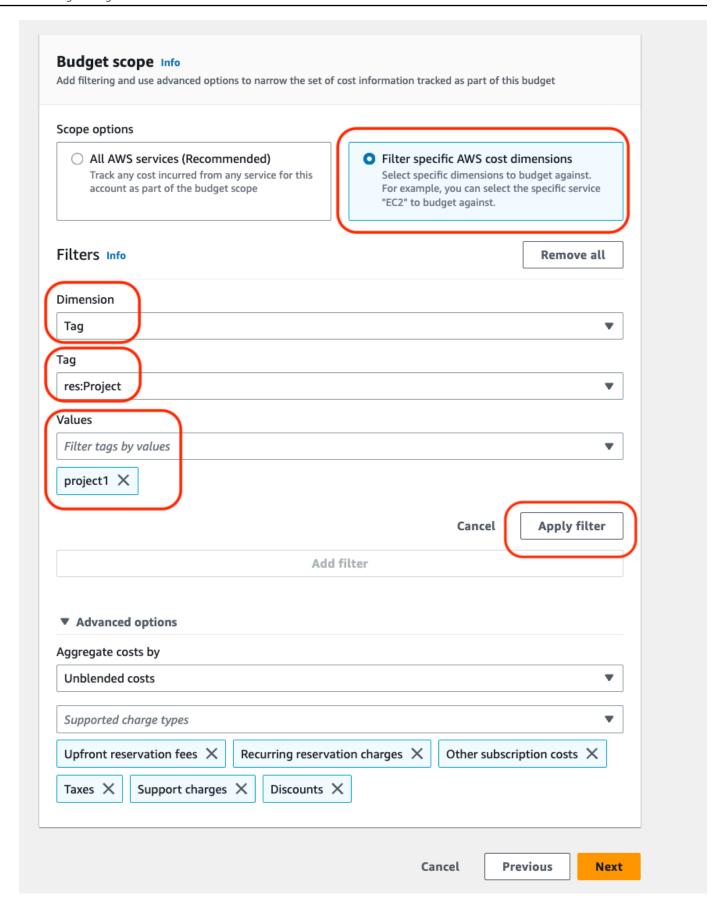
- Under **Set budget amount**, enter the amount budgeted for your project. 7.
- Under **Budget scope**, choose **Filter specific AWS cost dimensions**. 8.
- Choose **Add filter**. 9.
- 10. Under **Dimension**, choose **Tag**.
- 11. Under Tag, select res:Project.



Note

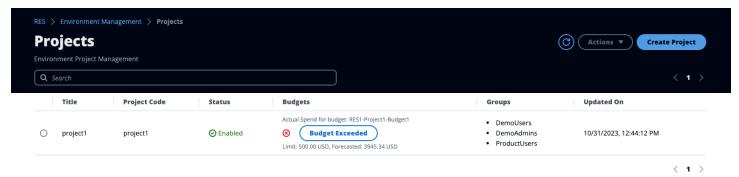
It may take up to two days for tags and values to become available. You can create a budget once the project name becomes available.

- 12. Under Values, select the project name.
- 13. Choose **Apply filter** to attach the project filter to the budget.
- 14. Choose Next.



- 15. (Optional.) Add an alert threshold.
- 16. Choose Next.
- 17. (Optional.) If an alert was configured, use **Attach actions** to configure desired actions with the
- Choose Next.
- Review the budget configuration and confirm the correct tag was set under Additional budget parameters.
- 20. Choose Create budget.

Now that the budget has been created, you can enable the budget for projects. To turn on budgets for a project, see <u>the section called "Edit a project"</u>. Virtual desktops will be blocked from launching if the budget is exceeded. If the budget is exceeded while a desktop is launched, the desktop will continue to operate.



If you need to change your budget, return to the console to edit the budget amount. It may take up to fifteen minutes for the change to take effect within RES. Alternatively, you may edit a project to disable a budget.

Session management

Session management provides a flexible and interactive environment for developing and testing sessions. As an administrative user, you can permit users to create and manage interactive sessions within their project environments.

Topics

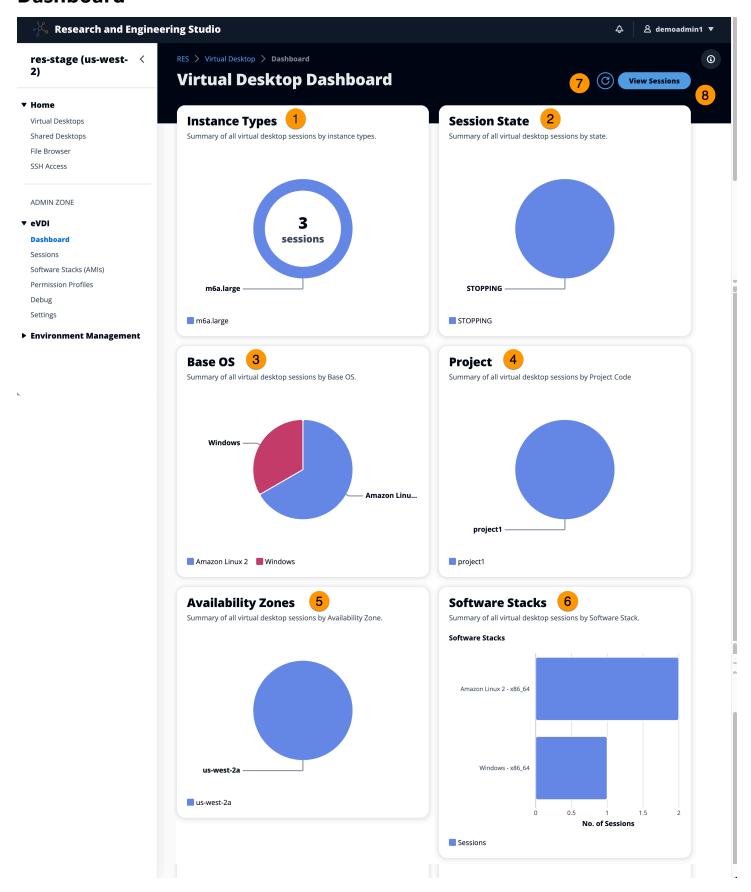
- Dashboard
- Sessions

Session management 97

- Software Stacks (AMIs)
- <u>Debugging</u>
- Desktop settings

Session management 98

Dashboard



Dashboard 99

The Session Management Dashboard provides administrators with a quick view into:

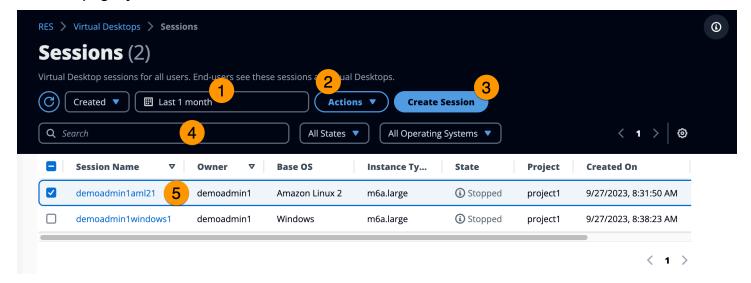
- 1. Instance types
- 2. Session states
- 3. Base OS
- 4. Projects
- 5. Availability zones
- 6. Software stacks

Additionally, administrators can:

- 7. Refresh the dashboard to update information.
- 8. Choose **View Sessions** to navigate to Sessions.

Sessions

Sessions displays all virtual desktops created within Research and Engineering Studio. From the Sessions page, you can filter and view session information or create a new session.



- 1. Use the menu to filter results by sessions created or updated within a specified time frame.
- 2. Select a session and use the Actions menu to:
 - a. Resume Session(s)
 - b. Stop/Hibernate Session(s)

Sessions 100

- c. Force Stop/Hibernate Session(s)
- d. Terminate Session(s)
- e. Force Terminate Session(s)
- f. Session(s) Health
- g. Create Software Stack
- 3. Choose **Create Session** to create a new session.
- 4. Search for a session by name and filter by state and operating system.
- 5. Select the **Session Name** to view more details.

Create a session

- 1. Choose Create Session. The Launch New Virtual Desktop modal opens.
- 2. Enter details for the new session.
- 3. (Optional.) Turn on **Show Advanced Options** to provide additional details such as subnet ID and DCV session type.
- 4. Choose Submit.

Sessions 101

Session Name	
Enter a name for the virtual desktop	
Session Name is required. Use any characters and form a name of length bo characters, inclusive.	etween 3 and 24
U-a-r	
User Select the user to create the session for	
Q	
Project	
Select the project under which the session will get created	
	▼
Select the operating system for the virtual desktop	
Select the operating system for the virtual desktop Amazon Linux 2	▼
Amazon Linux 2	▼
Amazon Linux 2 Software Stack	▼
Amazon Linux 2 Software Stack Select the software stack for your virtual desktop	▼
Amazon Linux 2 Software Stack Select the software stack for your virtual desktop	•
Amazon Linux 2 Software Stack Select the software stack for your virtual desktop Enable Instance Hibernation	
Amazon Linux 2 Software Stack Select the software stack for your virtual desktop Enable Instance Hibernation Hibernation saves the contents from the instance memory (RAM) to your Ar	
Amazon Linux 2 Software Stack Select the software stack for your virtual desktop Enable Instance Hibernation Hibernation saves the contents from the instance memory (RAM) to your Ar	
Amazon Linux 2 Software Stack Select the software stack for your virtual desktop Enable Instance Hibernation Hibernation saves the contents from the instance memory (RAM) to your Ar Store (Amazon EBS) root volume. You can not change instance type if you en	
Amazon Linux 2 Software Stack Select the software stack for your virtual desktop Enable Instance Hibernation Hibernation saves the contents from the instance memory (RAM) to your Ar Store (Amazon EBS) root volume. You can not change instance type if you enable of the contents of the	
Amazon Linux 2 Software Stack	

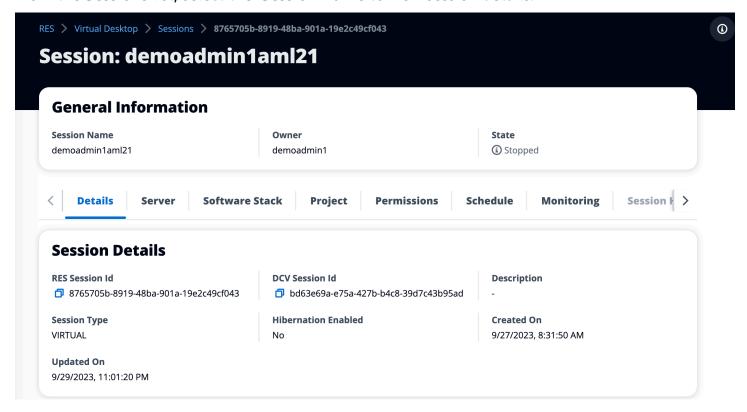
Sessions

Enter the storage size for your virtual desktop in GBs

10

Session details

From the Sessions list, select the Session Name to view session details.



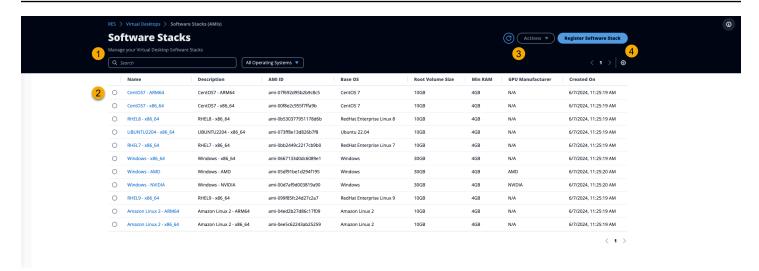
Software Stacks (AMIs)



To run the provided CentSO7 software stack in AWS GovCloud (US), you will need to subscribe to the AMI within AWS Marketplace using your linked standard account.

From the Software Stacks page, you can configure Amazon Machine Images (AMIs) or manage existing ones.

Software Stacks (AMIs) 103



- 1. To search for an existing software stack, use the operating system drop-down to filter by OS.
- 2. Select the name of a software stack to view details about the stack.
- 3. Once you select a software stack, use the **Actions** menu to edit the stack and assign the stack to a project.
- 4. The **Register Software Stack** button lets you create a new stack:
 - 1. Choose Register Software Stack.
 - 2. Enter details for the new software stack.
 - 3. Choose **Submit**.

Software Stacks (AMIs) 104

lame	
nter a name for the software stack	
lse any characters and form a name of length between 3 and 24 characters, ir	nclusive.
Description	
nter a user friendly description for the software stack	
AMI Id	
nter the AMI Id	
MI ld must start with ami-xxx	
in id mast start with ann ww	
Operating System	
elect the operating system for the software stack	
elect the operating system for the software stack Amazon Linux 2	
elect the operating system for the software stack	
Amazon Linux 2 SPU Manufacturer	
Amazon Linux 2 SPU Manufacturer elect the GPU Manufacturer for the software stack N/A	
Amazon Linux 2 SPU Manufacturer elect the GPU Manufacturer for the software stack	
Amazon Linux 2 SPU Manufacturer elect the GPU Manufacturer for the software stack N/A Min. Storage Size (GB)	
Amazon Linux 2 SPU Manufacturer elect the GPU Manufacturer for the software stack N/A Min. Storage Size (GB) nter the min. storage size for your virtual desktop in GBs	
Amazon Linux 2 SPU Manufacturer elect the GPU Manufacturer for the software stack N/A Min. Storage Size (GB) nter the min. storage size for your virtual desktop in GBs	

Assign software stack to a project

When you are creating a new software stack, you can assign the stack to projects. If you need to add the stack to a project after the initial creation, do the following:



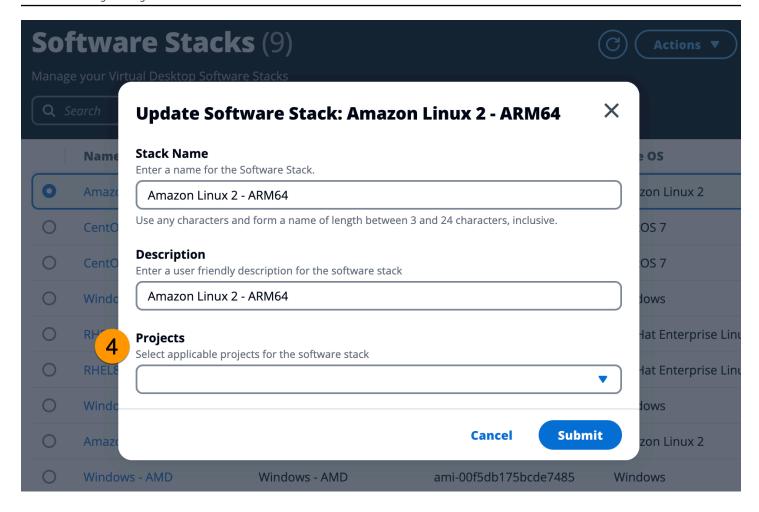
Note

You can only assign software stacks to projects of which you are a member.

- Select the software stack you need to add to a project from the Software Stacks page. 1.
- Choose Actions. 2.
- Choose Edit. 3.
- Use the **Projects** drop-down to select the project. 4.
- Choose Submit. 5.

You can also edit the software stack from the stack details page.

Software Stacks (AMIs) 106



View software stack details

From the **Software Stacks** list, select the **Software Stack Name** to view details. From the details page, you can also choose **Edit** to edit the software stack.

Debugging

The debugging panel displays message traffic associated with the virtual desktops. You can use this panel to observe activity between hosts. The VD Host tab displays instance specific activity, and the VD Sessions tab displays in-progress session activity.

Debugging 107



Desktop settings

You can use the Desktop Settings page to configure resources associated with virtual desktops. The **Server** tab provides access to settings such as:

DCV session idle timeout

The time after which the DCV session will be automatically disconnected. This does not change the state of the desktop session, it only closes the session from either the DCV client or the web browser.

Idle timeout warning

The time after which an idle warning will be provided to the client.

CPU utilization threshold

The CPU utilization to be considered idle.

Allowed sessions per user

The number of VDI sessions that an individual user can have at a given time. If a user meets or exceeds this value, this will prevent them from launching new sessions from the **My Virtual Desktops** page. The ability to launch sessions through the **Sessions** page is not impacted by this value.

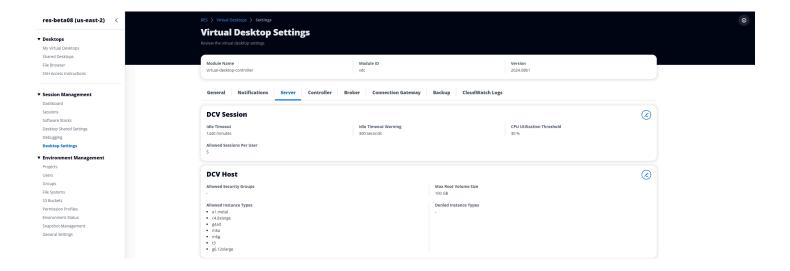
Max root volume size

The default size of the root volume on virtual desktop sessions.

Desktop settings 108

Allowed instance types

The list of instance families and sizes that can be launched for this RES environment. Instance family and instance size combinations are both accepted. For example, if you specify 'm7a', all sizes of the m7a family will be available to launch as VDI sessions. If you specify 'm7a.24xlarge', only m7a.24xlarge will be available to launch as a VDI session. This list affects all projects in the environment.



Environment management

From the Environment management section of RES, administrative users can create and manage isolated environments for their research and engineering projects. These environments can include compute resources, storage, and other necessary components, all within a secure environment. Users can configure and customize these environments to meet the specific requirements of their projects, making it easier to experiment, test, and iterate on their solutions without impacting other projects or environments.

Topics

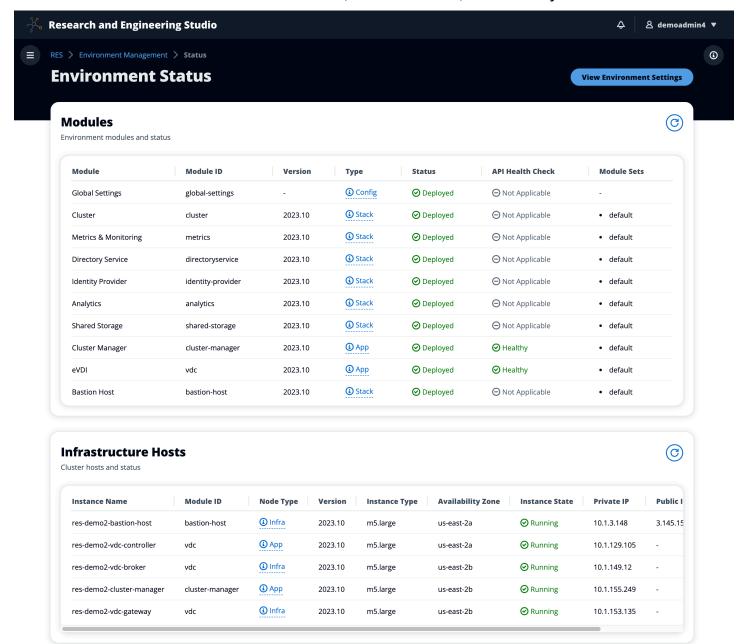
- Environment status
- Environment settings
- Users
- Groups
- Projects
- Permission policy

Environment management 109

- File Systems
- Snapshot management
- Amazon S3 buckets

Environment status

The **Environment Status** page displays the deployed software and hosts within the product. It includes information such as software version, module names, and other system information.



Environment status 110

Environment settings

The **Environment settings** page displays product configuration details, such as:

General

Displays information such as the Administrator Username and email for the user who provisioned the product. You can edit the web portal title and copyright text.

• Identity Provider

Displays information such as Single Sign-On status.

Network

Displays VPC ID, Prefix list IDs for access.

Directory Service

Displays active directory settings and service account secrets manager ARN for username and password.

Users

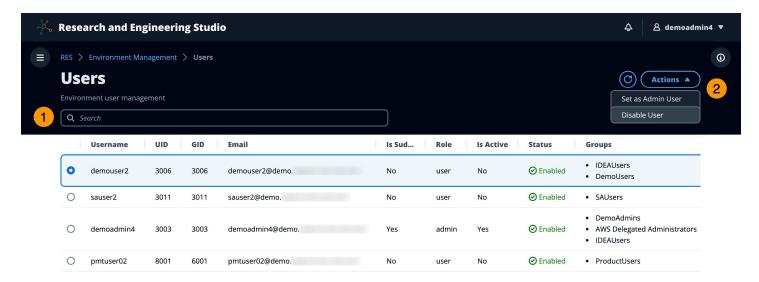
All users synced from your active directory will appear on the Users page. Users are synced by the cluster-admin user during configuration of the product. For more information on initial user configuration, see the Configuration guide.



Note

Administrators can only create sessions for active users. By default, all users will be in an inactive state until they sign in to the product environment. If a user is inactive, ask them to sign in prior to creating a session for them.

Environment settings 111



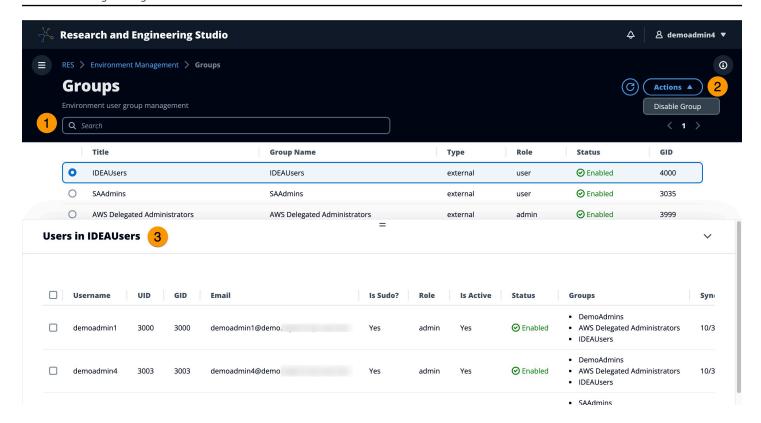
From the **Users** page, you can:

- 1. Search for users.
- 2. When a username is selected, use the **Actions** menu to:
 - a. Set as Admin user
 - b. Disable user

Groups

All Groups synced from the active directory appear on the Groups page. For more information on group configuration and management, see the *Configuration guide*.

Groups 112



From the **Groups** page, you can:

- 1. Search for user groups.
- 2. When a user group is selected, use the **Actions** menu to disable or enable a group.
- 3. When a user group is selected, you can expand the **Users** pane at the bottom of the screen to view users in the group.

Projects

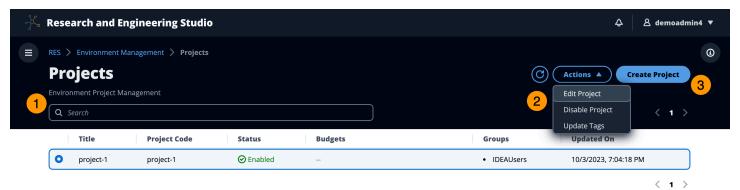
Projects form a boundary for virtual desktops, teams, and budgets. When you create a project, you define its settings, such as the name, description, and environment configuration. Projects typically include one or more environments, which can be customized to meet the specific requirements of your project, such as the type and size of the compute resources, the software stack, and the networking configuration.

Topics

- View projects
- Create a project
- Edit a project

- Add or remove tags from a project
- View file systems associated with a project
- Add a launch template

View projects



The Projects dashboard provides a list of projects available to you. From the Projects dashboard, you can:

- 1. You can use the search field to find projects.
- 2. When a project is selected, you can use the **Actions** menu to:
 - a. Edit a project
 - b. Disable or enable a project
 - c. Update project tags
- 3. You can choose **Create Project** to create a new project.

Create a project

- 1. Choose **Create Project**.
- 2. Enter project details.

The Project ID is a resource tag that can be used to track cost allocation in AWS Cost Explorer Service. For more information, see Activating user-defined cost allocation tags.



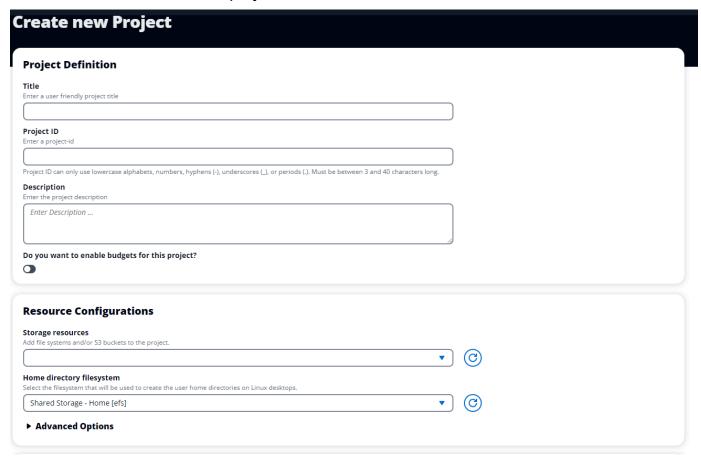
Important

The project ID cannot be changed after creation.

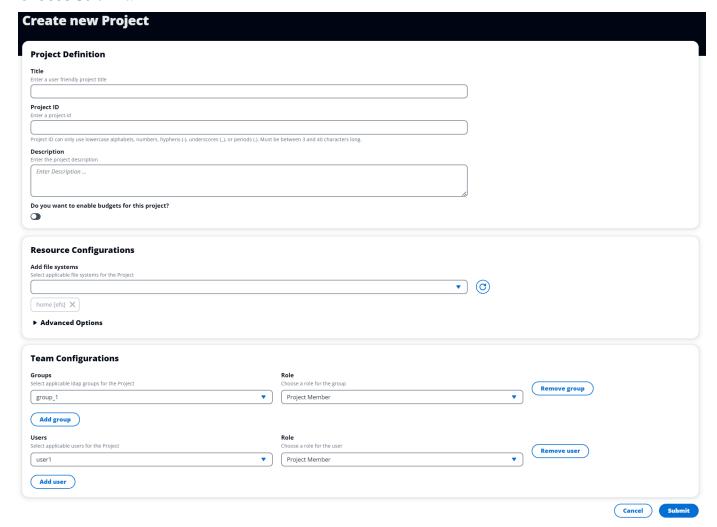
For information on **Advanced Options**, see Add a launch template.

- (Optional) Turn on budgets for the project. For more information on budgets, see Cost monitoring and control.
- The home directory filesystem may either use the Shared Home Filesystem (default), EFS, FSx for Lustre, FSx NetApp ONTAP, or EBS volume storage.

It is important to note that the shared home filesystem, EFS, FSx for Lustre, and FSx NetApp ONTAP can be shared across multiple projects and VDIs. However, the EBS volume storage option will require every VDI in that project to have their own home directory that is not shared between other VDIs or projects.

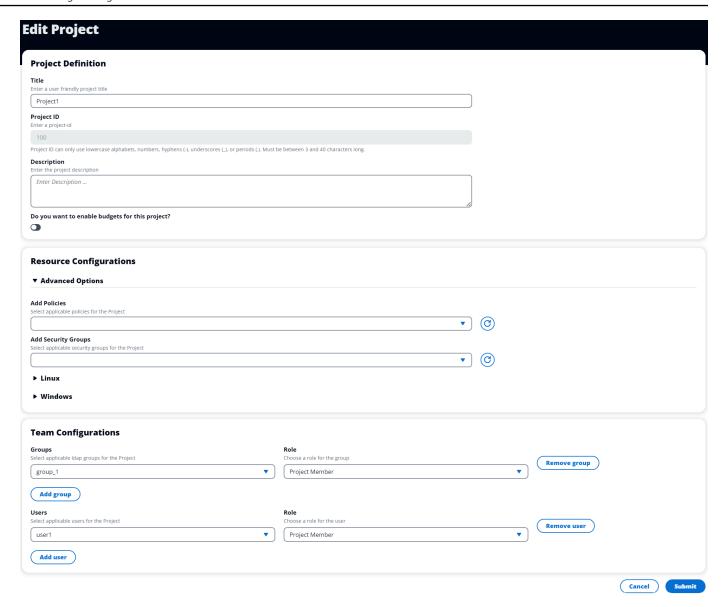


- 5. Assign users and/or groups the appropriate role ("Project Member" or "Project Owner"). See Default permissions profiles for the actions each role can take.
- 6. Choose Submit.



Edit a project

- 1. Select a project in the project list.
- 2. From the **Actions** menu, choose **Edit Project**.
- 3. Enter your updates. If you intend to enable budgets, see <u>Cost monitoring and control</u> for more information. For information on **Advanced Options**, see Add a launch template.
- Choose Submit.



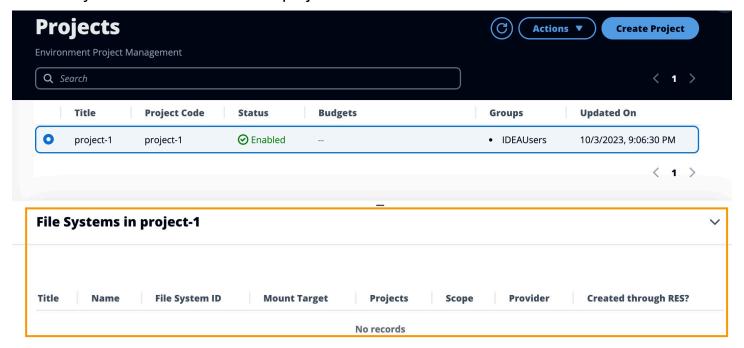
Add or remove tags from a project

Project tags will assign tags to all instances created under that project.

- 1. Select a project in the project list.
- 2. From the **Actions** menu, choose **Update Tags**.
- 3. Choose Add Tags and enter a value for Key.
- 4. To remove tags, choose **Remove** next to the tag you want to remove.

View file systems associated with a project

When a project is selected, you can expand the **File Systems** pane at the bottom of the screen to view file systems associated with the project.



Add a launch template

When creating or editing a project, you can add launch templates using the **Advanced Options** within the project configuration. Launch templates provide additional configurations, such as security groups, IAM policies, and launch scripts to all VDI instances within the project.

Add policies

You can add an IAM policy to control VDI access for all instances deployed under your project. To onboard a policy, tag the policy with the following key-value pair:

```
res:Resource/vdi-host-policy
```

For more information on IAM roles, see Policies and permissions in IAM.

Add security groups

You can add a security group to control the egress and ingress data for all VDI instances under your project. To onboard a security group, tag the security group with the following key-value pair:

res:Resource/vdi-security-group

For more information on security groups, see <u>Control traffic to your AWS resources using security</u> groups in the *Amazon VPC User Guide*.

Add launch scripts

You can add launch scripts that will initiate on all VDI sessions within your project. RES supports script initiation for Linux and Windows. For script initiation, you can choose either:

Run Script When VDI Starts

This option initiates the script at the beginning of a VDI instance before any RES configurations or installations run.

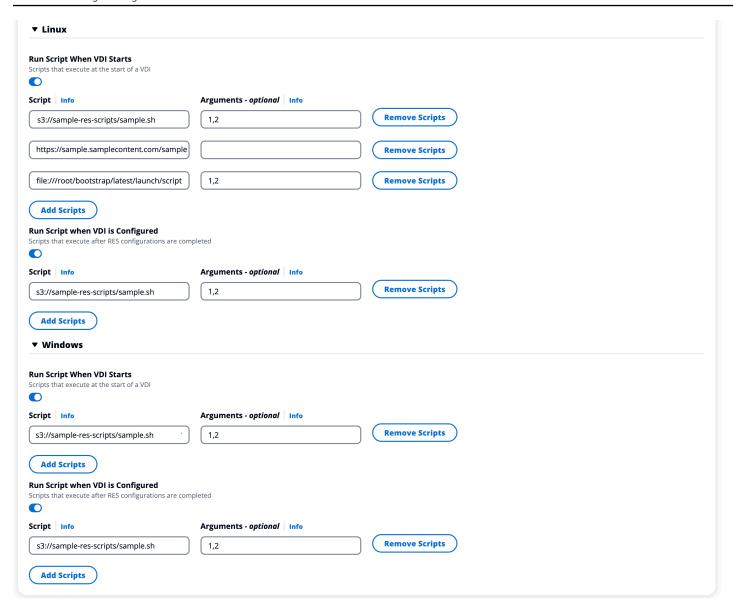
Run Script when VDI is Configured

This option initiates the script after RES configurations complete.

Scripts support the following options:

Script configuration	Example
S3 URI	s3://bucketname/script.sh
HTTPS URL	https://sample.samplecontent.com/sample
Local file	file:///user/scripts/example.sh

For **Arguments**, provide any arguments separated by a comma.



Example of a project configuration

Permission policy

Research and Engineering Studio (RES) allows an administrative user to create custom permission profiles that grant selected users additional permissions to manage the project that they are part of. Each project comes with two <u>default permission profiles</u>- "Project Member" and "Project Owner" that can be customized after deployment.

Currently, administrators can grant two collections of permissions using a permission profile:

- 1. Project management permissions which consist of "Update project membership" that allows a designated user to add other users and groups to, or remove them from, a project, and "Update project status" that allows a designated user to enable or disable a project.
- 2. VDI session management permissions which consist of "Create Session" that allows a designated user to create a VDI session within their project, and "Create/Terminate another user's session" that allows a designated user to create or terminate the sessions of other users within a project.

In this way, administrators can delegate project-based permissions to non-administrators in their environment.

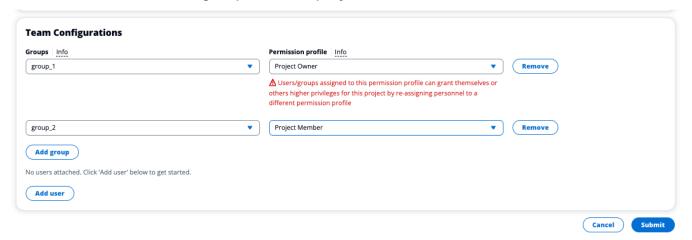
Topics

- Project management permissions
- · VDI session management permissions
- Managing permission profiles
- Default permissions profiles
- Environment Boundaries
- Desktop sharing profiles

Project management permissions

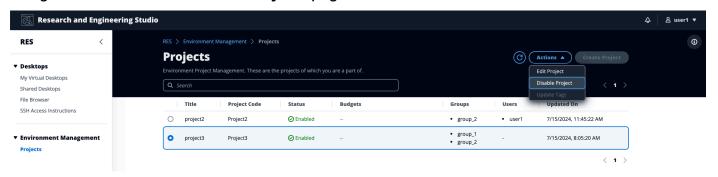
Update project membership

This permission allows non-admin users who have been granted it to add and remove users or groups from a project. It also allows them to set the permission profile and decide the access level for all other users and groups for that project.



Update project status

This permission allows non-admin users who have been granted it to enable or disable a project using the **Actions** button on the **Projects** page.

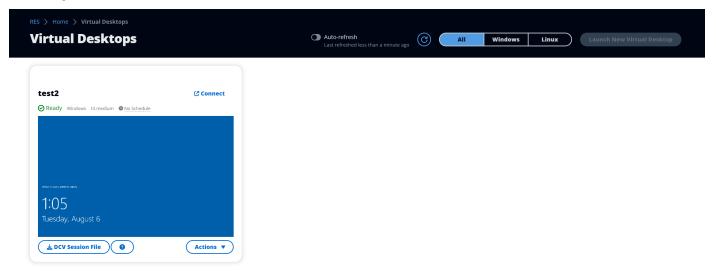


VDI session management permissions

Create a session

Controls whether or not a user is allowed to launch their own VDI session from the **My Virtual Desktops** page. Disable this to deny non-admin users the ability to launch their own VDI sessions. Users can always stop and terminate their own VDI sessions.

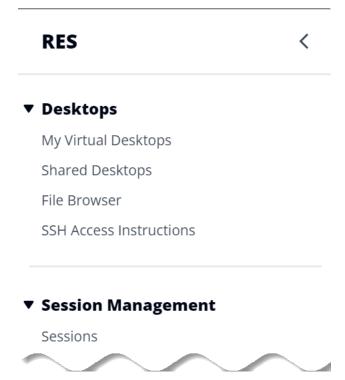
If a non-admin user does not have permissions to create a session, the **Launch New Virtual Desktop** button will be disabled for them as shown here:



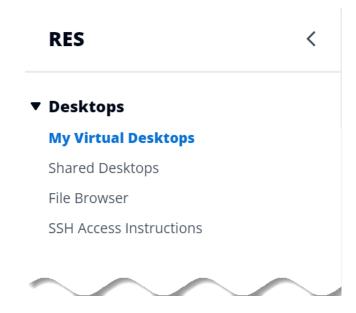
Create or Terminate the sessions of others

Allows non-admin users to access the **Sessions** page from the left-hand navigation pane. These users will be able to launch VDI sessions for other users in the projects where they have been granted this permission.

If a non-admin user has permission to launch sessions for other users, their left-hand navigation pane will display the **Sessions** link under **Session Management** as shown here:



If a non-admin user does not have permission to create sessions for others, their left-hand navigation pane will not display **Session Management** as shown here:

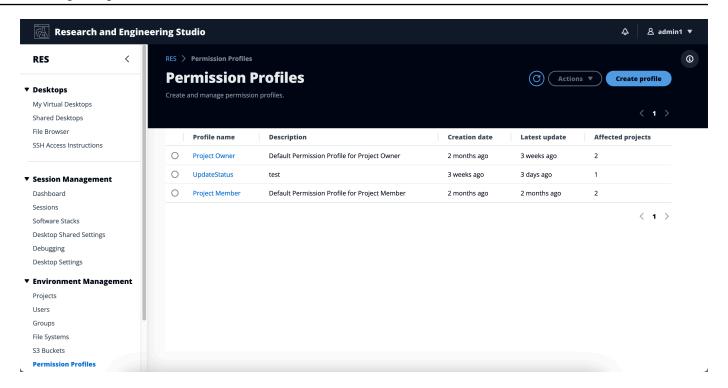


Managing permission profiles

As a RES administrator, you can perform the following actions to manage permission profiles.

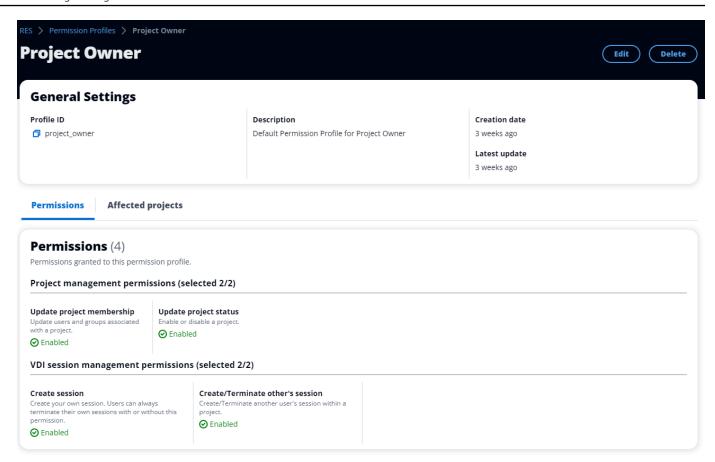
List permission profiles

• From the Research and Engineering Studio console page, choose **Permission Profiles** in the left-hand navigation pane. From this page you can create, update, list, view and delete permission profiles.

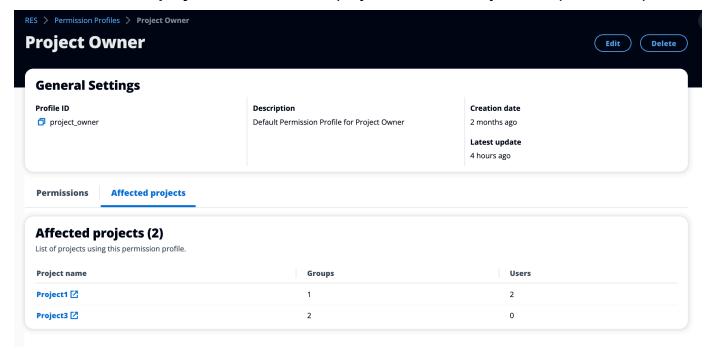


View permission profiles

1. On the main **Permission Profiles** page, select the name of the permission profile you want to view. From this page you can edit or delete the selected permission profile.

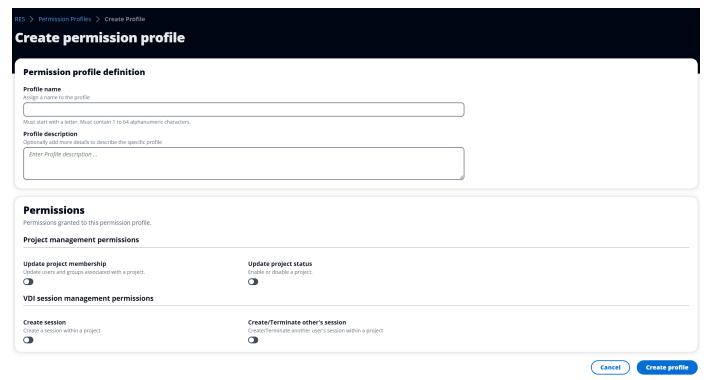


2. Select the Affected projects tab to view the projects that currently use the permission profile.



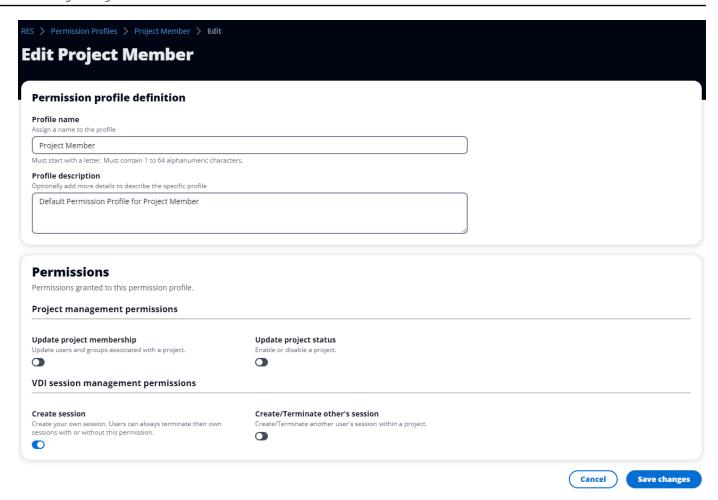
Create permission profiles

- 1. On the main **Permission Profiles** page, choose **Create profile** to create a permission profile.
- 2. Enter a permission profile name and description, then select the permissions to grant to the users or groups that you assign to this profile.



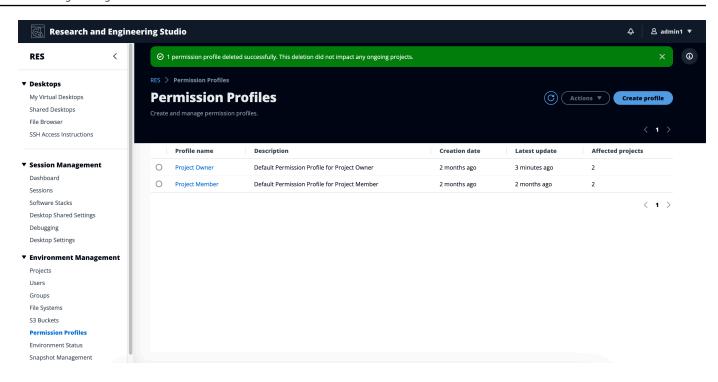
Edit permission profiles

On the main Permission Profiles page, select a profile by clicking the circle next to it, choose
 Actions, then choose Edit profile to update that permission profile.



Delete permission profiles

• On the main **Permission Profiles** page, select a profile by clicking the circle next to it, choose **Actions**, then choose **Delete profile**. You cannot delete a permission profile that is used by any existing project.



Default permissions profiles

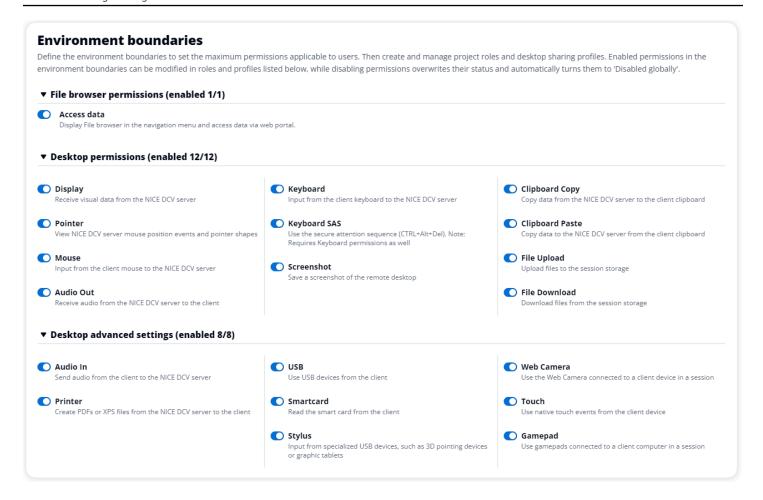
Every RES project comes with two default permission profiles that Global Administrators can configure. (In addition, Global Administrators can create and modify new permission profiles for a project.) The following table shows the allowed permissions for the default permission profiles-"Project Member" and "Project Owner". Permission profiles, and the permissions they grant to select users of a project, only apply to the project that they belong to; Global Administrators are super users who have all the permissions below across all projects.

Permissions	Description	Project Member	Project Owner
Create Session	Create your own session. Users can always stop and terminate their own sessions with or without this permission.	X	X

Permissions	Description	Project Member	Project Owner
Create/te rminate others' sessions	Create or terminate another user's session within a project.		X
Update Project membership	Update users and groups associated with a project.		X
Update Project Status	Enable or disable a project.		X

Environment Boundaries

Environment boundaries allow administrators to configure permissions that will take effect globally for all users. This includes permissions such as File Browser access and Desktop Permissions.



Configuring File browser access

Administrators can toggle **Access data** on or off under **File browser permissions**. If **Access data** is turned off, users will not see **File Browser** navigation in their web portal and cannot upload or download data attached to their global file system. When **Access data** is enabled, users have access to **File Browser** navigation in their web portal which allows them to upload or download data that is attached to their global file system.

When the **Access data** feature is turned on and then later turned off, users who are already logged in to the web portal will be unable to upload or download files, even if they are on the corresponding page. Additionally, the navigation menu will disappear when they refresh the page.

Configuring Desktop Permissions

Administrators can toggle **Desktop permissions** on or off to globally manage the VDI functionality of all owners. All of these permissions, or a subset, can be used to create **Desktop sharing profiles** that determine which actions can be performed by users with whom a desktop is shared. If any desktop permission is disabled, this will automatically disable the corresponding permissions in

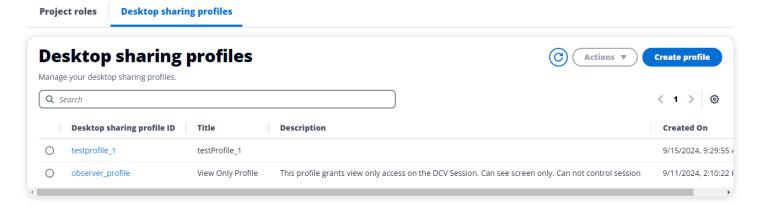
the **Desktop sharing profiles**. These permissions will be labeled as "Disabled Globally". Even if the administrator enables this desktop permission again, the permission in the desktop sharing profile will remain disabled until the administrator manually enables it.

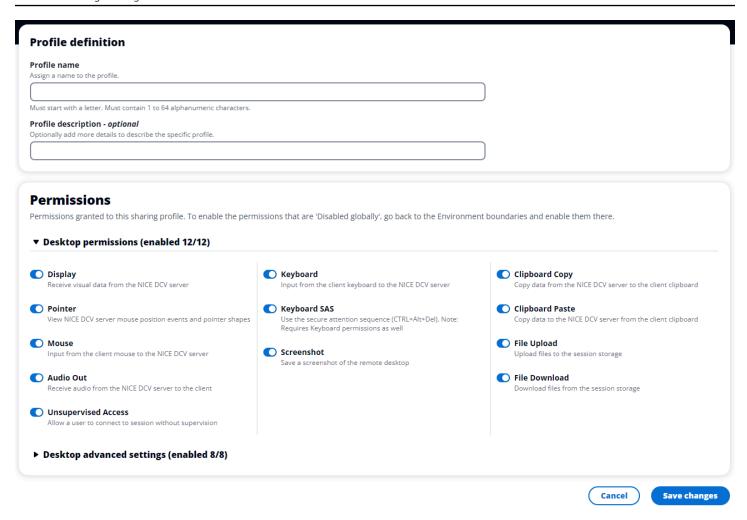
Desktop sharing profiles

Administrators can create new profiles and customize them. These profiles can be accessed by all users and are used when sharing a session with others. The maximum permissions granted within these profiles cannot exceed the desktop permissions allowed globally.

Create Profile

Administrators can choose **Create profile** to create a new profile. Then they can enter a **Profile name**, a **Profile Description**, set the desired permissions, and **Save** their changes.





Edit Profile

To edit a profile:

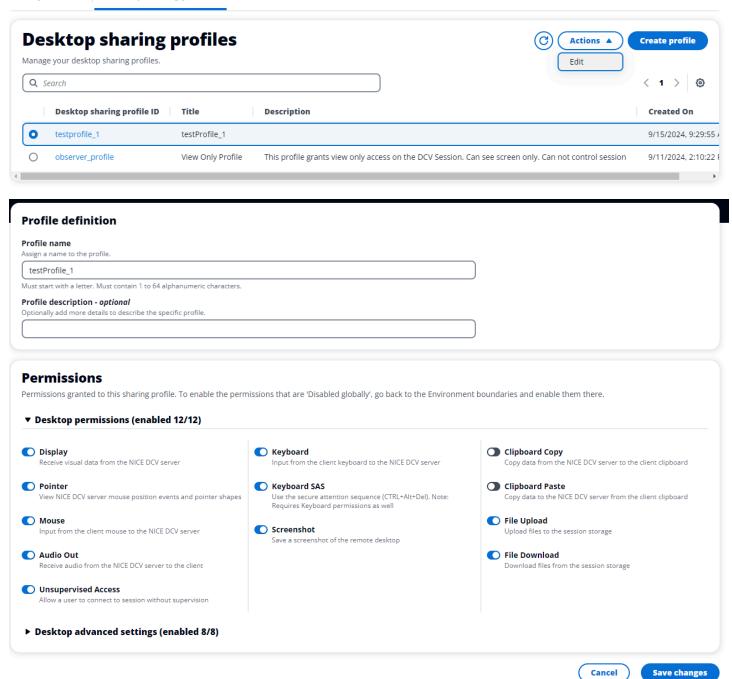
- 1. Select the desired profile.
- 2. Choose **Actions**, then select **Edit** to modify the profile.
- 3. Adjust the permissions as needed.
- 4. Choose Save changes.

Any changes made to the profile will be immediately applied to the current open sessions.

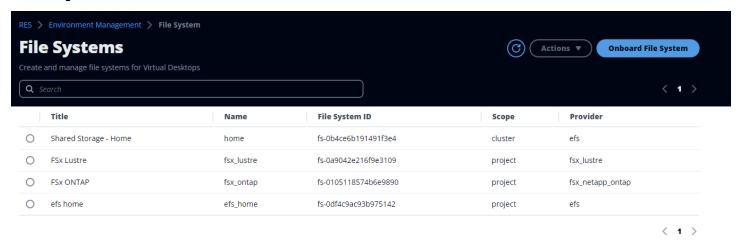
Cancel

Project roles

Desktop sharing profiles



File Systems



From the File Systems page, you can:

- 1. Search for file systems.
- 2. When a file system is selected, use the **Actions** menu to:
 - a. Add the file system to a project.
 - Remove the file system from a project
- 3. Onboard a new file system.
- 4. When a file system is selected, you can expand the pane at the bottom of the screen to view file system details.

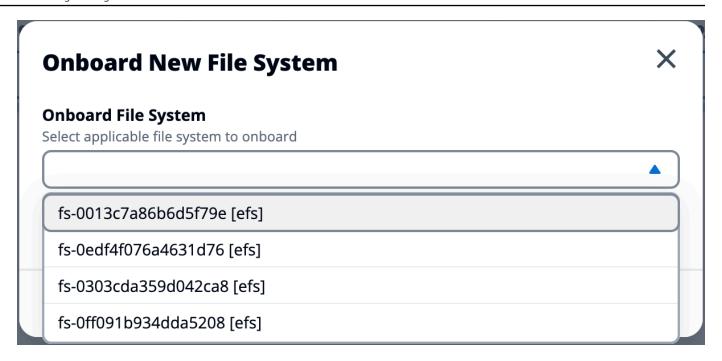
Topics

Onboard a file system

Onboard a file system

- 1. Choose **Onboard File System**.
- 2. Select a file system from the drop down. The modal will expand with additional detail entries.

File Systems 135



3. Enter file system details.



By default, administrators and project owners have the ability to choose a home filesystem when creating a new project, which cannot be edited afterwards. File systems intended to be used as home directories on projects must be onboarded by setting their **Mount Directory** path to /home. This will populate the onboarded filesystem on the home directory filesystem dropdown options. This feature helps to keep the data isolated across projects since only users associated with the project will have access to the filesystem through their VDIs. VDIs will mount the filesystem at the mount point selected during onboarding of a filesystem.

4. Choose Submit.

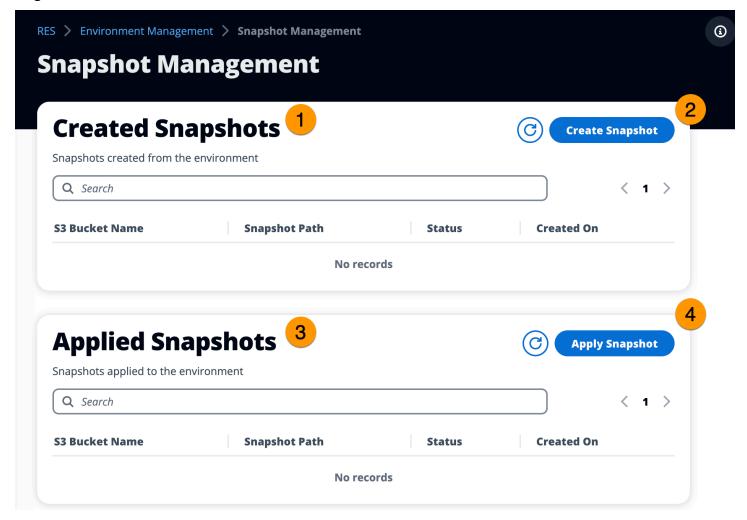
File Systems 136

Onboard New File System Onboard File System Select applicable file system to onboard fs-0edf4f076a4631d76 [efs] Title Enter a user friendly file system title **File System Name** Enter a file system name File System name cannot contain white spaces or special characters. Only use lowercase alphabets, numbers and underscore (_). Must be between 3 and 18 characters long. **Mount Directory** Enter directory to mount the file system Mount directory cannot contain white spaces or special characters. Only use lowercase alphabets, numbers, and hyphens (-). Must be between 3 and 18 characters long. Eg. /efs-01 **Submit** Cancel

File Systems 137

Snapshot management

Snapshot management simplifies the process of saving and migrating data between environments, ensuring consistency and accuracy. With snapshots, you can save your environment state and migrate data into a new environment with the same state.



From the **Snapshot management** page, you can:

- 1. View all created snapshots and their status.
- 2. Create a snapshot. Before you can create a snapshot, you will need to create a bucket with the appropriate permissions.
- 3. View all applied snapshots and their status.
- 4. Apply a snapshot.

Topics

Snapshot management 138

- Create a snapshot
- Apply a snapshot

Create a snapshot

Before you can create a snapshot, you must provide an Amazon S3 bucket with the necessary permissions. For information on creating a bucket, see Creating a bucket. We recommend enabling bucket versioning and server access logging. These settings can be enabled from the bucket's **Properties** tab after provisioning.



Note

This Amazon S3 bucket's lifecycle will not be managed within the product. You will need to manage the bucket lifecycle from the console.

To add permissions to the bucket:

- 1. Select the bucket you created from the **Buckets** list.
- Select the **Permissions** tab. 2.
- 3. Under Bucket policy, choose Edit.
- 4. Add the following statement to the bucket policy. Replace these values with your own:
 - AWS_ACCOUNT_ID
 - RES_ENVIRONMENT_NAME
 - AWS_REGION
 - S3_BUCKET_NAME



Important

There are limited version strings supported by AWS. For more information, see https://docs.aws.amazon.com/IAM/latest/UserGuide/ reference_policies_elements_version.html.

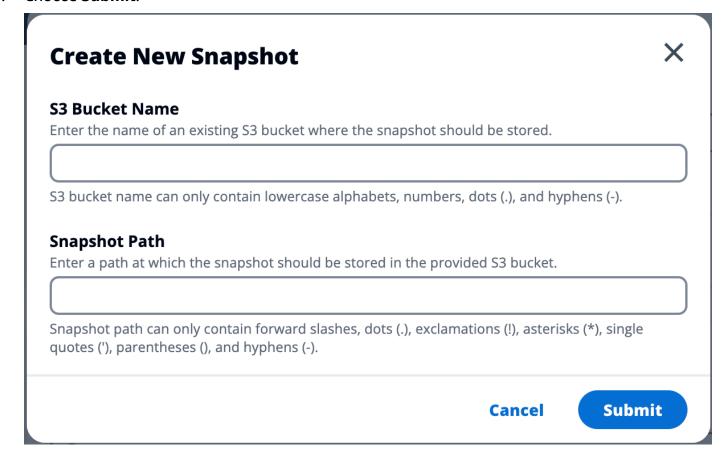
JSON

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "Export-Snapshot-Policy",
            "Effect": "Allow",
            "Principal": {
                "AWS":
 "arn:aws:iam::111122223333:role/{RES_ENVIRONMENT_NAME}-cluster-manager-
role-{AWS_REGION}}"
            },
            "Action": [
                "s3:GetObject",
                "s3:ListBucket",
                "s3:AbortMultipartUpload",
                "s3:PutObject",
                "s3:PutObjectAcl"
            ],
            "Resource": [
                "arn:aws:s3:::{S3_BUCKET_NAME}",
                "arn:aws:s3:::{S3_BUCKET_NAME}/*"
            ]
        },
        {
            "Sid": "AllowSSLRequestsOnly",
            "Action": "s3:*",
            "Effect": "Deny",
            "Resource": [
                "arn:aws:s3:::{S3_BUCKET_NAME}",
                "arn:aws:s3:::{S3_BUCKET_NAME}/*"
            ],
            "Condition": {
                "Bool": {
                    "aws:SecureTransport": "false"
                }
            },
            "Principal": "*"
        }
```

}

To create the snapshot:

- 1. Choose **Create Snapshot**.
- 2. Enter the name of the Amazon S3 bucket you created.
- Enter the path where you would like the snapshot stored within the bucket. For example, october2023/23.
- 4. Choose **Submit**.



5. After five to ten minutes, choose **Refresh** on the Snapshots page to check the status. A snapshot will not be valid until the status changes from IN_PROGRESS to COMPLETED.

Apply a snapshot

Once you have created a snapshot of an environment, you can apply that snapshot to a new environment to migrate data. You will need to add a new policy to the bucket allowing the environment to read the snapshot.

Applying a snapshot copies data such as user permissions, projects, software stacks, permission profiles, and file systems with their associations to a new environment. User sessions will not be replicated. When the snapshot is applied, it checks the basic information of each resource record to determine if it already exists. For duplicate records, the snapshot skips resource creation in the new environment. For records that are similar, such as share a name or key, but other basic resource information varies, it will create a new record with a modified name and key using the following convention: RecordName_SnapshotRESVersion_ApplySnapshotID. The ApplySnapshotID looks like a timestamp and identifies each attempt to apply a snapshot.

During the snapshot application, the snapshot checks for the availability of resources. Resource not available to the new environment will not be created. For resources with a dependent resource, the snapshot checks for the availability of the dependent resource. If the dependent resource is not available, it will create the main resource without the dependent resource.

If the new environment is not as expected or fails, you can check the CloudWatch logs found in the log group /res-<env-name>/cluster-manager for details. Each log will have the [apply snapshot] tag. Once you have applied a snapshot, you can check its status from the the section called "Snapshot management" page.

To add permissions to the bucket:

- 1. Select the bucket you created from the **Buckets** list.
- 2. Select the **Permissions** tab.
- 3. Under **Bucket policy**, choose **Edit**.
- 4. Add the following statement to the bucket policy. Replace these values with your own:
 - AWS_ACCOUNT_ID
 - RES_ENVIRONMENT_NAME
 - AWS_REGION
 - S3_BUCKET_NAME

JSON

```
"Sid": "Export-Snapshot-Policy",
            "Effect": "Allow",
            "Principal": {
                "AWS":
 "arn:aws:iam::111122223333:role/{RES_ENVIRONMENT_NAME}-cluster-manager-
role-{AWS_REGION}}"
            },
            "Action": [
                "s3:GetObject",
                "s3:ListBucket"
            ],
            "Resource": [
                "arn:aws:s3:::{S3_BUCKET_NAME}",
                "arn:aws:s3:::{S3_BUCKET_NAME}/*"
            ]
        },
        {
            "Sid": "AllowSSLRequestsOnly",
            "Action": "s3:*",
            "Effect": "Deny",
            "Resource": [
                "arn:aws:s3:::{S3_BUCKET_NAME}",
                "arn:aws:s3:::{S3_BUCKET_NAME}/*"
            ],
            "Condition": {
                "Bool": {
                     "aws:SecureTransport": "false"
            },
            "Principal": "*"
        }
    ]
}
```

To apply snapshot:

- Choose Apply snapshot.
- 2. Enter the name of the Amazon S3 bucket containing the snapshot.
- 3. Enter the file path to the snapshot within the bucket.
- Choose Submit.

Apply a Snapshot	
S3 Bucket Name	
Enter the name of the S3 bucke	et where the snapshot to be applied is stored.
S3 bucket name can only conta	in lowercase alphabets, numbers, dots (.), and hyphens (-).
Snapshot Path	
•	apshot to be applied is stored in the provided S3 bucket.
•	apshot to be applied is stored in the provided S3 bucket.
Enter the path at which the sna	n forward slashes, dots (.), exclamations (!), asterisks (*), single
Enter the path at which the sna Snapshot path can only contain	n forward slashes, dots (.), exclamations (!), asterisks (*), single

5. After five to ten minutes, choose **Refresh** on the Snapshot management page to check the status.

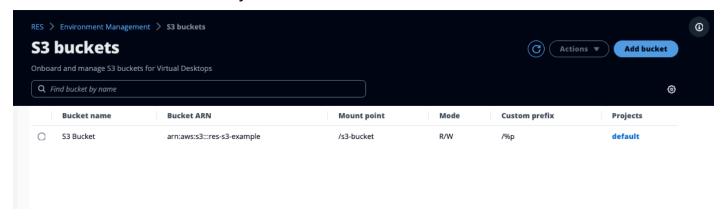
Amazon S3 buckets

Research and Engineering Studio (RES) supports mounting <u>Amazon S3 buckets</u> to Linux Virtual Desktop Infrastructure (VDI) instances. RES Administrators can onboard S3 buckets to RES, attach them to projects, edit their configuration, and remove buckets in the S3 buckets tab under **Environment Management**.

The S3 buckets dashboard provides a list of onboarded S3 buckets available to you. From the S3 buckets dashboard, you can:

- 1. Use Add bucket to onboard an S3 bucket to RES.
- 2. Select an S3 bucket and use the **Actions** menu to:
 - Edit a bucket
 - Remove a bucket

3. Use the search field to search by Bucket name and find onboarded S3 buckets.



The following sections describe how to manage Amazon S3 buckets in your RES projects.

Topics

- Amazon S3 bucket prerequisites for isolated VPC deployments
- Add an Amazon S3 bucket
- Edit an Amazon S3 bucket
- Remove an Amazon S3 bucket
- Data Isolation
- Cross account bucket access
- Preventing data exfiltration in a private VPC
- Troubleshooting
- Enabling CloudTrail

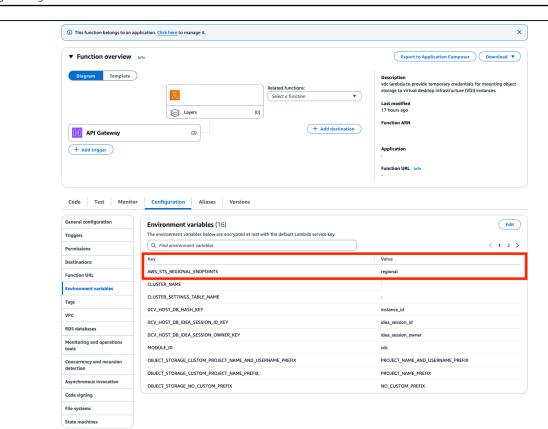
Amazon S3 bucket prerequisites for isolated VPC deployments

If you're deploying Research and Engineering Studio in an isolated VPC, follow these steps to update the lambda configuration parameters after you deploy RES in your AWS account.

- 1. Log into the Lambda Console of the AWS account where Research and Engineering Studio is deployed.
- 2. Find and navigate to the Lambda function named <RES-EnvironmentName</pre>-vdc-customcredential-broker-lambda.
- 3. Select the **Configuration** tab of the function.

(i)

0



- 4. On the left hand side, choose **Environment variables** to view that section.
- 5. Choose **Edit** and add the following new environment variable to the function:
 - Key: AWS_STS_REGIONAL_ENDPOINTS
 - · Value: regional
- 6. Choose Save.

Add an Amazon S3 bucket

To add an S3 bucket to your RES environment:

- 1. Choose Add bucket.
- 2. Enter the bucket details such as bucket name, ARN, and mount point.

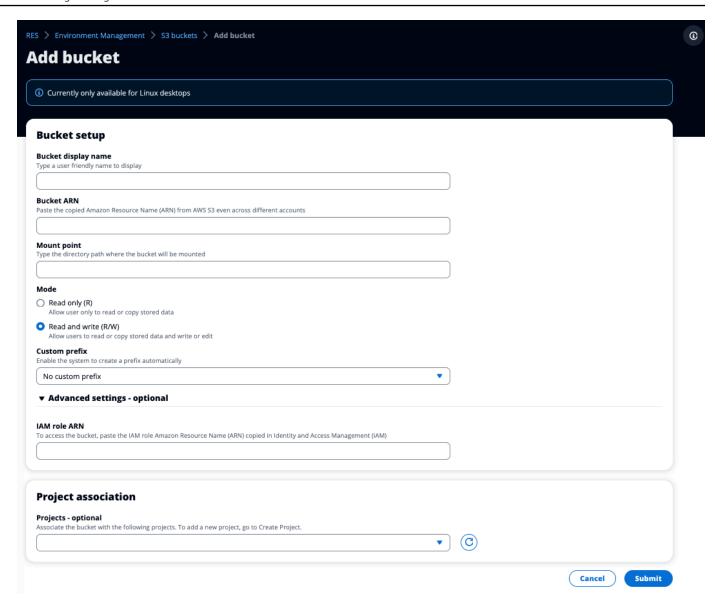
▲ Important

• The bucket ARN, mount point, and mode provided cannot be changed after creation.

- The bucket ARN can contain a prefix which will isolate the onboarded S3 bucket to that prefix.
- 3. Select a mode in which to onboard your bucket.

A Important

- See <u>Data Isolation</u> for more information related to data isolation with specific modes.
- 4. Under **Advanced Options**, you may provide an IAM role ARN to mount the buckets for cross account access. Follow the steps in <u>Cross account bucket access</u> to create the required IAM role for cross account access.
- 5. (Optional) Associate the bucket with projects, which can be changed later. However, an S3 bucket cannot be mounted to a project's existing VDI sessions. Only sessions launched after the project has been associated with the bucket will mount the bucket.
- 6. Choose **Submit**.



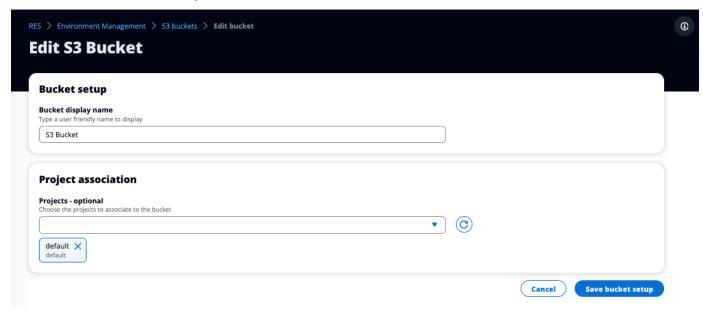
Edit an Amazon S3 bucket

- 1. Select an S3 bucket in the S3 buckets list.
- 2. From the **Actions** menu, select **Edit**.
- 3. Enter your updates.

▲ Important

 Associating a project with an S3 bucket will not mount the bucket to that project's existing virtual desktop infrastructure (VDI) instances. The bucket will only be

- mounted to VDI sessions launched in a project after the bucket has been associated with that project.
- Disassociating a project from an S3 bucket will not impact the data in the S3 bucket, but will result in desktop users losing access to that data.
- Choose Save bucket setup.



Remove an Amazon S3 bucket

- Select an S3 bucket in the S3 buckets list.
- 2. From the **Actions** menu, select **Remove**.

Important

- You must first remove all project associations from the bucket.
- The remove operation does not impact the data in the S3 bucket. It only removes the S3 bucket's association with RES.
- Removing a bucket will cause existing VDI sessions to lose access to the contents of that bucket at the expiration of that session's credentials (~1 hour).

Data Isolation

When you add an S3 bucket to RES, you have options to isolate the data within the bucket to specific projects and users. On the **Add Bucket** page, you can select a mode of Read Only (R) or Read and Write (R/W).

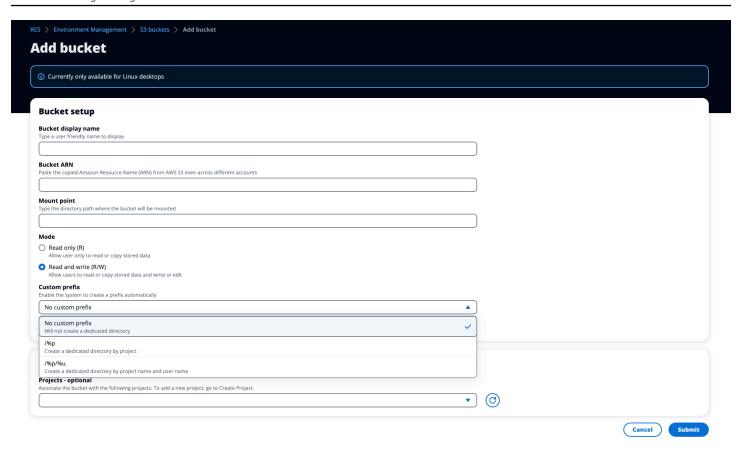
Read Only

If Read Only (R) is selected, data isolation is enforced based on the prefix of the bucket ARN (Amazon Resource Name). For example, if an admin adds a bucket to RES using the ARN arn: aws:s3:::bucket-name/example-data/ and associates this bucket with Project A and Project B, then users launching VDIs from within Project A and Project B can only read the data located in bucket-name under the path /example-data. They will not have access to data outside of that path. If there is no prefix appended to the bucket ARN, the entire bucket will be made available to any project associated with it.

Read and Write

If Read and Write (R/W) is selected, data isolation is still enforced based on the prefix of the bucket ARN, as described above. This mode has additional options to allow admins to provide variable-based prefixing for the S3 bucket. When Read and Write (R/W) is selected, a Custom Prefix section becomes available that offers a dropdown menu with the following options:

- No custom prefix
- /%p
- /%p/%u



No custom data isolation

When No custom prefix is selected for **Custom Prefix**, the bucket is added without any custom data isolation. This allows any projects associated with the bucket to have read and write access. For example, if an admin adds a bucket to RES using the ARN arn:aws:s3:::bucket-name with No custom prefix selected and associates this bucket with Project A and Project B, users launching VDIs from within Project A and Project B will have unrestricted read and write access to the bucket.

Data isolation on a per-project level

When /%p is selected for **Custom Prefix**, data in the bucket is isolated to each specific project associated with it. The %p variable represents the project code. For example, if an admin adds a bucket to RES using the ARN arn:aws:s3:::bucket-name with /%p selected and a **Mount Point** of /bucket, and associates this bucket with Project A and Project B, then User A in Project A can write a file to /bucket. User B in Project A can also see the file that User A wrote in /bucket. However, if User B launches a VDI in Project B and looks in /bucket, they will not see the file that User A wrote, as the data is isolated by project. The file User A wrote is found

in the S3 bucket under the prefix /ProjectA while User B can only access /ProjectB when using their VDIs from Project B.

Data isolation on a per-project, per-user level

When /%p/%u is selected for **Custom Prefix**, data in the bucket is isolated to each specific project and user associated with that project. The %p variable represents the project code, and %u represents the username. For example, an admin adds a bucket to RES using the ARN arn:aws:s3:::bucket-name with /%p/%u selected and a Mount Point of /bucket. This bucket is associated with Project A and Project B. User A in Project A can write a file to /bucket. Unlike the prior scenario with only %p isolation, User B in this case will not see the file User A wrote in Project A in /bucket, as the data is isolated by both project and user. The file User A wrote is found in the S3 bucket under the prefix /ProjectA/UserA while User B can only access /ProjectA/UserB when using their VDIs in Project A.

Cross account bucket access

RES has the ability to mount buckets from other AWS accounts, provided these buckets have the right permissions. In the following scenario, a RES environment in Account A wants to mount an S3 bucket in Account B.

Step 1: Create an IAM Role in the account that RES is deployed in (this will be referred to as Account A):

- 1. Sign in to the AWS Management Console for the RES account that needs access to the S3 bucket (Account A).
- 2. Open the IAM Console:
 - a. Navigate to the IAM dashboard.
 - b. In the navigation pane, choose **Policies**.
- 3. Create a Policy:
 - a. Choose Create policy.
 - b. Select the **JSON** tab.
 - c. Paste the following JSON policy (replace < BUCKET-NAME > with the name of the S3 bucket located in Account B):

JSON

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": [
                 "s3:GetObject",
                 "s3:PutObject",
                 "s3:ListBucket",
                 "s3:DeleteObject",
                 "s3:AbortMultipartUpload"
            ],
            "Resource": [
                 "arn:aws:s3:::<BUCKET-NAME>",
                 "arn:aws:s3:::<BUCKET-NAME>/*"
            ]
        }
    ]
}
```

- d. Choose Next.
- 4. Review and create the policy:
 - a. Provide a name for the policy (for example, "S3AccessPolicy").
 - b. Add an optional description to explain the purpose of the policy.
 - c. Review the policy and choose Create policy.
- 5. Open the IAM Console:
 - a. Navigate to the IAM dashboard.
 - b. In the navigation pane, choose **Roles**.
- 6. Create a Role:
 - a. Choose Create role.
 - b. Choose **Custom trust policy** as the type of trusted entity.

c. Paste the following JSON policy (replace <<u>ACCOUNT_ID</u>> with the actual account ID of Account A, <<u>ENVIRONMENT_NAME</u>> with the environment name of the RES deployment, and <<u>REGION</u>> with the AWS region RES is deployed to):

JSON

- d. Choose Next.
- 7. Attach Permissions Policies:
 - a. Search for and select the policy you created earlier.
 - b. Choose Next.
- 8. Tag, Review, and Create the Role:
 - a. Enter a role name (for example, "S3AccessRole").
 - b. Under Step 3, choose **Add Tag**, then enter the following key and value:
 - Key: res: Resource
 - Value: s3-bucket-iam-role
 - c. Review the role and choose **Create role**.
- 9. Use the IAM Role in RES:
 - a. Copy the IAM role ARN that you created.
 - b. Log into the RES console.

- c. In the left navigation pane, choose **S3 Bucket**.
- d. Choose Add Bucket and fill out the form with the cross-account S3 bucket ARN.
- e. Choose the **Advanced settings optional** dropdown.
- f. Enter the role ARN in the IAM role ARN field.
- q. Choose **Add Bucket**.

Step 2: Modify the bucket policy in Account B

- 1. Sign in to the AWS Management Console for Account B.
- 2. Open the S3 Console:
 - a. Navigate to the S3 dashboard.
 - b. Select the bucket you want to grant access to.
- 3. Edit the Bucket Policy:
 - a. Select the **Permissions** tab and choose **Bucket policy**.
 - b. Add the following policy to grant the IAM role from Account A access to the bucket (replace <<u>AccountA_ID</u>> with the actual account ID of Account A and <<u>BUCKET-NAME</u>> with the name of the S3 bucket):

JSON

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Principal": {
                "AWS": "arn:aws:iam::111122223333:role/S3AccessRole"
            },
            "Action": [
                "s3:GetObject",
                "s3:PutObject",
                "s3:ListBucket",
                "s3:DeleteObject",
                "s3:AbortMultipartUpload"
            ],
            "Resource": [
```

c. Choose Save.

Preventing data exfiltration in a private VPC

To prevent users from exfiltrating data from secure S3 buckets into their own S3 buckets in their account, you can attach a VPC endpoint to secure your private VPC. The following steps show how to create a VPC endpoint for the S3 service that supports access to S3 buckets within your account, as well as any additional accounts that have cross-account buckets.

- 1. Open the Amazon VPC Console:
 - a. Sign in to the AWS Management Console.
 - b. Open the Amazon VPC console at https://console.aws.amazon.com/vpc/.
- 2. Create a VPC Endpoint for S3:
 - a. In the left navigation pane, choose **Endpoints**.
 - b. Choose **Create Endpoint**.
 - c. For **Service category**, ensure that **AWS services** is selected.
 - d. In the **Service Name** field, enter com.amazonaws.<a hr
 - e. Select the S3 service from the list.
- 3. Configure Endpoint Settings:
 - a. For **VPC**, select the VPC where you want to create the endpoint.
 - b. For **Subnets**, select both the private subnets used for the VDI Subnets during deployment.
 - c. For **Enable DNS name**, ensure that the option is checked. This allows the private DNS hostname to be resolved to the endpoint network interfaces.
- 4. Configure the Policy to Restrict Access:
 - a. Under **Policy**, choose **Custom**.

b. In the policy editor, enter a policy that restricts access to resources within your account or a specific account. Here's an example policy (replace *mybucket* with your S3 bucket name and 111122223333 and 444455556666 with the appropriate AWS account IDs that you want to have access):

JSON

```
}
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Principal": "*",
            "Action": "s3:*",
            "Resource": [
                "arn:aws:s3:::mybucket",
                "arn:aws:s3:::mybucket/*"
            ],
            "Condition": {
                "StringEquals": {
                     "aws:PrincipalAccount": [
                         "111122223333",
                                           // Your Account ID
                         "444455556666"
                                           // Another Account ID
                    ]
                }
            }
        }
    ]
}
```

- 5. Create the Endpoint:
 - a. Review your settings.
 - b. Choose Create endpoint.
- 6. Verify the Endpoint:
 - a. Once the endpoint is created, navigate to the **Endpoints** section in the VPC console.
 - b. Select the newly created endpoint.
 - c. Verify that the **State** is **Available**.

By following these steps, you create a VPC endpoint that allows S3 access that is restricted to resources within your account or a specified account ID.

Troubleshooting

How to check if a bucket fails to mount on a VDI

If a bucket fails to mount on a VDI, there are a few locations where you can check for errors. Follow the steps below.

- Check the VDI Logs:
 - a. Log into the AWS Management Console.
 - b. Open the EC2 Console and navigate to **Instances**.
 - c. Select the VDI instance you launched.
 - d. Connect to the VDI via the Session Manager.
 - e. Run the following commands:

```
sudo su
cd ~/bootstrap/logs
```

Here, you'll find the bootstrap logs. The details of any failure will be located in the configure.log. {time} file.

Additionally, check the /etc/message log for more details.

- 2. Check Custom Credential Broker Lambda CloudWatch Logs:
 - a. Log into the AWS Management Console.
 - b. Open the CloudWatch Console and navigate to **Log groups**.
 - c. Search for the log group /aws/lambda/<stack-name>-vdc-custom-credentialbroker-lambda.
 - d. Examine the first available log group and locate any errors within the logs. These logs will contain details regarding potential issues providing temporary custom credentials for mounting S3 buckets.
- 3. Check Custom Credential Broker API Gateway CloudWatch Logs:
 - a. Log into the AWS Management Console.
 - b. Open the CloudWatch Console and navigate to **Log groups**.

- c. Search for the log group <stack-name>-vdc-custom-credential-brokerlambdavdccustomcredentialbrokerapigatewayaccesslogs<nonce>.
- d. Examine the first available log group and locate any errors within the logs. These logs will contain details regarding any requests and responses to the API Gateway for custom credentials needed to mount the S3 buckets.

How to edit a bucket's IAM role configuration after onboarding

- 1. Sign in to the AWS DynamoDB Console.
- 2. Select the Table:
 - a. In the left navigation pane, choose **Tables**.
 - b. Find and select <<u>stack-name</u>>.cluster-settings.
- 3. Scan the Table:
 - a. Choose **Explore table items**.
 - b. Ensure **Scan** is selected.
- 4. Add a Filter:
 - a. Choose **Filters** to open the filter entry section.
 - b. Set the filter to match your key-
 - Attribute: Enter the key.
 - Condition: Select Begins with.
 - Value: Enter shared-storage. <filesystem_id>.s3_bucket.iam_role_arn replacing <filesystem_id> with the value of the filesystem that needs to be modified.
- 5. Execute the Scan:

Choose Run to run the scan with the filter.

6. Check the value:

If the entry exists, ensure the value is correctly set with the right IAM role ARN.

If the entry does not exist:

a. Choose Create item.

- b. Enter the item details:
 - For the key attribute, enter sharedstorage. <filesystem_id>.s3_bucket.iam_role_arn.
 - Add the correct IAM role ARN.
- c. Choose **Save** to add the item.
- 7. Restart the VDI instances:

Reboot the instance to ensure the VDIs that are affected by the incorrect IAM role ARN are mounted again.

Enabling CloudTrail

To enable CloudTrail in your account using the CloudTrail console, follow the instructions provided in <u>Creating a trail with the CloudTrail console</u> in the *AWS CloudTrail User Guide*. CloudTrail will log the access to S3 buckets by recording the IAM role that accessed it. This can be linked back to an instance ID, which is linked to a project or user.

Use the product

This section offers guidance to users on using virtual desktops to collaborate with other users.

Topics

- SSH access
- Virtual desktops
- Shared desktops
- File browser

SSH access

To use SSH to access the bastion host:

- 1. From the RES menu, choose SSH access.
- 2. Follow the onscreen directions to use either SSH or PuTTY for access.

Virtual desktops

The virtual desktop interface (VDI) module allows users create and manage Windows or Linux virtual desktops on AWS. Users can launch Amazon EC2 instances with their favorite tools and application pre-installed and configured.

Supported operating systems

RES currently supports launching virtual desktops using the following operating systems:

- Amazon Linux 2 (x86 and ARM64)
- Ubuntu 22.04.03 (x86)
- RHEL 8 (x86), and 9 (x86)
- Windows 2019, 2022 (x86)

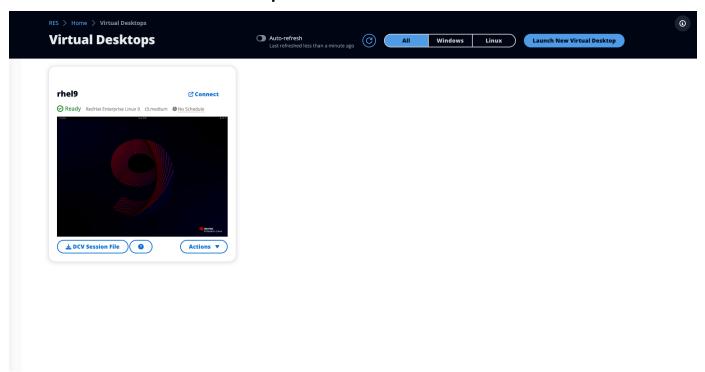
Topics

SSH access 161

- Launch a new desktop
- Access your desktop
- Control your desktop state
- Modify a virtual desktop
- Retrieve session information
- Schedule virtual desktops
- Virtual desktop interface autostop

Launch a new desktop

- 1. From the menu, choose My Virtual Desktops.
- 2. Choose Launch New Virtual Desktop.



- 3. Enter the details for your new desktop.
- 4. Choose Submit.

A new card with your desktop information appears instantly, and your desktop will be ready to use within 10-15 minutes. Startup time depends on the selected image. RES detects GPU instances and installs the relevant drivers.

Launch a new desktop 162

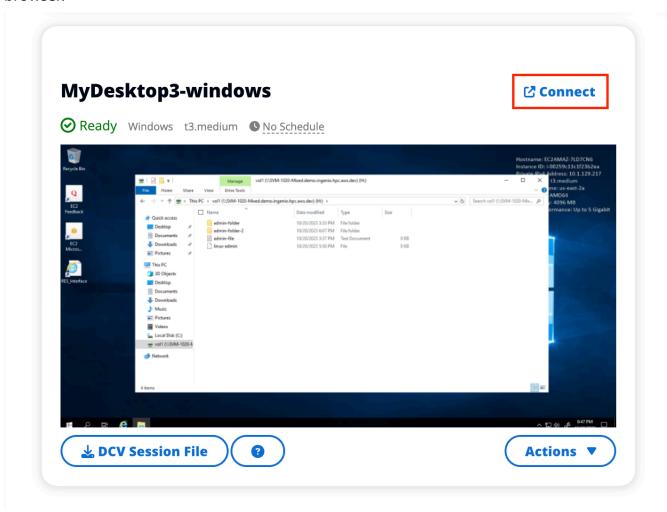
Access your desktop

To access a virtual desktop, choose the card for the desktop and connect using either the web or a DCV client.

Web connection

Accessing your desktop through the web browser is the easiest method of connection.

 Choose Connect, or choose the thumbnail to access your desktop directly through your browser.

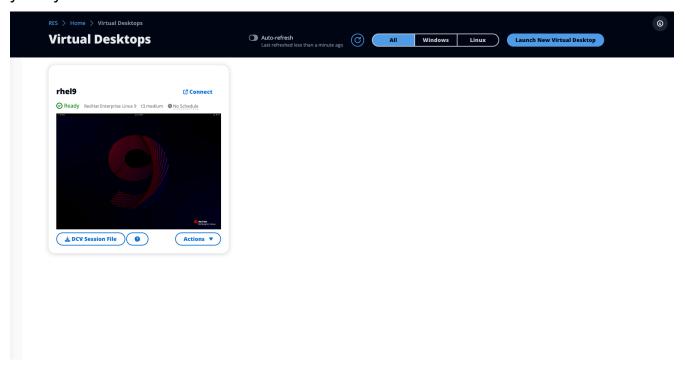


DCV connection

Accessing your desktop through a DCV client offers the best performance. To access via DCV:

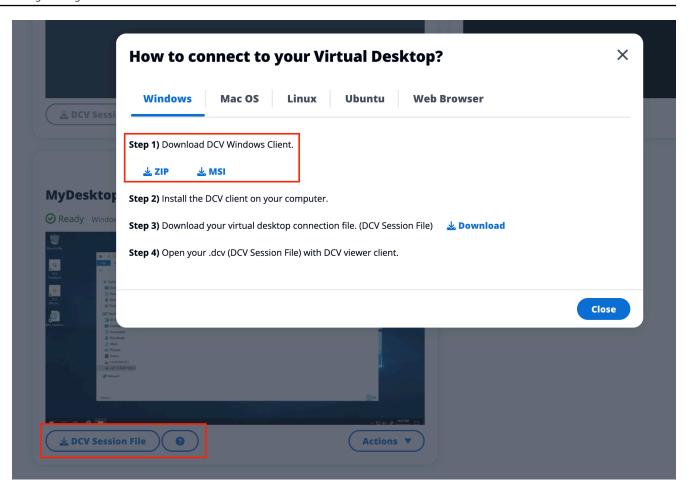
Access your desktop 163

1. Choose **DCV Session File** to download the .dcv file. You will need a DCV client installed on your system.



2. For installation instructions, choose the ? icon.

Access your desktop 164

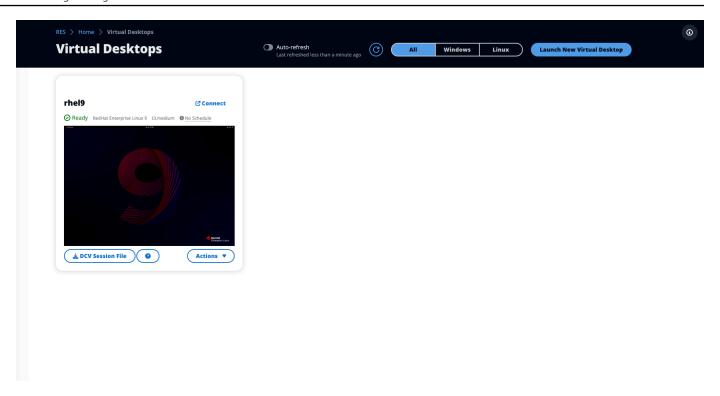


Control your desktop state

To control your desktop's state:

Choose Actions.

Control your desktop state 165



2. Choose **Virtual Desktop State**. You have four states to select from:

Stop

A stopped session will not suffer data loss, and you can restart a stopped session at any time.

Reboot

Reboots current session.

Terminate

Permanently ends a session. Terminating a session may cause data loss if you are using ephemeral storage. You should backup your data to the RES filesystem before terminating.

Hibernate

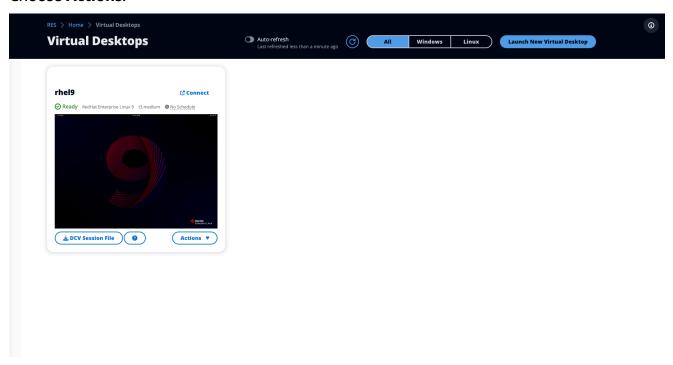
Your desktop state will be saved in memory. When you restart the desktop, your applications will resume but any remote connections may be lost. Not all instances support hibernation, and the option is only available if it was enabled during instance creation. To verify if your instance supports this state, see Hibernation prerequisites.

Control your desktop state 166

Modify a virtual desktop

You can update the hardware of your virtual desktop or change the session name.

- 1. Before making changes to the instance size, you must stop the session:
 - a. Choose Actions.



- b. Choose **Virtual Desktop State**.
- c. Choose **Stop**.



You cannot update the desktop size for hibernated sessions.

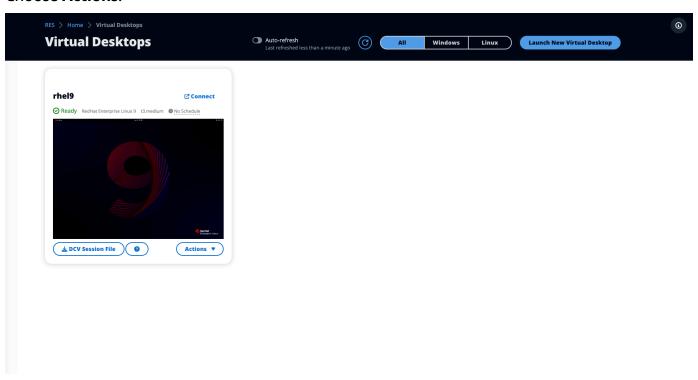
- 2. Once you have confirmed the desktop has stopped, choose **Actions** and then choose **Update Session**.
- 3. Change the session name or choose the desktop size you would like.
- 4. Choose Submit.
- 5. Once your instances updates, restart your desktop:
 - a. Choose Actions.

Modify a virtual desktop 167

- b. Choose Virtual Desktop State.
- c. Choose Start.

Retrieve session information

1. Choose Actions.



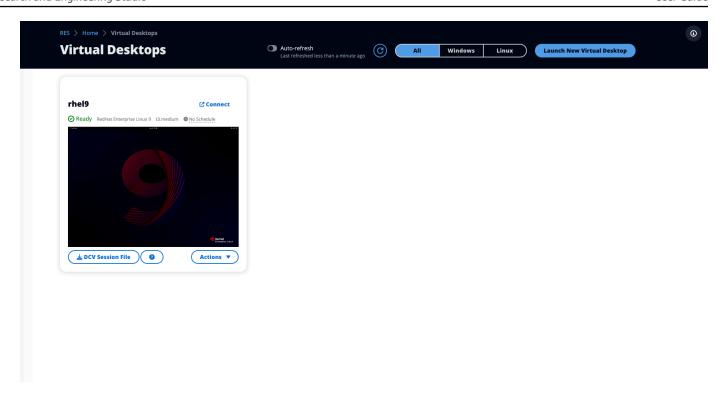
2. Choose Show Info.

Schedule virtual desktops

By default, virtual desktops do not have a schedule and will stay active until you stop or terminate the session. Desktops also stop if idle to prevent accidental stops. An idle state is determined by no active connection and CPU usage below 15% for at least 15 minutes. You can configure a schedule to automatically start and stop your desktop.

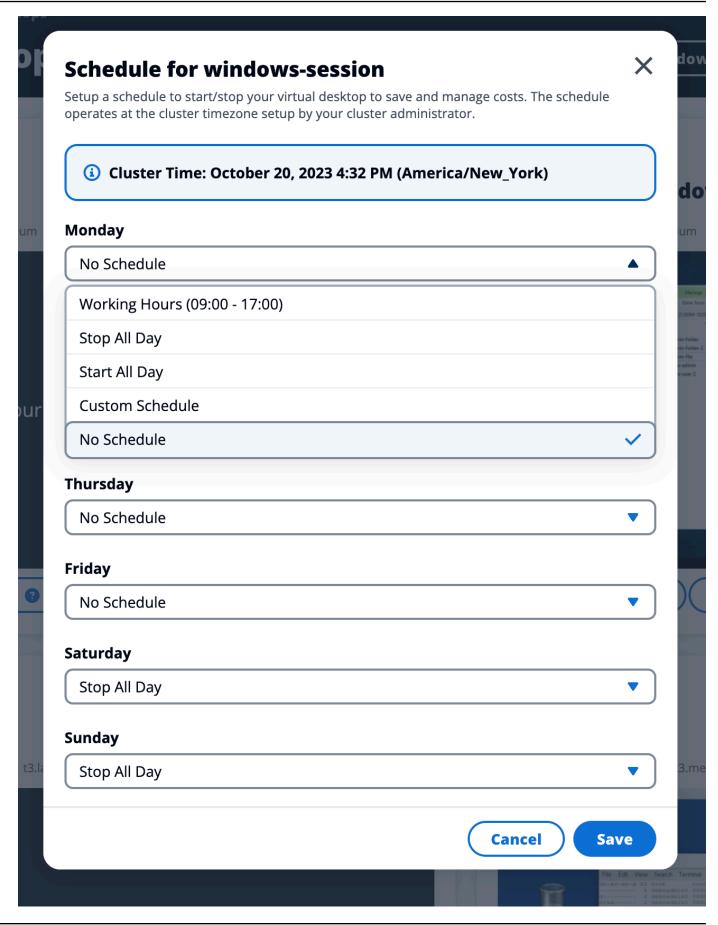
1. Choose Actions.

Retrieve session information 168



- 2. Choose **Schedule**.
- 3. Set your schedule for each day.
- 4. Choose **Save**.

Schedule virtual desktops 169



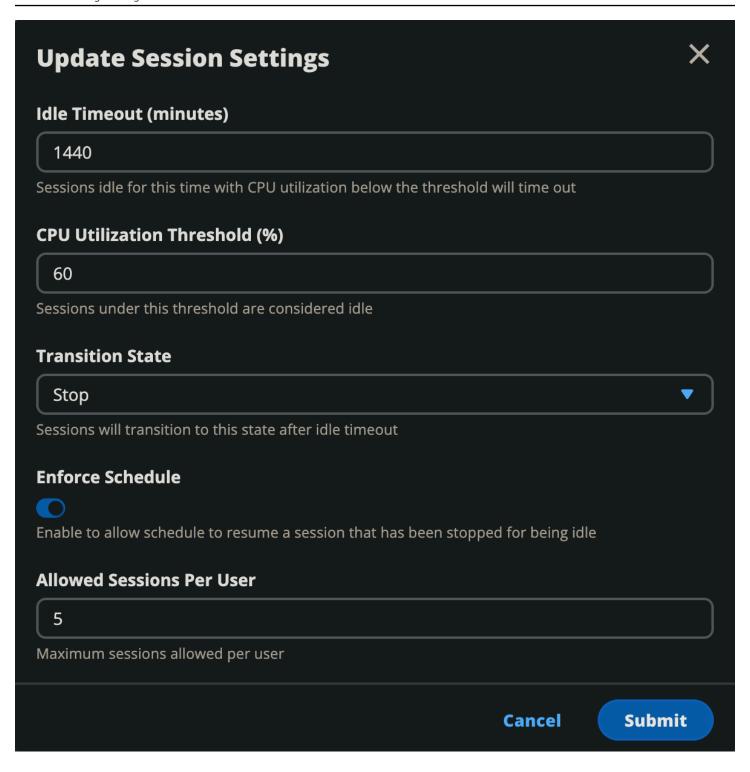
Schedule virtual desktops 170

Virtual desktop interface autostop

Administrators can configure settings to allow idle VDIs to be Stopped or Terminated. There are 4 configurable settings:

- 1. Idle Timeout: Sessions idle for this time with CPU utilization below the threshold will time out.
- 2. CPU Utilization Threshold: Sessions with no interaction and under this threshold are considered idle. If this is set to 0, then sessions will never be considered idle.
- 3. Transition State: After idle timeout, sessions will transition to this state (stopped or terminated).
- 4. Enforce Schedule: If selected, a session that has been stopped for being idle can be resumed by its daily schedule.

VDI autostop 171



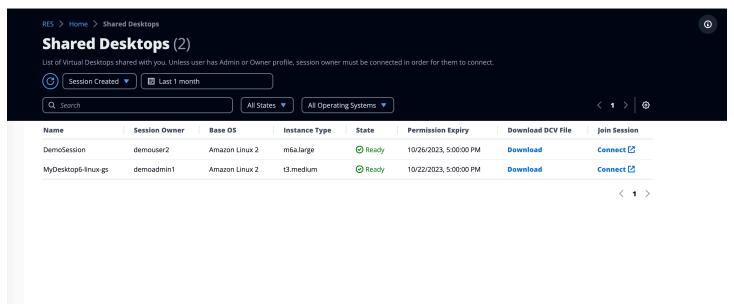
These settings are present on the **Desktop Settings** page under the **Server** tab. Once you update the settings according to your requirements, click on **Submit** to save the settings. New sessions will use the updated settings, but note that existing sessions will still use the settings which they had when they were launched.

VDI autostop 172

After they time out, sessions will either terminate or transition into the STOPPED_IDLE state based on their configuration. Users will have the ability to start STOPPED_IDLE sessions from the UI.

Shared desktops

On Shared Desktops, you can see the desktops that have been shared with you. In order to connect to a desktop, the session owner must be connected as well unless you are an Admin or Owner.



While sharing a session, you can configure permissions for your collaborators. For example, you can give read-only access to a teammate with whom you are collaborating.

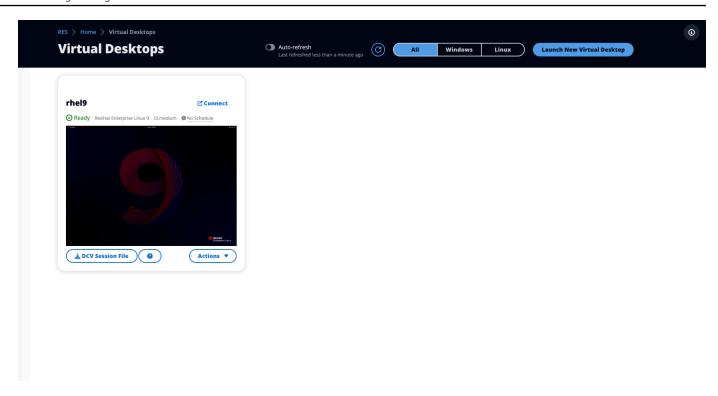
Topics

- · Share a desktop
- Access a shared desktop

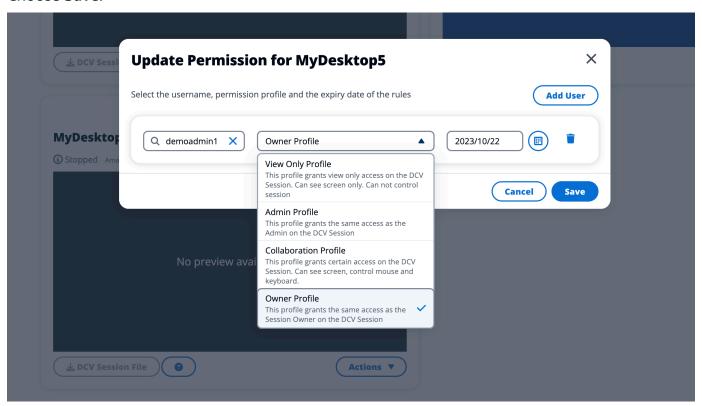
Share a desktop

1. From your desktop session, choose **Actions**.

Shared desktops 173



- 2. Select Session Permissions.
- 3. Select the user and permission level. You may also set an expiration time.
- 4. Choose Save.



Share a desktop 174

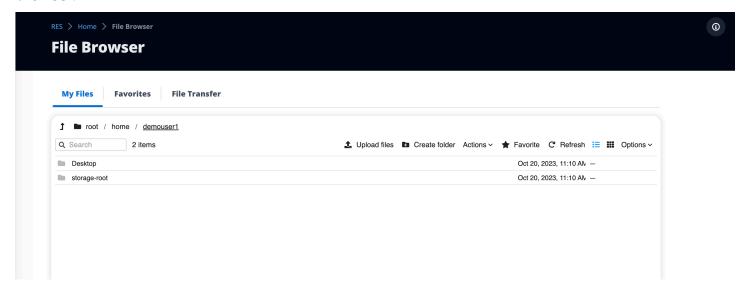
For more information on permissions, see the section called "Permission policy".

Access a shared desktop

From Shared Desktops, you can view the desktops shared with you and connect to an instance. You can join by either web browser or DCV. To connect, follow the directions in Access your desktop.

File browser

File browser allows you to access filesystems through the web portal. You can manage all available files you have permission to access on the underlying filesystem. Backend storage (Amazon EFS) is available for all Linux nodes. For Linux and Windows nodes, FSx for ONTAP is available. Updating files on your virtual desktop is the same as updating a file through the terminal or web-based file browser.



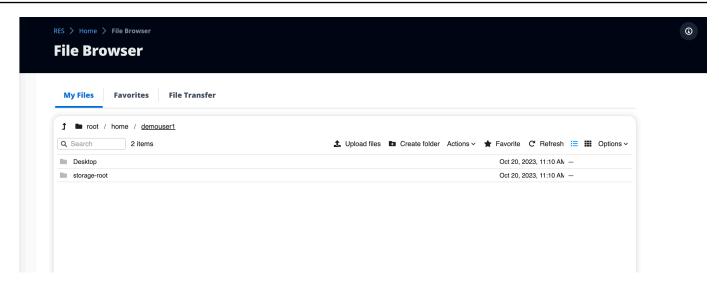
Topics

- Upload file(s)
- Delete file(s)
- Manage favorites
- Edit files
- Transfer files

Upload file(s)

1. Choose **Upload files**.

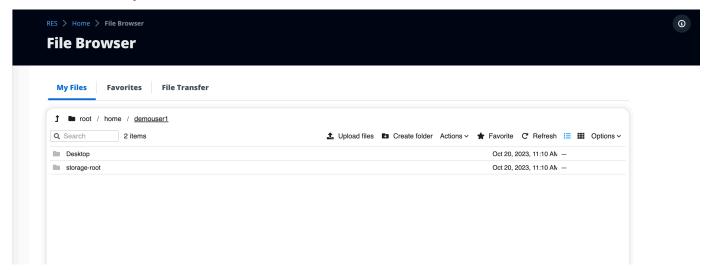
Access a shared desktop 175



- 2. Either drop files or browse for files to upload.
- 3. Choose Upload (n) files.

Delete file(s)

1. Select the file(s) you want to delete.



- 2. Choose Actions.
- 3. Select **Delete files**.

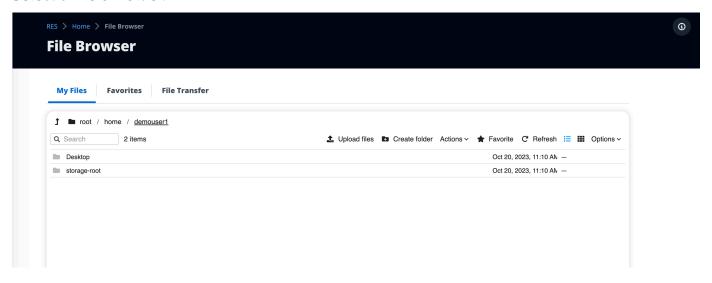
Alternatively, you can also right-click any file or folder and select **Delete files**.

Delete file(s) 176

Manage favorites

To pin important files and folders, you can add them to Favorites.

Select a file or folder.



Choose Favorite.

Alternatively, you can right-click any file or folder and select **Favorite**.



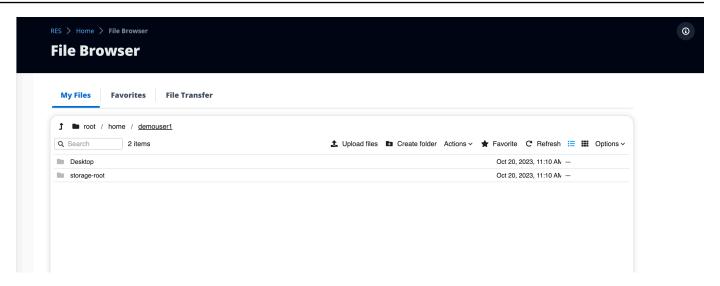
Favorites are stored to the local browser. If you change your browser or clear the cache, you will need to re-pin your favorites.

Edit files

You can edit the content of text-based files within the web portal.

1. Select the file you want to update. A modal will open with the file's content.

Manage favorites 177



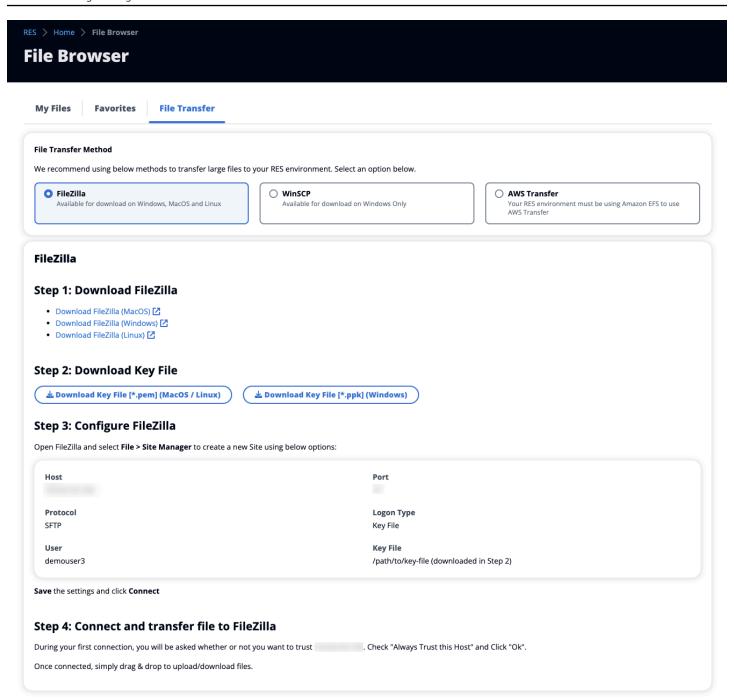
2. Make your updates and choose **Save**.

Transfer files

Use File Transfer to use external file transfer applications to transfer files. You can select from the following applications and follow the on-screen directions to transfer files.

- FileZilla (Windows, MacOS, Linux)
- WinSCP (Windows)
- AWS Transfer for FTP (Amazon EFS)

Transfer files 178



Transfer files 179

Troubleshooting

This section contains information about how to monitor the system and how to troubleshoot specific issues that may occur.

Topics

- General Debugging and Monitoring
- Issue RunBooks
- Known Issues

Detailed contents:

- · General Debugging and Monitoring
 - Useful log and event information sources
 - Log files on the environment Amazon EC2 instances
 - CloudFormation Stacks
 - System failures due to an issue and reflected by Amazon EC2 Auto Scaling Group Activity
 - Typical Amazon EC2 Console Appearance
 - Infrastructure hosts
 - Infrastructure hosts and virtual desktops
 - Hosts in a terminated state
 - Useful Active Directory (AD) related commands for reference
 - Windows DCV debugging
 - Find Amazon DCV Version Information
- Issue RunBooks
 - Installation issues
 - I want to set up custom domains after I install RES
 - AWS CloudFormation stack fails to create with message "WaitCondition received failed message. Error:States.TaskFailed"
 - Email notification not received after AWS CloudFormation stacks created successfully
 - Instances cycling or vdc-controller in failed state
 - Environment CloudFormation stack fails to delete due to dependent object error

- Error encountered for CIDR block parameter during environment creation
- CloudFormation stack creation failure during environment creation
- Creation of external resources (demo) stack fails with AdDomainAdminNode CREATE_FAILED
- Identity management issues
 - I am not authorized to perform iam:PassRole
 - I want to allow people outside of my AWS account to access my Research and Engineering Studio on AWS resources
 - When logging into the environment, I immediately return to the login page
 - "User not found" error when trying to log in
 - User added in Active Directory, but missing from RES
 - User unavailable when creating a session
 - Size limit exceeded error in CloudWatch cluster-manager log
- Storage
 - I created file system through RES but it doesn't mount on the VDI hosts
 - I onboarded a file system through RES but it doesn't mount on the VDI hosts
 - I am not able to read/write on from VDI hosts
 - Example permission handling use cases
 - I created Amazon FSx for NetApp ONTAP from RES but it did not join my domain
- Snapshots
 - A Snapshot has a status of Failed
 - A Snapshot fails to apply with logs indicating that the tables could not be imported.
- Infrastructure
 - Load balancer target groups without healthy instances
- Launching Virtual Desktops
 - A virtual desktop that was previously working is no longer able to connect successfully
 - I am only able to launch 5 virtual desktops
 - <u>Desktop Windows connect attempts fail with "The connection has been closed. Transport error"</u>
 - VDIs stuck in Provisioning state
 - VDIs get into Error state after launching
- Virtual Desktop Component

- Amazon EC2 instance is repeatedly showing terminated in the console
- vdc-controller instance is cycling due to failing to join AD / eVDI module shows Failed API
 Health Check
- Project does not appear in the pull down when editing the Software Stack to add it
- <u>cluster-manager Amazon CloudWatch log shows "<user-home-init> account not available</u>
 yet. waiting for user to be synced" (where the account is a user name)
- Windows desktop on login attempt says "Your account has been disabled. Please see your administrator"
- DHCP Options issues with external/customer AD configuration
- Firefox error MOZILLA PKIX ERROR REQUIRED TLS FEATURE MISSING
- · Env deletion
 - res-xxx-cluster stack in "DELETE_FAILED" state and cannot be deleted manually due to "Role is invalid or cannot be assumed" error
 - Collecting Logs
 - Downloading VDI Logs
 - Downloading logs from Linux EC2 instances
 - Downloading logs from Windows EC2 instances
 - Collecting ECS logs for the WaitCondition error
- Demo environment
 - Demo environment login error when handling authentication request to identity provider
 - Demo stack keycloak not working
- Known Issues 2024.x
 - Known Issues 2024.x
 - (2024.08) Virtual desktops fail to mount read/write Amazon S3 bucket with root bucket ARN and custom prefixing
 - (2024.06) Apply snapshot fails when the AD group name contains spaces
 - (2024.04-2024.04.02) Provided IAM Permission Boundary not attached to the VDI instances' role
 - (2024.04.02 and earlier) Windows NVIDIA instances in ap-southeast-2 (Sydney) fail to launch
 - (2024.04 and 2024.04.01) RES delete failure in GovCloud

- (2024.04 2024.04.02) Linux virtual desktop may be stuck in the "RESUMING" status on reboot
- (2024.04.02 and earlier) Fails to sync AD users whose SAMAccountName attribute includes capital letters or special characters
- (2024.04.02 and earlier) Private key for accessing the bastion host is invalid
- (2024.06 and earlier) Group members not synced to RES during AD sync
- (2024.06 and earlier) CVE-2024-6387, RegreSSHion, Security Vulnerability in RHEL9 and Ubuntu VDIs

General Debugging and Monitoring

This section contains information about where information can be found within RES.

- Useful log and event information sources
 - Log files on the environment Amazon EC2 instances
 - CloudFormation Stacks
 - System failures due to an issue and reflected by Amazon EC2 Auto Scaling Group Activity
- Typical Amazon EC2 Console Appearance
 - Infrastructure hosts
 - Infrastructure hosts and virtual desktops
 - Hosts in a terminated state
 - Useful Active Directory (AD) related commands for reference
- Windows DCV debugging
- Find Amazon DCV Version Information

Useful log and event information sources

There are various sources of information retained that can be referenced for troubleshooting and monitoring uses.

Log files on the environment Amazon EC2 instances

Log files exist on the Amazon EC2 instances in use by RES. The SSM Session Manager can be used to open a session to the instance for examining these files.

On infrastructure instances such as the cluster-manager and vdc-controller, application and other logs can be found at the following locations.

- /opt/idea/app/logs/application.log
- /root/bootstrap/logs/
- /var/log/
- /var/log/sssd/
- /var/log/messages
- /var/log/user-data.log
- /var/log/cloud-init.log
- /var/log/cloud-init-output.log

On a Linux virtual desktop, the following contain useful log files

- /var/log/dcv/
- /root/bootstrap/logs/userdata.log
- /var/log/messages

On Windows virtual desktop instances logs can be found at

- PS C:\ProgramData\nice\dcv\log
- PS C:\ProgramData\nice\DCVSessionManagerAgent\log

On Windows, some applications logging can be found at:

• PS C:\Program Files\NICE\DCV\Server\bin

On Windows, the NICE DCV certificate files can be found in:

C:\Windows\System32\config\systemprofile\AppData\Local\NICE\dcv\

Amazon CloudWatch Log Groups

The Amazon EC2 and AWS Lambda compute resources log information to Amazon CloudWatch Log Groups. The log entries within them can provide useful information when troubleshooting potential issues or for general information.

Those groups are named as follows:

- /aws/lambda/<envname>-/ lambda related
- /<envname>/
 - cluster-manager/ main infrastructure host
 - vdc/ virtual desktop related
 - dcv-broker/ desktop related
 - dcv-connection-gateway/ desktop related
 - controller/ main desktop controller host
 - dcv-session/ desktop session related

When examining log groups, it can be helpful to filter using upper and lower case strings such as the following. This will output only those messages containing the noted strings.

```
?"ERROR" ?"error"
```

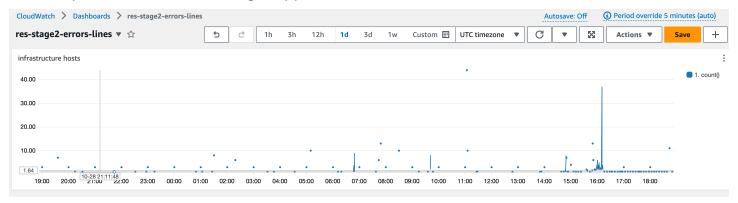
Another method of monitoring for issues is to create Amazon CloudWatch Dashboards that contain widgets that display the data of interest.

An example is to create a widget that counts the occurrence of the strings error and ERROR and graph them as lines. This method makes it easier to detect the occurrence of potential issues or trends indicating a pattern change has occurred.

The following is an example of that for the infrastructure hosts. To use this, concatenate the query lines and replace the <envname> and <region> attributes with the appropriate values.

```
"height": 6,
            "properties": {
                "query": "SOURCE '/<envname>/vdc/controller' |
                    SOURCE '/<envname>/cluster-manager' |
                    SOURCE '/<envname>/vdc/dcv-broker' |
                   SOURCE '/<envname>/vdc/dcv-connection-gateway' |
                    fields @timestamp, @message, @logStream, @log\n|
                    filter @message like /(?i)(error|ERROR)/\n|
                    sort @timestamp desc|
                    stats count() by bin(30s)",
                "region": "<region>",
                "title": "infrastructure hosts",
                "view": "timeSeries",
                "stacked": false
            }
        }
    ]
}
```

An example of the Dashboard might appear as follows:



CloudFormation Stacks

The CloudFormation stacks created during environment creation contain resources, event, and output information associated with the configuration of the environment.

For each of the stacks, the Events, Resources, and Outputs tab can be referred to for information about the stacks.

RES stacks:

- <envname>-bootstrap
- <envname>-cluster

- <envname>-metrics
- <envname>-directoryservice
- <envname>-identity-provider
- <envname>-shared-storage
- <envname>-cluster-manager
- <envname>-vdc
- <envname>-bastion-host

Demo Environment Stack (If you are deploying a demo environment and do not have these external resources available, you can use AWS High Performance Compute recipes to generate resources for a demo environment.)

- <envname>
- <envname>-Networking
- <envname>-DirectoryService
- <envname>-Storage
- <envname>-WindowsManagementHost

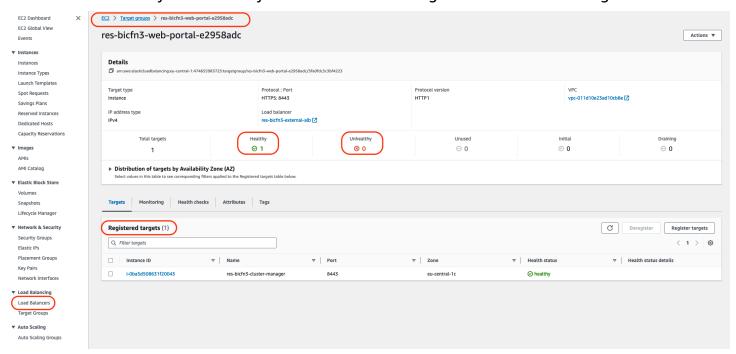
System failures due to an issue and reflected by Amazon EC2 Auto Scaling Group Activity

If the RES UIs indicate server errors, the cause may be an application software or other issue.

Each of the infrastructure Amazon EC2 instance autoscaling groups (ASGs) contains an Activity tab that can be useful for detecting scaling activity for the instances. If UI pages note any errors or are not accessible, check the Amazon EC2 console for multiple terminated instances and check the Auto Scaling Group Activity tab for the related ASG to determine if Amazon EC2 instances are cycling.

If so, use the related Amazon CloudWatch log group for the instance to determine if errors are being logged that might indicate the cause of the issue. It may also be possible to use the SSM Session console to open a session to a running instance of that type and examine the log files on the instance to determine a cause before the instance is marked as unhealthy and terminated by the ASG.

The ASG console may show activity similar to the following if this issue is occurring.

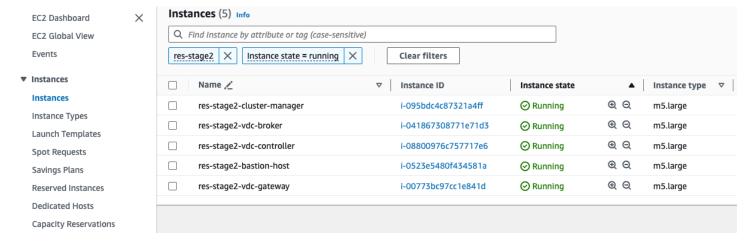


Typical Amazon EC2 Console Appearance

This section contains screenshots of the system operating in various states.

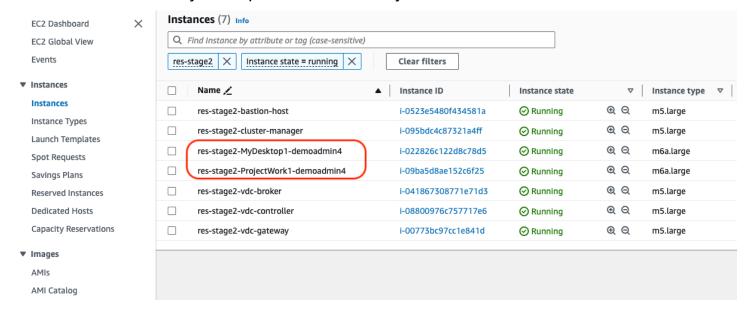
Infrastructure hosts

The Amazon EC2 console, when no desktops are running, typically looks similar to the following. The instances that are shown are the RES infrastructure Amazon EC2 hosts. The prefix in an instance name is the RES environment name.



Infrastructure hosts and virtual desktops

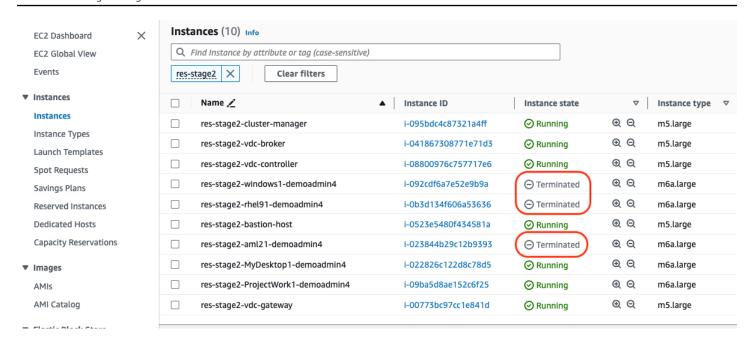
In the Amazon EC2 console, when virtual desktops are running, they appear similar to the following. In this case, the virtual desktops are noted in red. The suffix to the instance name is the user that created the desktop. The name in the center is the Session Name set at launch time and is either be the default "MyDesktop" or the name set by the user.



Hosts in a terminated state

When the Amazon EC2 console shows terminated instances, they are generally desktop hosts that have been terminated. If the console includes infrastructure hosts in a terminated state, particularly if there are multiple of the same type, that may indicate a system issue in progress.

The following image shows desktop instances that have been terminated.



Useful Active Directory (AD) related commands for reference

The following are examples of Idap related commands that can be entered on infrastructure hosts to view AD configuration related information. The domain and other parameters used should reflect those entered at environment creation time.

```
ldapsearch "(cn=AWS Delegated Add Workstations To Domain Users)" -x -h corp.res.com
   -b "DC=corp,DC=res,DC=com" -D "CN=Admin,OU=Users,OU=CORP,DC=corp,DC=res,DC=com"
   -w <password>
ldapsearch "(&(objectClass=group))" -x -h corp.res.com
   -b "DC=corp,DC=res,DC=com" -D "CN=Admin,OU=Users,OU=CORP,DC=corp,DC=res,DC=com"
   -w <password>
```

Windows DCV debugging

On a Windows desktop, you can list the session associated with it using the following:

```
PS C:\Windows\System32\config\systemprofile\AppData\Local\NICE\dcv> & 'C:\Program Files \NICE\DCV\Server\bin\dcv.exe'list-sessions
Session: 'a7953489-9dbf-492b-8135-7709dccc4cab' (owner:admin2 type:console name:windows1)
```

Windows DCV debugging 190

Find Amazon DCV Version Information

Amazon DCV is utilized for virtual desktop sessions. <u>AWS Amazon DCV</u>. The following examples show how to determine the version of the DCV software installed.

Linux

```
[root@ip-10-3-157-194 ~]# /usr/bin/dcv version

Amazon DCV 2023.0 (r14852)
Copyright (C) 2010-2023 NICE s.r.l.
All rights reserved.

This product is protected by copyright and licenses restricting use, copying, distribution, and decompilation.
```

Windows

```
PS C:\Windows\System32\config\systemprofile\AppData\Local\NICE\dcv> & 'C:\Program Files \NICE\DCV\Server\bin\dcv.exe' version

Amazon DCV 2023.0 (r15065)
Copyright (C) 2010-2023 NICE s.r.l.
All rights reserved.

This product is protected by copyright and licenses restricting use, copying, distribution, and decompilation.
```

Issue RunBooks

The following section contains issues that may occur, how to detect them, and suggestions on how to resolve the issue.

- Installation issues
 - I want to set up custom domains after I install RES
 - AWS CloudFormation stack fails to create with message "WaitCondition received failed message. Error:States.TaskFailed"
 - Email notification not received after AWS CloudFormation stacks created successfully
 - Instances cycling or vdc-controller in failed state

- Environment CloudFormation stack fails to delete due to dependent object error
- Error encountered for CIDR block parameter during environment creation
- CloudFormation stack creation failure during environment creation
- Creation of external resources (demo) stack fails with AdDomainAdminNode CREATE_FAILED
- Identity management issues
 - I am not authorized to perform iam:PassRole
 - I want to allow people outside of my AWS account to access my Research and Engineering Studio on AWS resources
 - When logging into the environment, I immediately return to the login page
 - "User not found" error when trying to log in
 - User added in Active Directory, but missing from RES
 - User unavailable when creating a session
 - Size limit exceeded error in CloudWatch cluster-manager log
- Storage
 - I created file system through RES but it doesn't mount on the VDI hosts
 - I onboarded a file system through RES but it doesn't mount on the VDI hosts
 - I am not able to read/write on from VDI hosts
 - Example permission handling use cases
 - I created Amazon FSx for NetApp ONTAP from RES but it did not join my domain
- Snapshots
 - A Snapshot has a status of Failed
 - A Snapshot fails to apply with logs indicating that the tables could not be imported.
- Infrastructure
 - Load balancer target groups without healthy instances
- Launching Virtual Desktops
 - A virtual desktop that was previously working is no longer able to connect successfully
 - I am only able to launch 5 virtual desktops
 - Desktop Windows connect attempts fail with "The connection has been closed. Transport error"
 - VDIs stuck in Provisioning state
 - VDIs get into Error state after launching

Issue RunBooks 192

- Virtual Desktop Component
 - Amazon EC2 instance is repeatedly showing terminated in the console
 - vdc-controller instance is cycling due to failing to join AD / eVDI module shows Failed API
 Health Check
 - · Project does not appear in the pull down when editing the Software Stack to add it
 - <u>cluster-manager Amazon CloudWatch log shows "<user-home-init> account not available yet.</u> waiting for user to be synced" (where the account is a user name)
 - Windows desktop on login attempt says "Your account has been disabled. Please see your administrator"
 - DHCP Options issues with external/customer AD configuration
 - Firefox error MOZILLA_PKIX_ERROR_REQUIRED_TLS_FEATURE_MISSING
- · Env deletion
 - res-xxx-cluster stack in "DELETE_FAILED" state and cannot be deleted manually due to "Role is invalid or cannot be assumed" error
 - Collecting Logs
 - Downloading VDI Logs
 - Downloading logs from Linux EC2 instances
 - Downloading logs from Windows EC2 instances
 - Collecting ECS logs for the WaitCondition error
- Demo environment
 - Demo environment login error when handling authentication request to identity provider
 - Demo stack keycloak not working

Installation issues

Topics

- I want to set up custom domains after I install RES
- AWS CloudFormation stack fails to create with message "WaitCondition received failed message.
 Error:States.TaskFailed"
- Email notification not received after AWS CloudFormation stacks created successfully
- Instances cycling or vdc-controller in failed state
- Environment CloudFormation stack fails to delete due to dependent object error

- Error encountered for CIDR block parameter during environment creation
- CloudFormation stack creation failure during environment creation
- Creation of external resources (demo) stack fails with AdDomainAdminNode CREATE_FAILED

I want to set up custom domains after I install RES



Note

Prerequisites: You must store Certificate and PrivateKey contents in a Secrets Manager secret before performing these steps.

Add certs to the web client

- Update the cert attached to the listener of the external-alb load balancer:
 - Navigate to the RES external load balancer in the AWS console under EC2 > Load **Balancing > Load Balancers.**
 - Search for the load balancer that follows the naming convention <env-name>external-alb.
 - Check the listeners attached to the load balancer.
 - Update the listener that has a Default SSL/TLS certificate attached with the new certificate details.
 - Save your changes.
- In the cluster-settings table: 2.
 - Find the cluster-settings table in DynamoDB -> Tables -> <env-name>.clustersettings.
 - b. Go to Explore Items and Filter by Attribute name "key", Type "string", condition "contains", and value "external_alb".
 - Set cluster.load_balancers.external_alb.certificates.provided to True. C.
 - d. Update the value of cluster.load_balancers.external_alb.certificates.custom_dns_name. This is the custom domain name for web user interface.

- e. Update the value of cluster.load_balancers.external_alb.certificates.acm_certificate_arn. This is the Amazon Resource Name (ARN) for the corresponding certificate stored in Amazon Certificate Manager (ACM).
- 3. Update the corresponding Route53 subdomain record you created for your web client to point to the DNS name of the external alb load balancer <env-name>-external-alb.
- 4. If SSO is already configured in the environment, re-configure SSO with the same inputs as you used initially from the **General Settings** > **Identity Provider** > **Single Sign On** > **Status** > **Edit** button in the RES web portal.

Add certs to the VDIs

- 1. Grant the RES application permission to perform a GetSecret operation on the secret by adding the following tags to the secrets:
 - res:EnvironmentName:<env-name>
 - res:ModuleName:virtual-desktop-controller
- 2. In the cluster-settings table:
 - a. Find the cluster-settings table in DynamoDB -> Tables -> <env-name>.cluster-settings.
 - b. Go to **Explore Items** and **Filter by Attribute** name "key", Type "string", condition "contains", and value "dcv_connection_gateway".
 - c. Set vdc.dcv_connection_gateway.certificate.provided to True.
 - d. Update the value of vdc.dcv_connection_gateway.certificate.custom_dns_name. This is the custom domain name for VDI access.
 - Update the value of vdc.dcv_connection_gateway.certificate.certificate_secret_arn. This is the ARN for the secret that holds the Certificate contents.
 - f. Update the value of vdc.dcv_connection_gateway.certificate.private_key_secret_arn. This is the ARN for the secret that holds the Private Key contents.
- 3. Update the launch template used for the gateway instance:

- Open the Auto Scaling group in the AWS Console under EC2 > Auto Scaling > Auto Scaling Groups.
- Select the gateway auto scaling group that corresponds to the RES environment. The name follows the naming convention <env-name>-vdc-gateway-asg.
- Find and open the Launch Template in the details section. c.
- d. Under **Details** > **Actions** > choose **Modify template** (Create new version).
- Scroll down to Advanced details. e.
- f. Scroll to the very bottom, to **User data**.
- Look for the words CERTIFICATE_SECRET_ARN and PRIVATE_KEY_SECRET_ARN. g. Update these values with the ARNs given to the secrets that hold the Certificate (see step 2.c) and Private Key (see step 2.d) contents.
- Ensure the Auto Scaling group is configured to use the recently created version of the launch template (from the Auto Scaling group page).
- Update the corresponding Route53 subdomain record you created for your virtual desktops to point to the DNS name of the external nlb load balancer: <env-name>-external-nlb.
- Terminate the existing dcv-gateway instance: <env-name>-vdc-gateway and wait for a new one to spin up.

........

AWS CloudFormation stack fails to create with message "WaitCondition received failed message. Error:States.TaskFailed"

To identify the issue, examine the Amazon CloudWatch log group named <stack-name>-InstallerTasksCreateTaskDefCreateContainerLogGroup<nonce>-<nonce>. If there are multiple log groups with the same name, examine the first available. An error message within the logs will provide more information on the issue.



Note

Confirm that the parameter values do not have spaces.

Installation issues 196

Email notification not received after AWS CloudFormation stacks created successfully

If an email invitation was not received after the AWS CloudFormation stacks were created successfully, verify the following:

- 1. Confirm the email address parameter was entered correctly.
 - If the email address is incorrect or cannot be accessed, delete and redeploy the Research and Engineering Studio environment.
- 2. Check Amazon EC2 console for evidence of cycling instances.
 - If there are Amazon EC2 instances with the <envname> prefix that appear as terminated and then are replaced with a new instance, there may be an issue with the network or Active Directory configuration.
- If you deployed the AWS High Performance Compute recipes to create your external resources, confirm that the VPC, private and public subnets, and other selected parameters were created by the stack.
 - If any of the parameters are incorrect, you may need to delete and redeploy the RES environment. For more information, see *Uninstall the product*.
- 4. If you deployed the product with your own external resources, confirm the networking and Active Directory match the expected configuration.
 - Confirming that infrastructure instances successfully joined the Active Directory is critical. Try the steps in the section called "Instances cycling or vdc-controller in failed state" to resolve the issue.

•••••

Instances cycling or vdc-controller in failed state

The most probable cause of this issue is the inability of resource(s) to connect or join the Active Directory.

To verify the issue:

 From the command line, start a session with SSM on the running instance of the vdccontroller.

- 2. Run sudo su -.
- Run systemctl status sssd.

If the status is inactive, failed, or you see errors in the logs, then the instance was unable to join Active Directory.

```
[root@ip-
                          ]# systemctl status sssd
  sssd.service - System Security Services Daemon
Loaded: loaded (/usr/lib/systemd/system/sssd.service; enabled; vendor preset: disabled)
   Active: active (running) since Tue 2023-11-14 12:12:19 UTC; 1 weeks 0 days ago
 Main PID: 31248 (sssd)
   CGroup: /system.slice/sssd.service
              31248 /usr/sbin/sssd -i --logger=files
              -31249 /usr/libexec/sssd/sssd_be --domain corp.res.com --uid 0 --gid 0 --logger=files
              31251 /usr/libexec/sssd/sssd_nss --uid 0 --gid 0 --logger=files
              -31252 /usr/libexec/sssd/sssd pam --uid 0 --gid 0 --logger=files
Nov 21 15:27:19 ip-10-3-144-194.ec2.internal sssd_be[31249]: GSSAPI client step 1
Nov 21 15:27:19 ip-10-3-144-194.ec2.internal sssd be[31249]: GSSAPI client step 2
Nov 21 15:42:19 ip-10-3-144-194.ec2.internal sssd be[31249]: GSSAPI client step
                                                                                             Might see errors
Nov 21 15:42:19 ip-10-3-144-194.ec2.internal sssd_be[31249]: GSSAPI client step
                                                                                              highlighted in
Nov 21 15:42:19 ip-10-3-144-194.ec2.internal sssd_be[31249]: GSSAPI client step
Nov 21 15:42:19 ip-10-3-144-194.ec2.internal sssd_be[31249]: GSSAPI client step Nov 21 15:57:19 ip-10-3-144-194.ec2.internal sssd_be[31249]: GSSAPI client step
                                                                                                RED here
Nov 21 15:57:19 ip-10-3-144-194.ec2.internal sssd be[31249]: GSSAPI client step
Nov 21 15:57:19 ip-10-3-144-194.ec2.internal sssd_be[31249]: GSSAPI client step
Nov 21 15:57:19 ip-10-3-144-194.ec2.internal sssd_be[31249]: GSSAPI client step
```

SSM error log

To solve the issue:

 From the same command line instance, run cat /root/bootstrap/logs/userdata.log to investigate the logs.

The issue could have one of three possible root causes.

Root cause 1: Incorrect Idap connection details entered

Review the logs. If you see the following repeated multiple times, the instance was unable to join the Active Directory.

```
+ local AD_AUTHORIZATION_ENTRY=
+ [[ -z '' ]]
+ [[ 0 -le 180 ]]
+ local SLEEP_TIME=34
+ log_info '(0 of 180) waiting for AD authorization, retrying in 34 seconds ...'
++ date '+%Y-%m-%d %H:%M:%S,%3N'
```

```
+ echo '[2024-01-16 22:02:19,802] [INFO] (0 of 180) waiting for AD authorization,
  retrying in 34 seconds ...'
[2024-01-16 22:02:19,802] [INFO] (0 of 180) waiting for AD authorization, retrying in
  34 seconds ...
+ sleep 34
+ (( ATTEMPT_COUNT++ ))
```

- 1. Verify the parameter values for the following were entered correctly during RES stack creation.
 - directoryservice.ldap_connection_uri
 - directoryservice.ldap_base
 - directoryservice.users.ou
 - directoryservice.groups.ou
 - directoryservice.sudoers.ou
 - directoryservice.computers.ou
 - · directoryservice.name
- 2. Update any incorrect values in the DynamoDB table. The table is found in the DynamoDB console under **Tables**. The table name should be <stack name>.cluster-settings.
- 3. After you update the table, delete the cluster-manager and vdc-controller currently running the environment instances. Auto scaling will start new instances using the latest values from the DynamoDB table.

Root cause 2: Incorrect ServiceAccount username entered

If the logs return Insufficient permissions to modify computer account, the ServiceAccount name entered during stack creation could be incorrect.

- 1. From the AWS Console, open Secrets Manager.
- 2. Search for directoryserviceServiceAccountUsername. The secret should be <stack name>-directoryservice-ServiceAccountUsername.
- 3. Open the secret to view the details page. Under **Secret Value**, choose **Retrieve secret value** and choose **Plaintext**.
- 4. If the value was updated, delete the currently running cluster-manager and vdc-controller instances of the environment. Auto scaling will start new instances using the latest value from Secrets Manager.

Root cause 3: Incorrect ServiceAccount password entered

If the logs display Invalid credentials, the ServiceAccount password entered during stack creation might be incorrect.

- 1. From the AWS Console, open Secrets Manager.
- Search for directoryserviceServiceAccountPassword. The secret should be <stack name>-directoryservice-ServiceAccountPassword.
- 3. Open the secret to view the details page. Under **Secret Value**, choose **Retrieve secret value** and choose **Plaintext**.
- 4. If you forgot the password or you are unsure if the entered password is correct, you can reset the password in Active Directory and Secrets Manager.
 - a. To reset the password in AWS Managed Microsoft AD:
 - i. Open the AWS Console and go to AWS Directory Service.
 - ii. Select the **Directory ID** for your RES directory, and choose **Actions**.
 - iii. Select **Reset user password**.
 - iv. Enter the ServiceAccount username.
 - v. Enter a new password, and choose **Reset password**.
 - b. To reset the password in Secrets Manager:
 - i. Open the AWS Console and go to Secrets Manager.
 - ii. Search for directoryserviceServiceAccountPassword. The secret should be
 <stack name>-directoryservice-ServiceAccountPassword.
 - iii. Open the secret to view the details page. Under **Secret Value**, choose **Retrieve secret value** then choose **Plaintext**.
 - iv. Choose Edit.
 - v. Set a new password for the ServiceAccount user and choose **Save**.
- 5. If you updated the value, delete the currently running cluster-manager and vdc-controller instances of the environment. Auto scaling will start new instances using the latest value.

•••••

Environment CloudFormation stack fails to delete due to dependent object error

If the deletion of the <env-name>-vdc CloudFormation stack fails due to a dependent object error such as the vdcdcvhostsecuritygroup, this could be due to an Amazon EC2 instance that was launched into a RES-created subnet or security group using the AWS Console.

To resolve the issue, find and terminate all Amazon EC2 instances launched in this manner. You can then resume the environment deletion.

.....

Error encountered for CIDR block parameter during environment creation

When creating an environment, an error appears for the CIDR block parameter with a response status of [FAILED].

Example of error:

```
Failed to update cluster prefix list:

An error occurred (InvalidParameterValue) when calling the

ModifyManagedPrefixList operation:

The specified CIDR (52.94.133.132/24) is not valid. For example, specify a CIDR in the following form: 10.0.0.0/16.
```

To resolve the issue, the expected format is x.x.x.0/24 or x.x.x.0/32.

.....

CloudFormation stack creation failure during environment creation

Creating an environment involves a series of resource creation operations. In some Regions, a capacity issue may occur which causes a CloudFormation stack creation to fail.

If this occurs, delete the environment and retry the creation. Alternatively, you can retry the creation in a different Region.

.....

Creation of external resources (demo) stack fails with AdDomainAdminNode CREATE FAILED

If the demo environment stack creation fails with the following error, it may be due to Amazon EC2 patching occurring unexpectedly during the provisioning after instance launch.

AdDomainAdminNode CREATE_FAILED Failed to receive 1 resource signal(s) within the specified duration

To determine the cause of failure:

- In the SSM State Manager, check if patching is configured and if it is configured for all instances.
- 2. In the SSM RunCommand/Automation execution history, check if execution of a patchingrelated SSM document coincides with an instance launch.
- In the log files for the environment's Amazon EC2 instances, review the local instance logging to determine if the instance rebooted during provisioning.

If the issue was caused by patching, delay patching for the RES instances at least 15 minutes post-launch.

•••••

Identity management issues

Most issues with single sign-on (SSO) and identity management occur due to misconfiguration. For information on setting up your SSO configuration, see:

- the section called "Setting up SSO with IAM Identity Center"
- the section called "Configuring your identity provider for SSO"

To troubleshoot other issues related to identity management, see the following troubleshooting topics:

Topics

- I am not authorized to perform iam:PassRole
- I want to allow people outside of my AWS account to access my Research and Engineering Studio on AWS resources
- When logging into the environment, I immediately return to the login page
- "User not found" error when trying to log in
- User added in Active Directory, but missing from RES

- User unavailable when creating a session
- Size limit exceeded error in CloudWatch cluster-manager log

.....

I am not authorized to perform iam:PassRole

If you receive an error that you're not authorized to perform the iam:PassRole action, your policies must be updated to allow you to pass a role to RES.

Some AWS services allow you to pass an existing role to that service instead of creating a new service role or service-linked role. To do this, you must have permissions to pass the role to the service.

The following example error occurs when an IAM user named marymajor tries to use the console to perform an action in RES. However, the action requires the service to have permissions that are granted by a service role. Mary does not have permissions to pass the role to the service.

User: arn:aws:iam::123456789012:user/marymajor is not authorized to perform:
 iam:PassRole

In this case, Mary's policies must be updated to allow her to perform the iam:PassRole action. If you need help, contact your AWS administrator. Your administrator is the person who provided you with your sign-in credentials.

.....

I want to allow people outside of my AWS account to access my Research and Engineering Studio on AWS resources

You can create a role that users in other accounts or people outside of your organization can use to access your resources. You can specify who is trusted to assume the role. For services that support resource-based policies or access control lists (ACLs), you can use those policies to grant people access to your resources.

To learn more, consult the following:

• To learn how to provide access to your resources across AWS accounts that you own, see Providing access to an IAM user in another AWS account that you own in the IAM User Guide.

- To learn how to provide access to your resources to third-party AWS accounts, see Providing access to AWS accounts owned by third parties in the IAM User Guide.
- To learn how to provide access through identity federation, see <u>Providing access to externally</u> authenticated users (identity federation) in the *IAM User Guide*.
- To learn the difference between using roles and resource-based policies for cross-account access, see How IAM roles differ from resource-based policies in the IAM User Guide.

.....

When logging into the environment, I immediately return to the login page

This issue occurs when your SSO integration is misconfigured. To determine the issue, check the controller instance logs and review the configuration settings for errors.

To check the logs:

- 1. Open the CloudWatch console.
- 2. From **Log groups**, find the group named /<<u>environment-name</u>>/cluster-manager.
- 3. Open the log group to search for any errors in the log streams.

To check the configuration settings:

- 1. Open the <u>DynamoDB console</u>
- 2. From **Tables**, find the table named <environment-name</pre>.cluster-settings.
- 3. Open the table and choose **Explore table items**.
- 4. Expand the filters section, and enter the following variables:
 - Attribute name key
 - Condition contains
 - Value sso
- 5. Choose Run.
- 6. In the returned string, verify that the SSO configuration values are correct. If they are incorrect, change the value of the sso_enabled key to **False**.

Edit item

You can add, remove, or edit the attributes of an item. You can nest attributes inside other attributes up to 32 levels deep. Learn more 🔀



Return to the RES user interface to reconfigure the SSO.

.....

"User not found" error when trying to log in

If a user receives the error "User not found" when they try to log in to the RES interface, and the user is present in Active Directory:

- If the user is not present in RES and you recently added the user to AD
 - It is possible that the user is not yet synced to RES. RES syncs hourly, so you may need to
 wait and check that the user was added after the next sync. To sync immediately, follow the
 steps in User added in Active Directory, but missing from RES.
- If the user is present in RES:
 - 1. Ensure the attribute mapping is configured correctly. For more information, see <u>Configuring</u> your identity provider for single sign-on (SSO).
 - 2. Ensure that the SAML subject and SAML email both map to the user's email address.

•••••

User added in Active Directory, but missing from RES

If you have added a user to the Active Directory but they are missing in RES, the AD sync needs to be triggered. The AD sync is performed hourly by a Lambda function that imports AD entries to the RES environment. Occasionally, there is a delay until the next sync process runs after you add new users or groups. You can initiate the sync manually from the Amazon Simple Queue Service.

Initiate the sync process manually:

- 1. Open the Amazon SQS console.
- 2. From Queues, select <environment-name>-cluster-manager-tasks.fifo.
- Choose Send and receive messages.
- 4. For Message body, enter:

```
{ "name": "adsync.sync-from-ad", "payload": {} }
```

- 5. For Message group ID, enter: adsync.sync-from-ad
- For Message deduplication ID, enter a random alpha-numeric string. This entry must be different from all calls made within the previous five minutes or the request will be ignored.

.....

User unavailable when creating a session

If you are an administrator creating a session, but find that a user who is in the Active Directory is not available when creating a session, the user may need to log in for the first time. Sessions can only be created for active users. Active users must log into the environment at least once.

.....

Size limit exceeded error in CloudWatch cluster-manager log

```
2023-10-31T18:03:12.942-07:00 ldap.SIZELIMIT_EXCEEDED: {'msgtype': 100, 'msgid': 11, 'result': 4, 'desc': 'Size limit exceeded', 'ctrls': []}
```

If you receive this error in the CloudWatch cluster-manager log, the ldap search may have returned too many user records. To fix this issue, increase your IDP's ldap search result limit.

•••••

Storage

Topics

- I created file system through RES but it doesn't mount on the VDI hosts
- I onboarded a file system through RES but it doesn't mount on the VDI hosts

- I am not able to read/write on from VDI hosts
- I created Amazon FSx for NetApp ONTAP from RES but it did not join my domain

.....

I created file system through RES but it doesn't mount on the VDI hosts

The file systems need to be in the "Available" state before they can be mounted by VDI hosts. Follow the steps below to validate the file system is in the required state.

Amazon EFS

- 1. Go to the Amazon EFS console.
- 2. Check that the File system state is Available.
- 3. If the file system state is not **Available**, wait before launching VDI hosts.

Amazon FSx ONTAP

- 1. Go to the Amazon FSx console.
- 2. Check that the **Status** is **Available**.
- 3. If Status is not Available, wait before launching VDI hosts.

.....

I onboarded a file system through RES but it doesn't mount on the VDI hosts

The file systems onboarded on RES should have the required security group rules configured to allow VDI hosts to mount the file systems. As these file systems are created externally to RES, RES doesn't manage the associated security group rules.

The security group associated with the onboarded file systems should allow the following inbound traffic:

- NFS traffic (port: 2049) from the linux VDC hosts
- SMB traffic (port: 445) from the windows VDC hosts

•••••

I am not able to read/write on from VDI hosts

ONTAP supports UNIX, NTFS and MIXED security style for the volumes. The security styles determine the type of permissions ONTAP uses to control data access and what client type can modify these permissions.

For example, if a volume uses UNIX security style, SMB clients can still access data (provided that they properly authenticate and authorize) due to the multi-protocol nature of ONTAP. However, ONTAP uses UNIX permissions that only UNIX clients can modify using native tools.

Example permission handling use cases

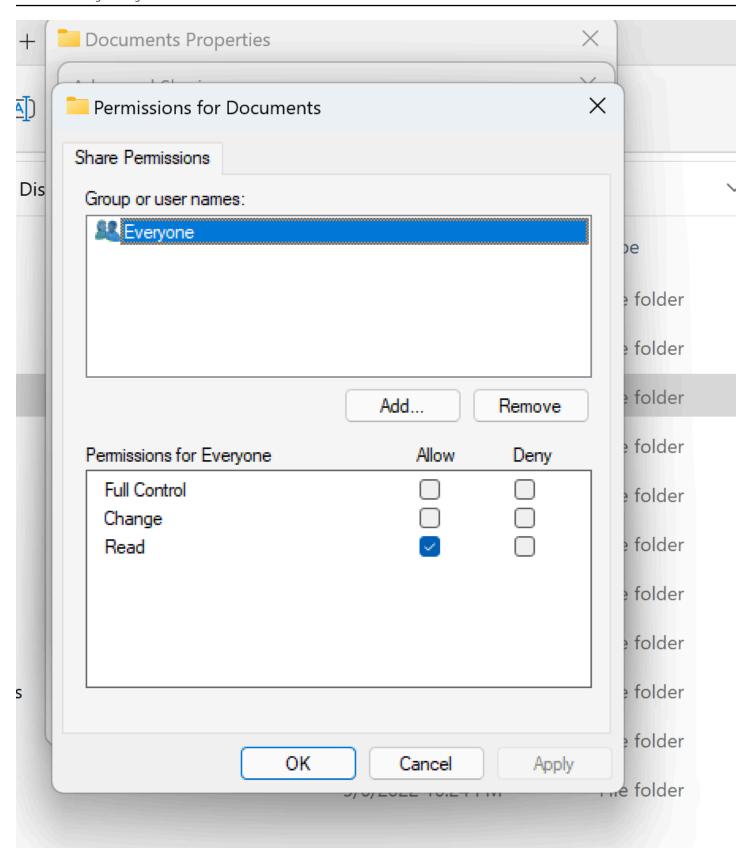
Using UNIX style volume with Linux workloads

Permissions can be configured by the sudoer for other users. For example, the following would give all members of <group-ID> full read/write permissions on the //project-name> directory:

```
sudo chown root:<group-ID> //sudo chmod 770 ///sudo chmod 770 ///
```

Using NTFS style volume with Linux and Windows workloads

Share Permissions can be configured using the share properties of a particular folder. For example, given a user user_01 and a folder myfolder, you can set permissions of Full Control, Change, or Read to Allow or Deny:



If the volume is going to be used by both Linux and Windows clients we need to set up a name mapping on SVM that will associate any Linux user name to the same user name with the NetBIOS domain name format of domain\username. This is needed to translate between Linux and Windows users. For reference see Enabling multiprotocol workloads with Amazon FSx for NetApp ONTAP.

.....

I created Amazon FSx for NetApp ONTAP from RES but it did not join my domain

Currently, when you create Amazon FSx for NetApp ONTAP from the RES console, the file system gets provisioned but it does not join the domain. To join the created ONTAP file system SVM to your domain, see <u>Joining SVMs to a Microsoft Active Directory</u> and follow the steps on the <u>Amazon FSx console</u>. Make sure required <u>permissions are delegated to the Amazon FSx Service Account</u> in AD. Once the SVM joins the domain successfully, go to SVM **Summary > Endpoints > SMB DNS name**, and copy the DNS name because you will need it later.

After it is joined to the domain, edit the SMB DNS config key in the cluster settings DynamoDB table:

- 1. Go to the Amazon DynamoDB console.
- 2. Choose **Tables**, then choose <stack-name>-cluster-settings.
- 3. Under **Explore table items**, expand **Filters**, and enter the following filter:
 - Attribute name key
 - Condition Equal to
 - Value-shared-storage.<file-system-name>.fsx_netapp_ontap.svm.smb_dns
- 4. Select the returned item, then **Actions**, **Edit item**.
- 5. Update the **value** with the SMB DNS name you copied earlier.
- 6. Choose Save and close.

In addition, ensure the security group associated with the file system allows traffic as recommended in <u>File System Access Control with Amazon VPC</u>. New VDI hosts that use the file system will now be able to mount the domain joined SVM and file system.

Alternatively, you may onboard an existing file system which is already joined to your domain using RES Onboard File System capability- from **Environment Management** choose **File Systems**, **Onboard File System**.

•••••

Snapshots

Topics

- A Snapshot has a status of Failed
- A Snapshot fails to apply with logs indicating that the tables could not be imported.

•••••

A Snapshot has a status of Failed

On the RES Snapshots page, if a snapshot has a status of Failed, the cause can be determined by going to the Amazon CloudWatch log group for the cluster-manager for the time that the error occurred.

```
[2023-11-19 03:39:20,208] [INFO] [snapshots-service] creating snapshot in S3 Bucket:
   asdf at path s31
[2023-11-19 03:39:20,381] [ERROR] [snapshots-service] An error occurred while
   creating the snapshot: An error occurred (TableNotFoundException)
   when calling the UpdateContinuousBackups operation:
   Table not found: res-demo.accounts.sequence-config
```

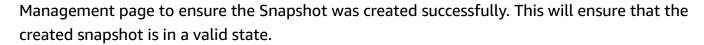
.....

A Snapshot fails to apply with logs indicating that the tables could not be imported.

If a snapshot taken from a previous env fails to apply in a new env, look into the CloudWatch logs for cluster-manager to identify the issue. If the issue mentions that the required tables cloud not be imported, verify that the snapshot is in a valid state.

- 1. Download the metadata.json file and verify that the ExportStatus for the various tables has status COMPLETED. Ensure the various tables have the ExportManifest field set. If you do not find the above fields set, the snapshot is in an invalid state and cannot be used with the apply snapshot functionality.
- 2. After initiating a snapshot creation, ensure that the Snapshot status turns to COMPLETED in RES. The Snapshot creation process takes up to 5 to 10 minutes. Reload or revisit the Snapshot

Snapshots 211



.....

Infrastructure

Topics

• Load balancer target groups without healthy instances

.....

Load balancer target groups without healthy instances

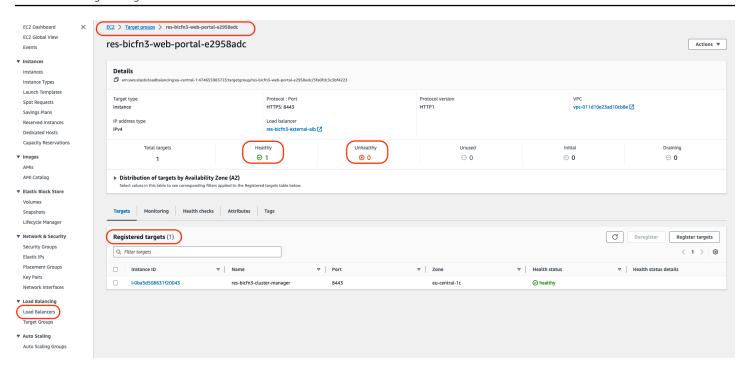
If issues such as server error messages are appearing in the UI or desktop sessions cannot connect, that may indicate an issue in the infrastructure Amazon EC2 instances.

The methods to determine the source of the issue are to first check the Amazon EC2 console for any Amazon EC2 instances that appear to be repeatedly terminating and being replaced by new instances. If that is the case, checking the Amazon CloudWatch logs may determine the cause.

Another method is check the load balancers in the system. An indication that there may be system issues is if any load balancers, found on the Amazon EC2 console, do not show any healthy instances registered.

An example of a normal appearance is shown here:

Infrastructure 212



If the Healthy entry is 0, that indicates that no Amazon EC2 instance is available to process requests.

If the Unhealthy entry is non-0, that indicates that an Amazon EC2 instance may be cycling. This can be due to the installed applications software not passing health checks.

If both Healthy and Unhealthy entries are 0, that indicates a potential network misconfiguration. For example, the public and private subnets might not have corresponding AZs. If this condition occurs, there may be additional text on the console indicating that network state exists.

•••••

Launching Virtual Desktops

Topics

- A virtual desktop that was previously working is no longer able to connect successfully
- I am only able to launch 5 virtual desktops
- Desktop Windows connect attempts fail with "The connection has been closed. Transport error"
- VDIs stuck in Provisioning state
- VDIs get into Error state after launching

•••••

A virtual desktop that was previously working is no longer able to connect successfully

If a desktop connection closes or you can no longer connect to it, the issue may be due to the underlying Amazon EC2 instance failing or the Amazon EC2 instance may have been terminated or stopped outside of the RES environment. The Admin UI status may continue to show a ready state but attempts to connect to it fail.

The Amazon EC2 Console should be used to determine if the instance has been terminated or stopped. If stopped, try starting it again. If the state is terminated, another desktop will have to be created. Any data that was stored on the user home directory should still be available when the new instance starts.

If the instance that failed previously still appears on the Admin UI, it may need to be terminated using the Admin UI.

•••••

I am only able to launch 5 virtual desktops

The default limit for the number of virtual desktops that a user can launch is 5. This can be changed by an admin using the Admin UI as follows:

- Go to Desktop Settings.
- Select the **Server** tab.
- In the **DCV Session** panel, click on the edit icon on the right.
- Change the value in **Allowed Sessions Per User** to the desired new value.
- Choose Submit.
- Refresh the page to confirm that the new setting is in place.

.....

Desktop Windows connect attempts fail with "The connection has been closed. Transport error"

If a Windows desktop connection fails with the UI error "The connection has been closed. Transport error", the cause can be due to an issue in the DCV server software related to certificate creation on the Windows instance.

The Amazon CloudWatch log group <envname>/vdc/dcv-connection-gateway may log the connection attempt error with messages similar to the following:

```
Nov 24 20:24:27.631 DEBUG HTTP:Splicer Connection{id=9}:
Websocket{session_id="1291e75f-7816-48d9-bbb2-7371b3b911cd"}:
Resolver lookup{client_ip=Some(52.94.36.19)
session_id="1291e75f-7816-48d9-bbb2-7371b3b911cd"
protocol_type=WebSocket extension_data=None}:NoStrictCertVerification:
Additional stack certificate (0): [s/n: 0E9E9C4DE7194B37687DC4D2C0F5E94AF0DD57E]

Nov 24 20:25:15.384 INFO HTTP:Splicer Connection{id=21}:Websocket{
session_id="d1d35954-f29d-4b3f-8c23-6a53303ebc3f"}:
Connection initiated error: unreachable, server io error Custom {
kind: InvalidData, error:
General("Invalid certificate: certificate has expired (code: 10)") }

Nov 24 20:25:15.384 WARN HTTP:Splicer Connection{id=21}:
Websocket{session_id="d1d35954-f29d-4b3f-8c23-6a53303ebc3f"}:
Error in websocket connection: Server unreachable: Server error: IO error: unexpected error: Invalid certificate: certificate has expired (code: 10)
```

If this occurs, a resolution may be to use the SSM Session Manager to open a connection to the Windows instance and remove the following 2 certificate related files:

The files should be automatically recreated and a subsequent connection attempt may be successful.

If this method resolves the issue and if new launches of Windows desktops produce the same error, use the Create Software Stack function to create a new Windows software stack of the fixed instance with the regenerated certificate files. That may produce a Windows software stack that can be used for successful launches and connections.

.....

VDIs stuck in Provisioning state

If a desktop launch remains in the provisioning state in the Admin UI, this may be due to several reasons.

To determine the cause, examine the log files on the desktop instance and look for errors that might be causing the issue. This document contains a list of log files and Amazon CloudWatch log groups that contain relevant information in the section labeled *Useful log and event information sources*.

The following are potential causes of this issue.

The AMI id used has been registered as a software-stack but is not supported by RES.

The bootstrap provisioning script failed to complete because the Amazon Machine Image (AMI) does not have the expected configuration or tooling required. The log files on the instance, such as /root/bootstrap/logs/ on a Linux instance, may contain useful information regarding this. AMIs ids taken from the AWS Marketplace may not work for RES desktop instances. They require testing to confirm if they are supported.

 User data scripts are not executed when the Windows virtual desktop instance is launched from a custom AMI.

By default, user data scripts run one time when an Amazon EC2 instance is launched. If you create an AMI from an existing virtual desktop instance, then register a software stack with the AMI and try to launch another virtual desktop with this software stack, user data scripts will not run on the new virtual desktop instance.

To fix the issue, open a PowerShell command window as Administrator on the **original** virtual desktop instance you used to create the AMI, and run the following command:

C:\ProgramData\Amazon\EC2-Windows\Launch\Scripts\InitializeInstance.ps1 -Schedule

Then create a new AMI from the instance. You can use the new AMI to register software stacks and launch new virtual desktops afterwards. Note that you may also run the same command on the instance that remains in the provisioning state and reboot the instance to fix the virtual desktop session, but you will run into the same issue again when launching another virtual desktop from the misconfigured AMI.

.....

VDIs get into Error state after launching

Possible issue 1: The home filesystem has directory for the user with different POSIX permissions.

This could be the issue you are facing if the following scenarios are true:

- 1. The RES Version deployed is 2024.01 or higher.
- 2. During deployment of the RES stack the attribute for EnableLdapIDMapping was set to True.
- 3. The home filesystem specified during the RES stack deployment was used in version prior to RES 2024.01 or was used in a previous environment with EnableLdapIDMapping set to False.

Resolution steps: Delete the user directories in the filesystem.

- 1. SSM to the cluster-manager host.
- 2. cd /home.
- 3. 1s should list directories with directory names that match usernames, such as admin1, admin2.. and so on.
- 4. Delete the directories, sudo rm -r 'dir_name'. **Do not delete the ssm-user and ec2-user directories**.
- 5. If the users are already synced to the new env, delete the user's from the user's DDB table (except clusteradmin).
- 6. Initiate AD sync run sudo /opt/idea/python/3.9.16/bin/resctl ldap sync-from-ad in the cluster-manager Amazon EC2.
- 7. Reboot the VDI instance in the Error state from the RES webpage. Validate that the VDI transitions into the Ready state in around 20 minutes.

•••••

Virtual Desktop Component

Topics

· Amazon EC2 instance is repeatedly showing terminated in the console

- vdc-controller instance is cycling due to failing to join AD / eVDI module shows Failed API Health
 Check
- · Project does not appear in the pull down when editing the Software Stack to add it
- <u>cluster-manager Amazon CloudWatch log shows "<user-home-init> account not available yet.</u> waiting for user to be synced" (where the account is a user name)
- Windows desktop on login attempt says "Your account has been disabled. Please see your administrator"
- DHCP Options issues with external/customer AD configuration
- Firefox error MOZILLA_PKIX_ERROR_REQUIRED_TLS_FEATURE_MISSING

•••••

Amazon EC2 instance is repeatedly showing terminated in the console

If an infrastructure instance is repeatedly showing as terminated in the Amazon EC2 console, the cause may be related to its configuration and depend on the infrastructure instance type. The following are methods to determine the cause.

If the vdc-controller instance shows repeated terminated states in the Amazon EC2 console, this can be due to an incorrect Secret tag. Secrets that are maintained by RES have tags that are used as a part of the IAM access control policies attached to the infrastructure Amazon EC2 instances. If the vdc-controller is cycling and the following error appears in the CloudWatch log group, the cause may be that a secret has not been tagged correctly. Note that the secret needs to be tagged with the following:

```
{
    "res:EnvironmentName": "<envname>" # e.g. "res-demo"
    "res:ModuleName": "virtual-desktop-controller"
}
```

The Amazon CloudWatch log message for this error will appear similar to the following:

```
An error occurred (AccessDeniedException) when calling the GetSecretValue operation: User: arn:aws:sts::160215750999:assumed-role/<envname>-vdc-gateway-role-us-east-1/i-043f76a2677f373d0 is not authorized to perform: secretsmanager:GetSecretValue on resource: arn:aws:secretsmanager:us-east-1:160215750999:secret:Certificate-res-bi-Certs-5W9SPUXF08IB-F1sNRv
```

because no identity-based policy allows the secretsmanager:GetSecretValue action

Check the tags on the Amazon EC2 instance and confirm that they match the above list.

•••••

vdc-controller instance is cycling due to failing to join AD / eVDI module shows Failed API Health Check

If the eVDI module is failing it's health check, it will show the following in the Environment Status section.

Modules

Environment modules and status



Module	Module ID	Version	Туре	Status	API Health Check	Module Sets
Global Settings	global-settings	-	Config	⊘ Deployed	O Not Applicable	-
Cluster	cluster	2023.10b1	Stack	⊘ Deployed	O Not Applicable	• default
Metrics & Monitoring	metrics	2023.10b1	Stack	⊘ Deployed	O Not Applicable	 default
Directory Service	directoryservice	2023.10b1	Stack	⊘ Deployed	O Not Applicable	 default
Identity Provider	identity-provider	2023.10b1	Stack	⊘ Deployed	O Not Applicable	 default
Analytics	analytics	2023.10b1	Stack	⊘ Deployed	O Not Applicable	• default
Shared Storage	shared-storage	2023.10b1	Stack	⊘ Deployed	O Not Applicable	 default
Cluster Manager	cluster-manager	2023.10b1	Арр	⊘ Deployed	⊘ Healthy	• default
eVDI	vdc	2023.10b1	③ Арр	⊘ Deployed	⊗ Failed	 default
Bastion Host	bastion-host	2023.10b1	Stack	⊘ Deployed	O Not Applicable	 default

In this case, the general path for debugging is to look into the **cluster-manager** <u>CloudWatch</u> logs. (Look for the log group named <env-name>/cluster-manager.)

Possible issues:

• If the logs contain the text Insufficient permissions, make sure the ServiceAccount username given when the res stack was created is spelled correctly.

Example log line:

```
Insufficient permissions to modify computer account:
   CN=IDEA-586BD25043,OU=Computers,OU=RES,OU=CORP,DC=corp,DC=res,DC=com:
   000020E7: AtrErr: DSID-03153943, #1: 0: 000020E7: DSID-03153943, problem 1005
   (CONSTRAINT_ATT_TYPE), data 0, Att 90008 (userAccountControl):len 4 >> 432 ms -
   request will be retried in 30 seconds
```

- You can access the ServiceAccount Username provided during RES deployment from the
 <u>SecretsManager console</u>. Find the corresponding secret in Secrets manager and choose
 Retrieve Plain text. If the Username is incorrect, choose Edit to update the secret value.
 Terminate the current cluster-manager and vdc-controller instances. The new instances will come up in a stable state.
- The username must be "ServiceAccount" if you are utilizing the resources created by the
 provided <u>external resources stack</u>. If the DisableADJoin parameter was set to False during
 your deployment of RES, ensure the "ServiceAccount" user has permissions to create **Computer**objects in the AD.
- If the username used was correct, but the logs contain the text Invalid credentials, then the password you entered might be **wrong** or have **expired**.

Example log line:

```
{'msgtype': 97, 'msgid': 1, 'result': 49, 'desc': 'Invalid credentials', 'ctrls': [],
'info': '80090308: LdapErr: DSID-0C090569, comment: AcceptSecurityContext error,
data 532, v4563'}
```

- You can read the password you entered during env creation by accessing the secret that stores the password in the <u>Secrets Manager console</u>. Select the secret (for example, <env_name>directoryserviceServiceAccountPassword) and choose **Retrieve plain** text.
- If the password in the secret is incorrect, choose **Edit** to update its value in the secret.

 Terminate the current cluster-manager and vdc-controller instances. The new instances will use the updated password and come up in a stable state.
- If the password is correct, it could be that the password has expired in the connected Active
 Directory. You'll have to first reset the password in the Active Directory and then update the
 secret. You can reset the user's password in the Active Directory from the <u>Directory Service</u>
 console:
 - 1. Choose the appropriate Directory ID

- 2. Choose **Actions**, **Reset user password** then fill out the form with the username (for example, "ServiceAccount") and the new password.
- 3. If the newly set password is different from the previous password, update the password in the corresponding Secret Manager secret (for example, <env_name>directoryserviceServiceAccountPassword.
- 4. Terminate the current cluster-manager and vdc-controller instances. The new instances will come up in a stable state.

.....

Project does not appear in the pull down when editing the Software Stack to add it

This issue may be related to the following issue associated with syncing the user account with AD. If this issue appears, check the cluster-manager Amazon CloudWatch log group for the error "<user-home-init> account not available yet. waiting for user to be synced" to determine if the cause is the same or related.

.....

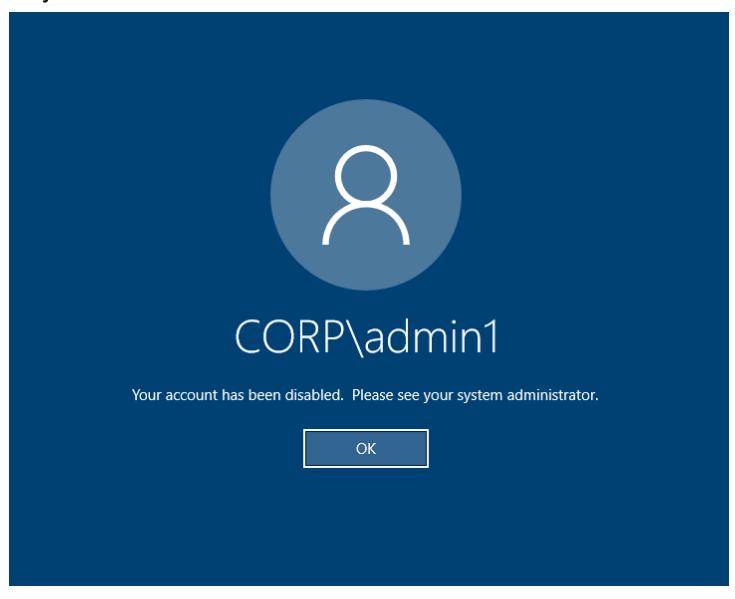
cluster-manager Amazon CloudWatch log shows "<user-home-init> account not available yet. waiting for user to be synced" (where the account is a user name)

The SQS subscriber is busy and stuck in an infinite loop because it cannot get to the user account. This code is triggered when trying to create a home filesystem for a user during user sync.

The reason it is not able to get to the user account may be that RES was not configured correctly for the AD in use. An example might be that the ServiceAccountCredentialsSecretArn parameter used at BI/RES environment creation was not the correct value.

.....

Windows desktop on login attempt says "Your account has been disabled. Please see your administrator"



If the user is unable to log back in to a locked screen, this may indicate that the user has been disabled in the AD configured for RES after having successfully signed on via SSO.

The SSO	login sł	าould f	fail if	the u	ser ac	count	has	been	disable	d in	AD.

DHCP Options issues with external/customer AD configuration

If you encounter an error stating "The connection has been closed. Transport error" with Windows virtual desktops when using RES with your own Active Directory, check the dcv-connection-gateway Amazon CloudWatch log for something similar to the following:

```
Oct 28 00:12:30.626 INFO HTTP:Splicer Connection{id=263}:
Websocket{session_id="96cffa6e-cf2e-410f-9eea-6ae8478dc08a"}: Connection initiated
error: unreachable, server io error Custom { kind: Uncategorized, error: "failed to
lookup address information: Name or service not known" }

Oct 28 00:12:30.626 WARN HTTP:Splicer Connection{id=263}:
Websocket{session_id="96cffa6e-cf2e-410f-9eea-6ae8478dc08a"}: Error in websocket
connection: Server unreachable: Server error: IO error: failed to lookup address
information: Name or service not known

Oct 28 00:12:30.627 DEBUG HTTP:Splicer Connection{id=263}: ConnectionGuard dropped
```

If you are using an AD domain controller for your DHCP Options for your own VPC, you need to:

- 1. Add AmazonProvidedDNS to the two domain controller IPs.
- 2. Set the domain name to ec2.internal.

A example is shown here. Without this configuration, the Windows desktop will give you **Transport error**, because RES/DCV looks for ip-10-0-x-xx.ec2.internal hostname.

Domain name

Domain name servers

1 10.0.2.168, 10.0.3.228,

AmazonProvidedDNS

Firefox error MOZILLA_PKIX_ERROR_REQUIRED_TLS_FEATURE_MISSING

When you use the Firefox web browser, you might encounter the error message type MOZILLA_PKIX_ERROR_REQUIRED_TLS_FEATURE_MISSING when you attempt to connect to a virtual desktop.

The cause is that the RES web server is set up with TLS + Stapling On but is not responding with Stapling Validation (see https://support.mozilla.org/en-US/questions/1372483.

You can fix this by following the instructions at: https://really-simple-ssl.com/ mozilla_pkix_error_required_tls_feature_missing.

.....

Env deletion

Topics

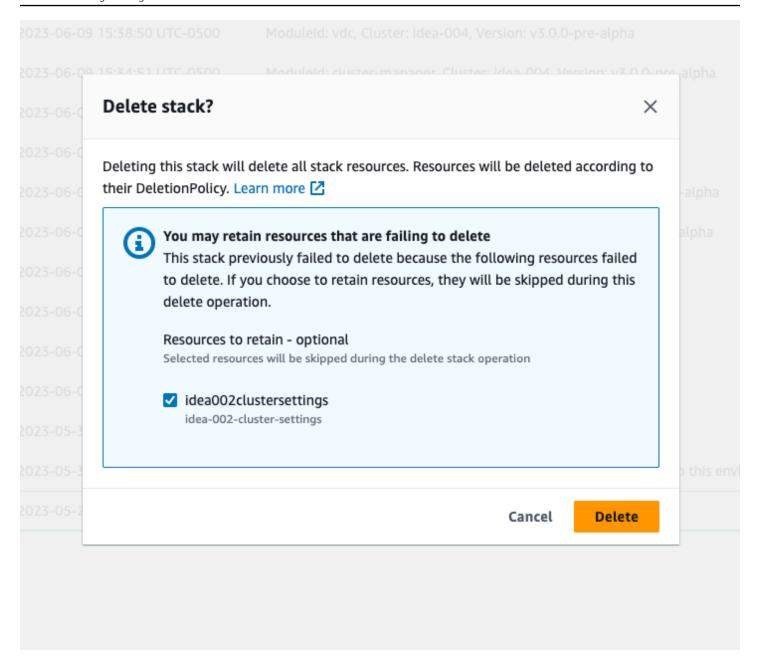
- res-xxx-cluster stack in "DELETE_FAILED" state and cannot be deleted manually due to "Role is invalid or cannot be assumed" error
- Collecting Logs
- Downloading VDI Logs
- Downloading logs from Linux EC2 instances
- Downloading logs from Windows EC2 instances
- Collecting ECS logs for the WaitCondition error

.....

res-xxx-cluster stack in "DELETE_FAILED" state and cannot be deleted manually due to "Role is invalid or cannot be assumed" error

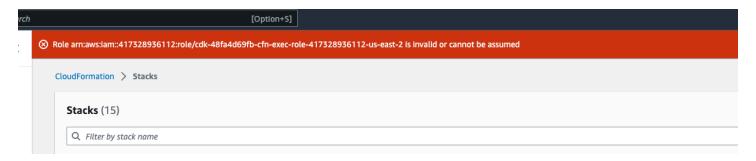
If you notice that the "res-xxx-cluster" stack is in "DELETE_FAILED" state and cannot be deleted manually, you can perform the following steps to delete it.

If you see the stack in a "DELETE_FAILED" state, first try to manually delete it. It may pop up a dialog confirming Delete Stack. Choose **Delete**.



Sometimes, even if you delete all the required stack resources, you may still see the message to select resources to retain. In that case, select all the resources as the "resources to retain" and choose **Delete**.

You may see an error that looks like Role: arn:aws:iam::... is Invalid or cannot be assumed

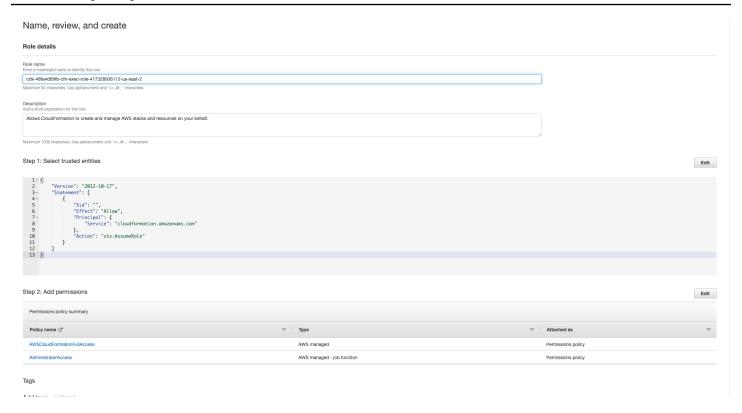


This means that the role required to delete the stack got deleted first before the stack. To get around this, copy the name of the role. Go to IAM console and create a role with that name using the parameters as shown here, which are:

- For Trusted entity type choose AWS service.
- For **Use case**, under Use cases for other AWS services choose CloudFormation.



Choose **Next**. Make sure you give the role 'AWSCloudFormationFullAccess' and 'AdministratorAccess' permissions. Your review page should look like this:



Then go back to the CloudFormation console and delete the stack. You should now be able to delete it since you created the role. Finally, go to IAM console and delete the role you created.

.....

Collecting Logs

Logging into an EC2 instance from the EC2 console

- Follow these instructions to login to your Linux EC2 instance.
- Follow <u>these instructions</u> to login to your Windows EC2 instance. Then open Windows PowerShell for running any commands.

Collecting Infrastructure host logs

- 1. Cluster-manager: Get logs for the cluster manager from the following places and attach them to the ticket.
 - a. All the logs from the CloudWatch log group <env-name>/cluster-manager.
 - b. All the logs under the /root/bootstrap/logs directory on the <env-name>-cluster-manager EC2 instance. Follow the instructions linked to from "Logging into an EC2 instance from the EC2 console" at the beginning of this section to login to your instance.

- 2. Vdc-controller: Get the logs for the vdc-controller from the following places and attach them to the ticket.
 - a. All the logs from the CloudWatch log group <env-name>/vdc-controller.
 - b. All the logs under the /root/bootstrap/logs directory on the <env-name>-vdc-controller EC2 instance. Follow the instructions linked to from "Logging into an EC2 instance from the EC2 console" at the beginning of this section to login to your instance.

One of the ways to get the logs easily is to follow the instructions in the <u>Downloading logs from</u> <u>Linux EC2 instances</u> section. The module name would be the instance name.

Collecting VDI logs

Identify the corresponding Amazon EC2 instance

If a user launched a VDI with session name VDI1, the corresponding name of the instance on the Amazon EC2 console would be <env-name>-VDI1-<user name>.

Collect Linux VDI logs

Log in to the corresponding Amazon EC2 instance from the Amazon EC2 console by following the instructions linked to in "Logging into an EC2 instance from the EC2 console" at the beginning of this section. Get all the logs under the /root/bootstrap/logs and /var/log/dcv/directories on the VDI Amazon EC2 instance.

One of the ways to get the logs would be to upload them to s3 and then download them from there. For that, you can follow these steps to get all the logs from one directory and then upload them:

1. Follow these steps to copy the dcv logs under the /root/bootstrap/logs directory:

```
sudo su -
cd /root/bootstrap
mkdir -p logs/dcv_logs
cp -r /var/log/dcv/* logs/dcv_logs/
```

2. Now, follow the steps listed in the next section- Downloading VDI Logs to download the logs.

Collect Windows VDI logs

Log in to the corresponding Amazon EC2 instance from the Amazon EC2 console by following the instructions linked to in "Logging into an EC2 instance from the EC2 console"

at the beginning of this section. Get all the logs under the \$env:SystemDrive\Users \Administrator\RES\Bootstrap\Log\ directory on the VDI EC2 instance.

One of the ways to get the logs would be to upload them to S3 and then download them from there. To do that, follow the steps listed in the next section- Downloading VDI Logs.

Downloading VDI Logs

- 1. Update the VDI EC2 instance IAM role to allow S3 access.
- 2. Go to the EC2 console and select your VDI instance.
- 3. Select the IAM role it is using.
- In the **Permission Policies** section from the **Add permissions** dropdown menu, choose **Attach Policies** then select the **AmazonS3FullAccess** policy.
- Choose **Add permissions** to attach that policy.
- After that, follow the steps listed below based on your VDI type to download the logs. The module name would be the instance name.
 - Downloading logs from Linux EC2 instances for Linux. a.
 - Downloading logs from Windows EC2 instances for Windows. b.
- Lastly, edit the role to remove the AmazonS3FullAccess policy. 7.



Note

All VDIs use the same IAM role which is <env-name>-vdc-host-role-<region>

Downloading logs from Linux EC2 instances

Login to the EC2 instance from which you want to download logs and run the following commands to upload all the logs to an s3 bucket:

sudo su -

```
ENV_NAME=<environment_name>
REGION=<region>
ACCOUNT=<aws_account_number>
MODULE=<module_name>

cd /root/bootstrap
tar -czvf ${MODULE}_logs.tar.gz logs/ --overwrite
aws s3 cp ${MODULE}_logs.tar.gz s3://${ENV_NAME}-cluster-${REGION}-${ACCOUNT}/
${MODULE}_logs.tar.gz
```

After this, go to the S3 console, select the bucket with name <environment_name>-cluster-<region>-<aws_account_number> and download the previously uploaded <module_name>_logs.tar.gz file.

•••••

Downloading logs from Windows EC2 instances

Login to the EC2 instance from which you want to download logs and run the following commands to upload all the logs to an S3 bucket:

```
$ENV_NAME="<environment_name>"
$REGION="<region>"
$ACCOUNT="<aws_account_number>"
$MODULE="<module_name>"

$logDirPath = Join-Path -Path $env:SystemDrive -ChildPath "Users\Administrator\RES\Bootstrap\Log"
$zipFilePath = Join-Path -Path $env:TEMP -ChildPath "logs.zip"
Remove-Item $zipFilePath
Compress-Archive -Path $logDirPath -DestinationPath $zipFilePath
$bucketName = "${ENV_NAME}-cluster-${REGION}-${ACCOUNT}"
$keyName = "${MODULE}_logs.zip"
Write-S30bject -BucketName $bucketName -Key $keyName -File $zipFilePath
```

After this, go to the S3 console, select the bucket with name <environment_name>-cluster-<region>-<aws_account_number> and download the previously uploaded <module_name>_logs.zip file.

.....

Collecting ECS logs for the WaitCondition error

- 1. Go to the deployed stack and select the **Resources** tab.
- Expand Deploy → ResearchAndEngineeringStudio → Installer → Tasks → CreateTaskDef → CreateContainer → LogGroup, and select the log group to open CloudWatch logs.
- 3. Grab the latest log from this log group.

.....

Demo environment

Topics

- Demo environment login error when handling authentication request to identity provider
- Demo stack keycloak not working

.....

Demo environment login error when handling authentication request to identity provider

Issue

If you attempt to log in and get an 'Unexpected error when handling authentication request to identity provider', your passwords might be expired. This could be either the password for the user you are trying to log in as or your Active Directory Service Account.

Mitigation

- 1. Reset the user and Service Account passwords in the <u>Directory service console</u>.
- 2. Update the Service Account passwords in <u>Secrets Manager</u> to match the new password you entered above:
 - for the Keycloak stack: **PasswordSecret**-...-**RESExternal**-...-**DirectoryService**-... with Description: Password for Microsoft Active Directory
 - for RES: **res-ServiceAccountPassword**-... with Description: Active Directory Service Account Password

Demo environment 231

 Go to the <u>EC2 console</u> and terminate the cluster-manager instance. Auto Scaling rules will automatically trigger deployment of a new instance.

•••••

Demo stack keycloak not working

Issue

If your keycloak server crashed and, when you restarted the server, the IP of the instance changed, this might have resulted in keycloak breaking—the login page of your RES portal either fails to load or becomes stuck in a loading state which never resolves.

Mitigation

You will need to delete the existing infrastructure and redeploy the Keycloak stack to restore Keycloak to a healthy state. Follow these steps:

- 1. Go to Cloudformation. You should see two keycloak related stacks there:
 - <env-name>-RESSsoKeycloak-<random characters> (Stack1)

```
<env-name>-RESSsoKeycloak-<random characters>-RESSsoKeycloak-* (Stack2)
```

- 2. Delete Stack1. If prompted to delete the nested stack, select **Yes** to delete the nested stack.
 - Make sure the stack has been deleted completely.
- 3. Download the RES SSO Keycloak stack template here.
- 4. Deploy this stack manually with the exact same parameter values as the deleted stack. Deploy it from the CloudFormation console by going to Create Stack → With new resources (standard) → Choose an existing template → Upload a template file. Fill in the required parameters using the same inputs as the deleted stack. You can find these inputs in your deleted stack by changing the filter on the CloudFormation console and going to the Parameters tab. Make sure that the environment name, key pair, and other parameters match the original stack parameters.
- 5. Once the stack is deployed, your environment is ready to be used again. You can find the ApplicationUrl in the **Outputs** tab of the deployed stack.

.....

Demo environment 232

Known Issues

- Known Issues 2024.x
 - (2024.08) Virtual desktops fail to mount read/write Amazon S3 bucket with root bucket ARN and custom prefixing
 - (2024.06) Apply snapshot fails when the AD group name contains spaces
 - (2024.04-2024.04.02) Provided IAM Permission Boundary not attached to the VDI instances' role
 - (2024.04.02 and earlier) Windows NVIDIA instances in ap-southeast-2 (Sydney) fail to launch
 - (2024.04 and 2024.04.01) RES delete failure in GovCloud
 - (2024.04 2024.04.02) Linux virtual desktop may be stuck in the "RESUMING" status on reboot
 - (2024.04.02 and earlier) Fails to sync AD users whose SAMAccountName attribute includes capital letters or special characters
 - (2024.04.02 and earlier) Private key for accessing the bastion host is invalid
 - (2024.06 and earlier) Group members not synced to RES during AD sync
 - (2024.06 and earlier) CVE-2024-6387, RegreSSHion, Security Vulnerability in RHEL9 and Ubuntu VDIs

Known Issues 2024.x

.....

(2024.08) Virtual desktops fail to mount read/write Amazon S3 bucket with root bucket ARN and custom prefixing

Bug description

Research and Engineering Studio 2024.08 fails to mount read/write S3 buckets on to a virtual desktop infrastructure (VDI) instance when using a root bucket ARN (that is, arn:aws:s3:::example-bucket) and a custom prefix (project name or project name and user name).

Bucket configurations that are **not affected** by this issue include:

· read-only buckets

Known Issues 233

- read/write buckets with a prefix as part of the bucket ARN (that is, arn:aws:s3:::example-bucket/example-folder-prefix) and custom prefixing (project name or project name and user name)
- read/write buckets with a root bucket ARN, but no custom prefixing

After you provision a VDI instance, the specified mount directory for that S3 bucket will not have the bucket mounted. Although the mount directory on the VDI will be present, the directory will be empty and will not contain the current contents of the bucket. When you write a file to the directory using the terminal, the error Permission denied, unable to write a file will be thrown and the file contents will not be uploaded to the corresponding S3 bucket.

Affected versions

2024.08

Mitigation

- 1. To download the patch script and patch file (patch.py and s3_mount_custom_prefix_fix.patch), run the following command, replacing <outputdirectory> with the directory where you want to download the patch script and patch file and <environment-name> with the name of your RES environment:
 - a. The patch only applies to RES 2024.08.
 - b. The patch script requires AWS CLI v2, Python 3.9.16 or above, and Boto3.
 - c. Configure the AWS CLI for the account and region where RES is deployed, and make sure that you have Amazon S3 permissions to write to the bucket created by RES.

```
OUTPUT_DIRECTORY=<output-directory>
ENVIRONMENT_NAME=<environment-name>

mkdir -p ${OUTPUT_DIRECTORY}
curl https://research-engineering-studio-us-east-1.s3.amazonaws.com/
releases/2024.08/patch_scripts/patch.py --output ${OUTPUT_DIRECTORY}/patch.py
curl https://research-engineering-studio-us-east-1.s3.amazonaws.com/
releases/2024.08/patch_scripts/patches/s3_mount_custom_prefix_fix.patch --output
${OUTPUT_DIRECTORY}/s3_mount_custom_prefix_fix.patch
```

2. Navigate to the directory where the patch script and patch file are downloaded. Run the following patch command:

```
python3 ${OUTPUT_DIRECTORY}/patch.py --environment-name ${ENVIRONMENT_NAME} --res-
version 2024.08 --module virtual-desktop-controller --patch ${OUTPUT_DIRECTORY}/
s3_mount_custom_prefix_fix.patch
```

3. To terminate the Virtual Desktop Controller (vdc-controller) instance for your environment, run the following commands. (You already set the ENVIRONMENT_NAME variable to the name of your RES environment in the first step.)

```
INSTANCE_ID=$(aws ec2 describe-instances \
    --filters \
    Name=tag:Name,Values=${ENVIRONMENT_NAME}-vdc-controller \
    Name=tag:res:EnvironmentName,Values=${ENVIRONMENT_NAME}\
    --query "Reservations[0].Instances[0].InstanceId" \
    --output text)

aws ec2 terminate-instances --instance-ids ${INSTANCE_ID}
```

Note

For private VPC setups, if you haven't already done so, for the <RES-EnvironmentName>-vdc-custom-credential-broker-lambda function make sure to add the Environment variable with name AWS_STS_REGIONAL_ENDPOINTS and value of regional. See Amazon S3 bucket prerequisites for isolated VPC deployments for more information.

4. After the target group starting with the name <RES-EnvironmentName>-vdc-ext becomes healthy, new VDIs will need to be launched that will have the read/write S3 buckets with root bucket ARN and custom prefixing mounted correctly.

•••••

(2024.06) Apply snapshot fails when the AD group name contains spaces

Issue

RES 2024.06 fails to apply snapshots from prior versions if the AD groups contain spaces in their names.

The cluster-manager CloudWatch logs (under the <environment-name>/cluster-manager log group) will include the following error during AD sync:

```
[apply-snapshot] authz.role-assignments/<Group name with spaces>:group#sprojectID>:project FAILED_APPLY because: [INVALID_PARAMS] Actor key doesn't match the regex pattern ^[a-zA-Z0-9_.][a-zA-Z0-9_.-]{1,20}:(user|group)$
```

The error results from RES only accepting group names that meet the following requirements:

- It can only contain lowercase and uppercase ASCII letters, digits, dash(-), period (.), and underscore (_)
- A dash (-) is not allowed as the first character
- It cannot contain spaces.

Affected versions

2024.06

Mitigation

- To download the patch script and patch file (<u>patch.py</u> and <u>groupname_regex.patch</u>), run the following command, replacing <output-directory> with the directory where you want to put the files, and <environment-name> with the name of your RES environment:
 - a. The patch only applies to RES 2024.06
 - b. The patch script requires AWS CLI v2, Python 3.9.16 or above, and Boto3.
 - c. Configure the AWS CLI for the account and region where RES is deployed, and make sure that you have S3 permissions to write to the bucket created by RES:

```
OUTPUT_DIRECTORY=<<u>output-directory</u>>
ENVIRONMENT_NAME=<<u>environment-name</u>>

mkdir -p ${OUTPUT_DIRECTORY}
curl https://research-engineering-studio-us-east-1.s3.amazonaws.com/
releases/2024.06/patch_scripts/patch.py --output ${OUTPUT_DIRECTORY}/patch.py
curl https://research-engineering-studio-us-east-1.s3.amazonaws.com/
releases/2024.06/patch_scripts/patches/groupname_regex.patch --output
${OUTPUT_DIRECTORY}/groupname_regex.patch
```

2. Navigate to the directory where the patch script and patch file are downloaded. Run the following patch command:

```
python3 patch.py --environment-name ${ENVIRONMENT_NAME} --res-version 2024.06 --
module cluster-manager --patch ${OUTPUT_DIRECTORY}/groupname_regex.patch
```

3. To restart the Cluster Manager instance for your environment, run the following commands: You may also terminate the instance from the Amazon EC2 Management Console.

```
INSTANCE_ID=$(aws ec2 describe-instances \
    --filters \
    Name=tag:Name,Values=${ENVIRONMENT_NAME}-cluster-manager \
    Name=tag:res:EnvironmentName,Values=${ENVIRONMENT_NAME}\
    --query "Reservations[0].Instances[0].InstanceId" \
    --output text)

aws ec2 terminate-instances --instance-ids ${INSTANCE_ID}
```

Note

The patch allows AD group names to contain lower case and uppercase ASCII letters, digits, dash(-), period (.), underscore (_), and spaces with a total length between 1 and 30, inclusive.

.....

(2024.04-2024.04.02) Provided IAM Permission Boundary not attached to the VDI instances' role

The issue

Virtual desktop sessions are not properly inheriting their project's permission boundary configuration. This is a result of the permissions boundary defined by the IAMPermissionBoundary parameter not being properly assigned to a project during that project's creation.

Affected versions

2024.04 - 2024.04.02

Mitigation

Follow these steps to allow VDIs to properly inherit the permissions boundary assigned to a project:

- To download the patch script and patch file (<u>patch.py</u> and <u>vdi_host_role_permission_boundary.patch</u>), run the following command, replacing <outputdirectory> with the local directory where you'd like to put the files:
 - a. The patch only applies to RES 2024.04.02. If you are on version 2024.04 or 2024.04.01, you can follow the steps listed in the public document for minor version updates to update your environment to 2024.04.02.
 - b. The patch script requires AWS CLI v2), Python 3.9.16 or above, and Boto3.
 - c. Configure the AWS CLI for the account and region where RES is deployed, and make sure that you have S3 permissions to write to the bucket created by RES.

```
OUTPUT_DIRECTORY=<output-directory>

curl https://research-engineering-studio-us-east-1.s3.amazonaws.com/
releases/2024.04.02/patch_scripts/patch.py --output ${OUTPUT_DIRECTORY}/patch.py

curl https://research-engineering-studio-us-east-1.s3.amazonaws.com/
releases/2024.04.02/patch_scripts/patches/vdi_host_role_permission_boundary.patch
--output ${OUTPUT_DIRECTORY}/vdi_host_role_permission_boundary.patch
```

2. Navigate to the directory where the patch script and patch file are downloaded. Run the following patch command, replacing <environment-name> with the name of your RES environment:

```
python3 patch.py --environment-name <environment-name> --res-version 2024.04.02 --
module cluster-manager --patch vdi_host_role_permission_boundary.patch
```

3. Restart the cluster-manager instance in your environment by running this command, replacing <environment-name> with the name of your RES environment. You may also terminate the instance from the Amazon EC2 Management Console.

```
ENVIRONMENT_NAME=<environment-name>
INSTANCE_ID=$(aws ec2 describe-instances \)
```

```
--filters \
Name=tag:Name,Values=${ENVIRONMENT_NAME}-cluster-manager \
Name=tag:res:EnvironmentName,Values=${ENVIRONMENT_NAME}\
--query "Reservations[0].Instances[0].InstanceId" \
--output text)

aws ec2 terminate-instances --instance-ids ${INSTANCE_ID}
```

.....

(2024.04.02 and earlier) Windows NVIDIA instances in ap-southeast-2 (Sydney) fail to launch

The issue

Amazon Machine Images (AMIs) are used to spin up virtual desktops (VDIs) in RES with specific configurations. Each AMI has an associated ID that differs per region. The AMI ID configured in RES to launch Windows Nvidia instances in ap-southeast-2 (Sydney) is currently incorrect.

AMI-ID ami-0e190f8939a996caf for this type of instance configuration is incorrectly listed in apsoutheast-2 (Sydney). AMI ID ami-027cf6e71e2e442f4 should be used instead.

Users will get the following error when trying to launch an instance with the default ami -0e190f8939a996caf AMI.

```
An error occured (InvalidAMIID.NotFound) when calling the RunInstances operation: The image id '[ami-0e190f8939a996caf]' does not exist
```

Steps to reproduce the bug, including an example configuration file:

- Deploy RES in the ap-southeast-2 region.
- Launch an instance using Windows-NVIDIA default software stack (AMI ID ami-0e190f8939a996caf).

Affected versions

All RES versions 2024.04.02 or earlier are impacted

Mitigation

The following mitigation has been tested on RES version 2024.01.01:

Register a new software stack with the following settings

AMI ID: ami-027cf6e71e2e442f4

• Operating System: Windows

GPU Manufacturer: NVIDIA

• Min. Storage Size (GB): 30

• Min. RAM (GB): 4

Use this software stack to launch Windows-NVIDIA instances

.....

(2024.04 and 2024.04.01) RES delete failure in GovCloud

The issue

During the RES delete workflow the UnprotectCognitoUserPool Lambda inactivates Deletion Protection for Cognito User Pools that will later be deleted. The Lambda execution is started by the InstallerStateMachine.

Because of default AWS CLI version differences between Commercial and GovCloud regions, the update_user_pool call in the Lambda will fail in GovCloud regions.

Customers will get the following error when trying to delete RES in GovCloud regions:

Parameter validation failed: Unknown parameter in input: \"DeletionProtection \", must be one of: UserPoolId, Policies, LambdaConfig, AutoVerifiedAttributes, SmsVerificationMessage, EmailVerificationMessage, EmailVerificationSubject, VerificationMessageTemplate, SmsAuthenticationMessage, MfaConfiguration, DeviceConfiguration, EmailConfiguration, SmsConfiguration, UserPoolTags, AdminCreateUserConfig, UserPoolAddOns, AccountRecoverySetting

Steps to reproduce the bug:

- Deploy RES in a GovCloud region
- Delete the RES stack

Affected versions

RES version 2024.04 and 2024.04.01

Mitigation

The following mitigation has been tested on RES version 2024.04:

- Open the UnprotectCognitoUserPool Lambda
 - Naming convention: <env-name>-InstallerTasksUnprotectCognitoUserPool-...
- Runtime Settings -> Edit -> Select Runtime Python 3.11 -> Save.
- Open CloudFormation.
- Delete RES stack -> leave Retain Installer Resource UNCHECKED -> Delete.

.....

(2024.04 - 2024.04.02) Linux virtual desktop may be stuck in the "RESUMING" status on reboot

The issue

Linux virtual desktops can get stuck in "RESUMING" status when restarting after a manual or scheduled stop.

After the instance is rebooted, the AWS Systems Manager doesn't run any remote commands to create a new DCV session and the following log message is missing in the vdc-controller CloudWatch logs (under the <environment-name>/vdc/controller CloudWatch log group):

Handling message of type DCV_HOST_REBOOT_COMPLETE_EVENT

Affected versions

2024.04 - 2024.04.02

Mitigation

To recover the virtual desktops that are stuck in the "RESUMING" state:

- 1. SSH into the problem instance from the EC2 console.
- 2. Run the following commands on the instance:

sudo su -

```
/bin/bash /root/bootstrap/latest/virtual-desktop-host-linux/
configure_post_reboot.sh
sudo reboot
```

Wait for the instance to reboot.

To prevent new virtual desktops from running into the same issue:

 To download the patch script and patch file (<u>patch.py</u> and <u>vdi_stuck_in_resuming_status.patch</u>), run the following command, replacing <outputdirectory> with the directory where you want to put the files:

Note

- The patch only applies to RES 2024.04.02.
- The patch script requires AWS CLI v2, Python 3.9.16 or above, and Boto3.
- Configure the AWS CLI for the account and region where RES is deployed, and make sure that you have S3 permissions to write to the bucket created by RES.

```
OUTPUT_DIRECTORY=<output-directory>

curl https://research-engineering-studio-us-east-1.s3.amazonaws.com/
releases/2024.04.02/patch_scripts/patch.py --output ${OUTPUT_DIRECTORY}/patch.py

curl https://research-engineering-studio-us-east-1.s3.amazonaws.com/
releases/2024.04.02/patch_scripts/patches/vdi_stuck_in_resuming_status.patch --output ${OUTPUT_DIRECTORY}/vdi_stuck_in_resuming_status.patch
```

2. Navigate to the directory where the patch script and patch file are downloaded. Run the following patch command, replacing <environment-name> with the name of your RES environment and <aws-region> with the region where RES is deployed:

```
python3 patch.py --environment-name <environment-name> --res-version 2024.04.02
  --module virtual-desktop-controller --patch vdi_stuck_in_resuming_status.patch --
region <aws-region>
```

3. To restart the VDC Controller instance for your environment, run the following commands, replacing <environment-name> with the name of your RES environment:

```
ENVIRONMENT_NAME=<environment-name>

INSTANCE_ID=$(aws ec2 describe-instances \
     --filters \
     Name=tag:Name,Values=${ENVIRONMENT_NAME}-vdc-controller \
     Name=tag:res:EnvironmentName,Values=${ENVIRONMENT_NAME}\
     --query "Reservations[0].Instances[0].InstanceId" \
     --output text)

aws ec2 terminate-instances --instance-ids ${INSTANCE_ID}
```

.....

(2024.04.02 and earlier) Fails to sync AD users whose SAMAccountName attribute includes capital letters or special characters

The issue

RES fails to sync AD users after SSO is set up for at least two hours (two AD sync cycles). The cluster-manager CloudWatch logs (under the <environment-name>/cluster-manager log group) include the following error during AD sync:

```
Error: [INVALID_PARAMS] Invalid params: user.username must match regex: ^(?=.{3,20}$)
(?![_.])(?!.*[_.]{2})[a-z0-9._]+(?<![_.])$</pre>
```

The error results from RES only accepting a SAMAccount username that meets the following requirements:

- It can only contain lower case ASCII letters, digits, period (.), underscore (_).
- A period or underscore is not allowed as the first or last character.
- It cannot contain two continuous periods or underscores (e.g. .., __, ._, _.).

Affected versions

2024.04.02 and earlier

Mitigation

To download the patch script and patch file (<u>patch.py</u> and <u>samaccountname_regex.patch</u>), run
the following command, replacing <output-directory> with the directory where you want
to put the files:

Note

- The patch only applies to RES 2024.04.02.
- The patch script requires AWS CLI v2, Python 3.9.16 or above, and Boto3.
- Configure the AWS CLI for the account and region where RES is deployed, and make sure that you have S3 permissions to write to the bucket created by RES.

```
OUTPUT_DIRECTORY=curl https://research-engineering-studio-us-east-1.s3.amazonaws.com/
releases/2024.04.02/patch_scripts/patch.py --output ${OUTPUT_DIRECTORY}/patch.py

curl https://research-engineering-studio-us-east-1.s3.amazonaws.com/
releases/2024.04.02/patch_scripts/patches/samaccountname_regex.patch --output
${OUTPUT_DIRECTORY}/samaccountname_regex.patch
```

2. Navigate to the directory where the patch script and patch file are downloaded. Run the following patch command, replacing <environment-name> with the name of your RES environment:

```
python3 patch.py --environment-name <environment-name> --res-version 2024.04.02 --
module cluster-manager --patch samaccountname_regex.patch
```

3. To restart the Cluster Manager instance for your environment, run the following commands, replacing <environment-name> with the name of your RES environment. You may also terminate the instance from the Amazon EC2 Management Console.

```
ENVIRONMENT_NAME=<environment-name>

INSTANCE_ID=$(aws ec2 describe-instances \
    --filters \
    Name=tag:Name, Values=${ENVIRONMENT_NAME}-cluster-manager \
    Name=tag:res:EnvironmentName, Values=${ENVIRONMENT_NAME}\
    --query "Reservations[0].Instances[0].InstanceId" \
```

```
--output text)
aws ec2 terminate-instances --instance-ids ${INSTANCE_ID}
```

User Guide

.....

(2024.04.02 and earlier) Private key for accessing the bastion host is invalid

The issue

When a user downloads the private key to access the bastion host from the RES web portal, the key is not well formatted—multiple lines are downloaded as a single line, which makes the key invalid. The user will get the following error when they attempt to access the bastion host with the downloaded key:

```
Load key "<downloaded-ssh-key-path>": error in libcrypto 
<user-name>@<bastion-host-public-ip>: Permission denied (publickey,gssapi-keyex,gssapi-with-mic)
```

Affected versions

2024.04.02 and earlier

Mitigation

We recommend using Chrome to download the keys, as this browser is unaffected.

Alternatively, the key file can be reformatted by creating a new line after ----BEGIN PRIVATE KEY---- and another new line just before ----END PRIVATE KEY----.

.....

(2024.06 and earlier) Group members not synced to RES during AD sync

Bug description

Group members will not properly sync to RES if the GroupOU differs from the UserOU.

RES creates an Idapsearch filter when attempting to sync users from an AD group. The current filter incorrectly utilizes the UserOU parameter instead of the GroupOU parameter. The result is that the search fails to return any users. This behavior only occurs in instances where the UsersOU and GroupOU differ.

Affected versions

This issue affects all RES versions 2024.06 or earlier

Mitigation

Follow these steps to resolve the issue:

To download the patch.py script and group_member_sync_bug_fix.patch file, run the following commands, replacing <output-directory> with the local directory where you'd like to download the files, and <res_version> with the version of RES you want to patch:

Note

- The patch script requires AWS CLI v2, Python 3.9.16 or above, and Boto3.
- Configure the AWS CLI for the account and region where RES is deployed, and make sure that you have S3 permissions to write to the bucket created by RES.
- The patch only supports RES versions 2024.04.02 and 2024.06. If you are using 2024.04 or 2024.04.01, you can follow the steps listed in Minor version updates to first update your environment to 2024.04.02 prior to applying the patch.
 - RES Version: RES 2024.04.02

Patch download link: 2024.04.02_group_member_sync_bug_fix.patch

• RES Version: RES 2024.06

Patch download link: 2024.06_group_member_sync_bug_fix.patch

```
OUTPUT_DIRECTORY=curl output-directory>
RES_VERSION=res_version>
mkdir -p ${OUTPUT_DIRECTORY}

curl https://research-engineering-studio-us-east-1.s3.amazonaws.com/releases/
${RES_VERSION}/patch_scripts/patch.py --output ${OUTPUT_DIRECTORY}/patch.py

curl https://research-engineering-studio-us-east-1.s3.amazonaws.com/releases/
${RES_VERSION}/patch_scripts/patches/${RES_VERSION}_group_member_sync_bug_fix.patch
--output ${OUTPUT_DIRECTORY}/${RES_VERSION}_group_member_sync_bug_fix.patch
```

2. Navigate to the directory where the patch script and patch file are downloaded. Run the following patch command, replacing <environment-name> with the name of your RES environment:

```
cd ${OUTPUT_DIRECTORY}
ENVIRONMENT_NAME=<environment-name>

python3 patch.py --environment-name ${ENVIRONMENT_NAME} --res-
version ${RES_VERSION} --module cluster-manager --patch $PWD/
${RES_VERSION}_group_member_sync_bug_fix.patch
```

3. To restart the cluster-manager instance for your environment, run the following commands:

```
INSTANCE_ID=$(aws ec2 describe-instances \
    --filters \
    Name=tag:Name,Values=${ENVIRONMENT_NAME}-cluster-manager \
    Name=tag:res:EnvironmentName,Values=${ENVIRONMENT_NAME}\
    --query "Reservations[0].Instances[0].InstanceId" \
    --output text)

aws ec2 terminate-instances --instance-ids ${INSTANCE_ID}
```

.....

(2024.06 and earlier) CVE-2024-6387, RegreSSHion, Security Vulnerability in RHEL9 and Ubuntu VDIs

Bug description

<u>CVE-2024-6387</u>, dubbed regreSSHion, has been identified in the OpenSSH server. This vulnerability enables remote, unauthenticated attackers to execute arbitrary code on the target server, presenting a severe risk to systems that utilize OpenSSH for secure communications.

For RES, the standard configuration is to go through the bastion host to SSH into virtual desktops, and the bastion host is unaffected by this vulnerability. However, the default AMI (Amazon Machine Image) we provide for RHEL9 and Ubuntu2024 VDIs (Virtual Desktop Infrastructure) in **ALL** RES versions utilizes an OpenSSH version which is vulnerable to the security threat.

This means that existing RHEL9 and Ubuntu2024 VDIs could be exploitable, but the attacker would require access to the bastion host.

More details about the issue can found here.

Affected versions

This issue affects all RES versions 2024.06 or earlier.

Mitigation

Both RHEL9 and Ubuntu have released patches for OpenSSH which fixes the security vulnerability. These can be pulled using the platform's respective package manager.

If you have existing RHEL9 or Ubuntu VDIs, we recommend following the **PATCH EXISTING VDIs** instructions below. To patch future VDIs, we recommend following the **PATCH FUTURE VDIs** instructions. These instructions describe how to run a script to apply the platform update on your VDIs.

PATCH EXISTING VDIs

- 1. Run the following command which will patch all existing Ubuntu and RHEL9 VDIs:
 - a. The patch script requires AWS CLI v2.
 - b. Configure the AWS CLI for the account and region where RES is deployed, and make sure that you have AWS Systems Manager permissions to send a Systems Manager Run Command.

```
aws ssm send-command \
    --document-name "AWS-RunRemoteScript" \
    --targets "Key=tag:res:NodeType,Values=virtual-desktop-dcv-host" \
    --parameters '{"sourceType":["S3"],"sourceInfo":["{\"path\":\"https://
    research-engineering-studio-us-east-1.s3.amazonaws.com/releases/2024.06/
    patch_scripts/scripts/patch_openssh.sh\"}"],"commandLine":["bash
    patch_openssh.sh"]}'
```

2. You can verify the script ran successfully on the <u>Run Command page</u>. Click on the **Command History** tab, select the most recent Command ID, and verify that all instance IDs have a SUCCESS message.

PATCH FUTURE VDIS

 To download the patch script and patch file (<u>patch.py</u> and <u>update_openssh.patch</u>) run the following commands, replacing <output-directory> with the directory where you want to download the files, and <environment-name> with the name of your RES environment:

Note

- The patch only applies to RES 2024.06.
- The patch script requires AWS CLI v2), Python 3.9.16 or above, and Boto3.
- Configure your copy of the AWS CLI for the account and region where RES is deployed, and make sure that you have S3 permissions to write to the bucket created by RES.

```
OUTPUT_DIRECTORY=<output-directory>
ENVIRONMENT_NAME=<environment-name>

curl https://research-engineering-studio-us-east-1.s3.amazonaws.com/
releases/2024.06/patch_scripts/patch.py --output ${OUTPUT_DIRECTORY}/patch.py

curl https://research-engineering-studio-us-east-1.s3.amazonaws.com/
releases/2024.06/patch_scripts/patches/update_openssh.patch --output
${OUTPUT_DIRECTORY}/update_openssh.patch
```

2. Run the following patch command:

```
python3 ${OUTPUT_DIRECTORY}/patch.py --environment-name ${ENVIRONMENT_NAME} --res-
version 2024.06 --module virtual-desktop-controller --patch ${OUTPUT_DIRECTORY}/
update_openssh.patch
```

3. Restart the VDC Controller instance for your environment with the following commands:

```
INSTANCE_ID=$(aws ec2 describe-instances \
    --filters \
    Name=tag:Name,Values=${ENVIRONMENT_NAME}-vdc-controller \
    Name=tag:res:EnvironmentName,Values=${ENVIRONMENT_NAME}\
    --query "Reservations[0].Instances[0].InstanceId" \
    --output text)
```

aws ec2 terminate-instances --instance-ids \${INSTANCE_ID}

Patching future VDIs is only supported on RES versions 2024.06 and later. To patch future VDIs in RES environments with versions earlier than 2024.06, first upgrade the RES environment to 2024.06 using the instructions at: Major version updates.

.....

Notices

Each Amazon EC2 instance comes with two Remote Desktop Services (Terminal Services) licenses for administration purposes. This <u>information</u> is available to help you provision these licenses for your administrators. You can also use <u>AWS Systems Manager Session Manager</u>, which enables remoting into Amazon EC2 instances without RDP and without a need for RDP licenses. If additional Remote Desktop Services licenses are needed, Remote Desktop user CALs should be purchased from Microsoft or a Microsoft license reseller. Remote Desktop users CALs with active Software Assurance have License Mobility benefits and can be brought to AWS default (shared) tenant environments. For information on bringing licenses without Software Assurance or License Mobility benefits, please see <u>this section</u> of the FAQ.

Customers are responsible for making their own independent assessment of the information in this document. This document: (a) is for informational purposes only, (b) represents AWS current product offerings and practices, which are subject to change without notice, and (c) does not create any commitments or assurances from AWS and its affiliates, suppliers or licensors. AWS products or services are provided "as is" without warranties, representations, or conditions of any kind, whether express or implied. AWS responsibilities and liabilities to its customers are controlled by AWS agreements, and this document is not part of, nor does it modify, any agreement between AWS and its customers.

Research and Engineering Studio on AWS is licensed under the terms of the Apache License Version 2.0 available at <u>The Apache Software Foundation</u>.

Revisions

For more information, see the $\underline{\text{CHANGELOG.md}}$ file in the GitHub repository.

Date	Change
October 2024	 Release version 2024.10: Added support for — Environment Boundaries. Desktop sharing profiles. Virtual desktop interface autostop.
August 2024	 Release version 2024.08: Added support for — mounting Amazon S3 buckets to Linux Virtual Desktop Infrastructure (VDI) instances. See Amazon S3 buckets. custom project permissions, an enhanced permission model that allows for cusomization of existing roles and the addition of custom roles. See Permission policy. User Guide: expanded the Troubleshooting section.
June 2024	 Release version 2024.06 — Ubuntu support, Project owner permissions. User Guide: added <u>Create a demo environme</u> nt
April 2024	Release version 2024.04 — RES-ready AMIs and project launch templates
March 2024	Additional troubleshooting topics, CloudWatc h Logs retention, uninstall minor versions

Date	Change
February 2024	Release version 2024.01.01 — updated deployment template
January 2024	Release version 2024.01
December 2023	GovCloud directions and templates added
November 2023	Initial release