



Veeam Backup for AWS

Version 8.0

User Guide

August, 2024

© 2024 Veeam Software.

All rights reserved. All trademarks are the property of their respective owners.

No part of this publication may be reproduced, transmitted, transcribed, stored in a retrieval system, or translated into any language in any form by any means, without written permission from Veeam Software (Veeam). The information contained in this document represents the current view of Veeam on the issue discussed as of the date of publication and is subject to change without notice. Veeam shall not be liable for technical or editorial errors or omissions contained herein. Veeam makes no warranties, express or implied, in this document. Veeam may have patents, patent applications, trademark, copyright, or other intellectual property rights covering the subject matter of this document. All other trademarks mentioned herein are the property of their respective owners. Except as expressly provided in any written license agreement from Veeam, the furnishing of this document does not give you any license to these patents, trademarks, copyrights, or other intellectual property.

NOTE

Read the End User Software License Agreement before using the accompanying software programs. Using any part of the software indicates that you accept the terms of the End User Software License Agreement.

Contents

CONTACTING VEEAM SOFTWARE	9
ABOUT THIS DOCUMENT	10
OVERVIEW	11
Integration with Veeam Backup & Replication	13
Solution Architecture	14
Backup Server	15
AWS Plug -In for Veeam Backup & Replication	16
Backup Appliances	17
Backup Repositories	19
Worker Instances	20
Additional Repositories and Tape Devices	28
Gateway Servers	29
Protecting EC2 Instances	30
EC2 Backup	32
EC2 Restore	41
Protecting RDS Resources	44
RDS Backup	45
RDS Restore	50
Protecting DynamoDB Tables	52
DynamoDB Backup	54
DynamoDB Restore	56
Protecting Redshift Clusters	57
Redshift Backup	60
Redshift Restore	62
Protecting EFS File Systems	63
EFS Backup	64
EFS Restore	67
Protecting FSx File Systems	68
FSx Backup	73
FSx Restore	75
Protecting VPC Configurations	76
VPC Configuration Backup	77
Exporting VPC Configuration	79
VPC Configuration Restore	80
Retention Policies	82
Immutability	83
Private Network Deployment	86

Backup Appliances in Private Environment	87
Worker Instances in Private Environment	94
Data Encryption.....	118
Backup Repository Encryption	119
AWS KMS Encryption	120
PLANNING AND PREPARATION	134
System Requirements	135
Ports.....	137
AWS Services.....	140
Plug-In Permissions	142
IAM Permissions	154
Worker IAM Permissions	155
Repository IAM Permissions.....	166
Backup IAM Permissions	168
Restore IAM Permissions	199
Full List of IAM Permissions	223
IAM Permissions Changelog	236
Considerations and Limitations	238
Sizing and Scalability Guidelines	244
Backup Appliance.....	245
Object Storage.....	251
Backup Policies	252
Worker Instances	254
DEPLOYMENT	255
Deploying Plug-In.....	256
Installing Plug-In	257
Installing and Uninstalling Plug-In in Unattended Mode.....	258
Upgrading Plug-In	260
Uninstalling Plug-In	261
Deploying Backup Appliance	262
Step 1. Launch New Veeam Backup for AWS Appliance Wizard	263
Step 2. Choose Deployment Mode	264
Step 3. Specify AWS Account	265
Step 4. Specify EC2 Instance Name and Description	266
Step 5. Specify Connection Type	267
Step 6. Specify Network Settings.....	268
Step 7. Specify User Credentials.....	270
Step 8. Track Progress	271
Step 9. Finish Working with Wizard	272
Failure and Recovery.....	273

LICENSING	274
Limitations	275
Scenarios	276
Viewing License Information	277
Revoking License Units	280
ACCESSING VEEAM BACKUP FOR AWS	282
Accessing Web UI from Console	283
Accessing Web UI from Workstation	284
CONFIGURING VEEAM BACKUP FOR AWS.....	287
Managing Backup Appliances	288
Adding Appliances	289
Editing Appliance Settings.....	302
Rescanning Appliances	304
Removing Appliances	305
Managing Backup Repositories	307
Adding Backup Repositories Using Console	308
Adding Backup Repositories Using Web UI.....	322
Editing Backup Repository Settings.....	336
Rescanning Backup Repositories	339
Removing Backup Repositories	340
Managing IAM Roles	342
Adding IAM Roles	343
Editing IAM Role Settings.....	356
Checking IAM Role Permissions.....	358
Removing IAM Roles	362
Managing User Accounts	363
Adding User Accounts	365
Editing User Account Settings	371
Changing User Passwords.....	372
Enabling Multi-Factor Authentication	373
Managing Database Accounts	374
Adding Database Accounts	375
Editing Database Accounts	380
Removing Database Accounts	381
Managing Worker Instances	382
Managing Worker Configurations	383
Managing Worker Profiles	393
Adding Worker Tags.....	397
Configuring General Settings.....	398
Configuring Private Network Deployment	399

Configuring Global Retention Settings	401
Configuring Global Notification Settings	403
Replacing Security Certificates	407
Changing Time Zone	409
Configuring SSO Settings.....	410
Performing Configuration Backup and Restore	412
Performing Configuration Backup	413
Performing Configuration Restore.....	417
VIEWING AVAILABLE RESOURCES.....	435
Adding Resources to Policy	437
PERFORMING BACKUP	438
Performing Backup Using Console	440
Creating Backup Policies.....	441
Editing Backup Policy Settings	442
Enabling and Disabling Backup Policies	443
Starting and Stopping Backup Policies	444
Deleting Backup Policies	445
Creating Backup Copy Jobs.....	446
Copying Backups to Tapes	447
Performing Backup Using Web UI	448
Performing EC2 Backup	449
Performing RDS Backup	491
Performing DynamoDB Backup	528
Performing Redshift Backup	559
Performing EFS Backup	585
Performing FSx Backup	615
Performing VPC Configuration Backup	644
Managing Backup Policies	655
MANAGING BACKED-UP DATA.....	660
Managing Backed-Up Data Using Console	661
Managing Backed-Up Data Using Web UI	664
EC2 Data	665
RDS Data	672
DynamoDB Data	676
Redshift Data	679
EFS Data.....	682
FSx Data.....	685
VPC Configuration Data	688
PERFORMING RESTORE.....	707
EC2 Restore.....	708

EC2 Restore Using Console	709
EC2 Restore Using Web UI	728
RDS Restore	770
RDS Restore Using Console	771
RDS Restore Using Web UI	798
DynamoDB Restore	824
DynamoDB Restore Using Console	825
DynamoDB Restore Using Web UI	826
Redshift Restore	838
Redshift Restore Using Console	839
Redshift Restore Using Web UI	840
EFS Restore	852
EFS Restore Using Console	853
EFS Restore Using Web UI	865
FSx Restore	891
FSx Restore Using Console	892
FSx Restore Using Web UI	893
VPC Configuration Restore	907
Performing VPC Configuration Restore Using Console	908
VPC Configuration Restore Using Web UI	909
Instant Recovery	925
Exporting Disks	927
Publishing Disks	928
Restoring to Microsoft Azure	929
Restoring to Google Cloud	931
Restoring to Nutanix AHV	932
REVIEWING DASHBOARD	934
VIEWING SESSION STATISTICS	936
COLLECTING OBJECT PROPERTIES	938
UPDATING VEEAM BACKUP FOR AWS	939
Updating Appliances Using Console	940
Upgrading to 8.0 from Version 6.0 or Earlier	942
Updating Appliances Using Web UI	944
Upgrading Appliances	945
Checking for Updates	946
Installing Updates	947
Updating IAM Roles	951
Viewing Updates History	953
Configuring Web Proxy	954
GETTING TECHNICAL SUPPORT	956

APPENDICES.....	959
Appendix A. Creating IAM Roles in AWS	960
Appendix B. Creating IAM Policies in AWS	962
Appendix C. Configuring Endpoints in AWS	963
Appendix D. Enabling Swap Partition	969
Appendix E. Uninstalling Backup Appliances Deployed from AWS Marketplace	972
Deleting CloudFormation Stack	974
Deleting AWS Resources	975
Appendix F. Installing and Removing License	976

Contacting Veeam Software

At Veeam Software we value feedback from our customers. It is important not only to help you quickly with your technical issues, but it is our mission to listen to your input and build products that incorporate your suggestions.

Customer Support

Should you have a technical concern, suggestion or question, visit the [Veeam Customer Support Portal](#) to open a case, search our knowledge base, reference documentation, manage your license or obtain the latest product release.

Company Contacts

For the most up-to-date information about company contacts and office locations, visit the [Veeam Contacts Webpage](#).

Online Support

If you have any questions about Veeam products, you can use the following resources:

- Full documentation set: veeam.com/documentation-guides-datasheets.html
- Veeam R&D Forums: forums.veeam.com

About This Document

This guide is intended for IT managers, cloud infrastructure administrators, and other personnel responsible for the product installation and operation.

The guide contains information on Veeam Backup for AWS configuration and provides a set of tasks that are required to perform data protection and disaster recovery operations.

Overview

Veeam Backup for Amazon Web Services (Veeam Backup for AWS) is a solution developed for protection and disaster recovery tasks for AWS environments: Amazon Elastic Compute Cloud (EC2), Amazon Relational Database Service (RDS), Amazon Redshift, Amazon DynamoDB, Amazon Elastic File System (EFS) and Amazon FSx File System. Veeam Backup for AWS also allows you to back up and restore Amazon Virtual Private Cloud (VPC) configurations.

With Veeam Backup for AWS, you can perform the following data protection operations:

- Create cloud-native snapshots of EC2 instances.
- Create cloud-native snapshots of RDS resources: DB instances and Amazon Aurora DB clusters (Aurora DB clusters).
- Replicate cloud-native snapshots to any AWS Region within any AWS account.
- Create image-level backups of EC2 instances and keep them in Amazon Simple Storage Service (Amazon S3) for high availability, cost-effective and long-term storage.
- [Available only for backup appliances managed by Veeam Backup & Replication] Create image-level backups of PostgreSQL DB instances and keep them in Amazon Simple Storage Service (Amazon S3) for high availability, cost-effective and long-term storage.
- [Available only for backup appliances managed by Veeam Backup & Replication] Create backups of DynamoDB tables and store them in any backup vault in the source AWS Region.
- [Available only for backup appliances managed by Veeam Backup & Replication] Create backup copies of DynamoDB tables and store them in any AWS Region within the same AWS account.
- [Available only for backup appliances managed by Veeam Backup & Replication] Create backups of Redshift clusters and store them in any backup vault in the source AWS Region.
- Create backups of EFS file systems and store them in any backup vault in the source AWS Region.
- Create backup copies of EFS file systems and store them in any AWS Region within the same AWS account.
- [Available only for backup appliances managed by Veeam Backup & Replication] Create backups of FSx file systems and store them in any backup vault in the specific AWS Regions.
- [Available only for backup appliances managed by Veeam Backup & Replication] Create backup copies of FSx file systems and store them in specific AWS Regions within the same AWS account.
- Create backups of VPC configurations and keep them in the Veeam Backup for AWS database and in Amazon S3.
- Create backups of the Veeam Backup for AWS configuration database.

To recover backed-up data, you can perform the following operations:

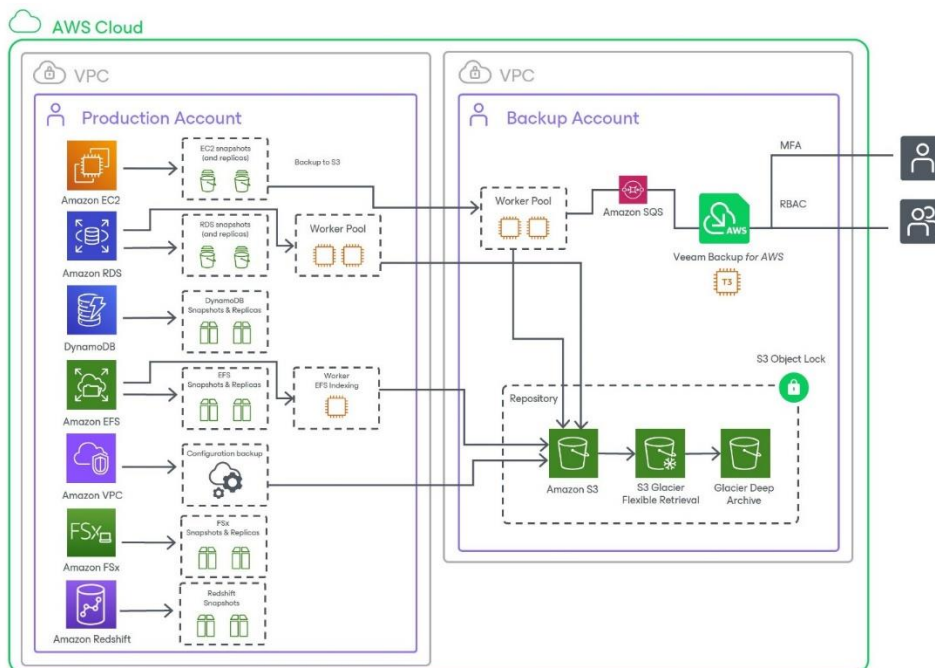
- Restore entire EC2 instances.
- Restore EC2 instance volumes.
- Restore EC2 instance files and folders.
- [Available only for backup appliances managed by Veeam Backup & Replication] Restore entire EC2 instances to Microsoft Azure, Google Cloud and Nutanix AHV.
- [Available only for backup appliances managed by Veeam Backup & Replication] Perform Instant Recovery of EC2 instances to VMware vSphere and Hyper-V environments, and to Nutanix AHV clusters.

- Restore RDS DB instances and Aurora DB clusters.
- [Available only for backup appliances managed by Veeam Backup & Replication] Restore PostgreSQL DB instances.
- [Available only for backup appliances managed by Veeam Backup & Replication] Restore DynamoDB tables.
- [Available only for backup appliances managed by Veeam Backup & Replication] Restore Red shift clusters.
- Restore entire EFS and FSx file systems.
- Restore EFS files and directories.
- Restore entire VPC configurations of AWS Regions.
- Restore specific items of VPC configurations of AWS Regions.
- Restore the Veeam Backup for AWS configuration database to the same or another backup appliance

IMPORTANT

Consider the following:

- Veeam Backup for AWS is available only in AWS Global and AWS GovCloud (US) regions.
- Starting from version 7.0, Veeam Backup for AWS is part of the Veeam Backup & Replication solution, and some new features are available only for backup appliances managed by Veeam Backup & Replication. For more information, see [Integration with Veeam Backup & Replication](#).



Integration with Veeam Backup & Replication

Starting from version 7.0, Veeam Backup for AWS is part of the Veeam Backup & Replication solution. AWS Plug-in for Veeam Backup & Replication extends the Veeam Backup & Replication functionality and allows you to add backup appliances to Veeam Backup & Replication. With AWS Plug-in for Veeam Backup & Replication, you can manage data protection and recovery operations for all these appliances from a single Veeam Backup & Replication console.

Versions 7.0 and 8.0 come with 4 major features – the ability to create image-level backups of PostgreSQL DB instances, as well as the ability to back up DynamoDB tables, Redshift clusters and FSx file systems. These features are available only for backup appliances managed by a Veeam Backup & Replication server. To unlock the full functionality, [install AWS Plug-in for Veeam Backup & Replication on the server](#) and [add your appliances](#) to the backup infrastructure.

IMPORTANT

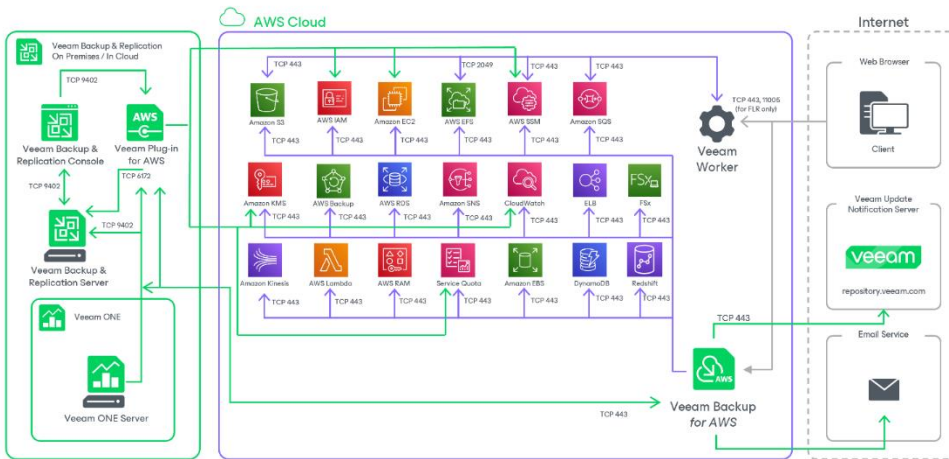
Consider the following:

- If you remove a backup appliance from the backup infrastructure, the following will happen:
- You will no longer be able to create image-level backups of DB instances, and the existing RDS backup policies configured to create these backups will start failing. To work around the issue, you can disable image-level backup when [editing backup policy settings](#).
- You will no longer be able to add and start DynamoDB, Redshift and FSx backup policies. Creating DynamoDB, Redshift and FSx backups manually will also be unavailable.
- If the connection between a backup appliance and the backup server is lost for more than 31 days, the appliance will enter the standalone mode, and you will no longer be able to back up DB instances and DynamoDB tables, as well as Redshift clusters and FSx file systems.

Solution Architecture

The Veeam Backup for AWS architecture includes the following components:

- Backup server
- AWS Plug-in for Veeam Backup & Replication
- Backup appliances
- Backup repositories
- Worker instances
- Additional repositories and tape devices
- Gateway servers



Backup Server

The backup server is a Windows-based physical or virtual machine on which Veeam Backup & Replication is installed. It is the core component in the backup infrastructure. For more information, see the Veeam Backup & Replication User Guide, section [Backup Server](#).

AWS Plug-In for Veeam Backup & Replication

Plug-in is an architecture component that enables integration between Veeam Backup & Replication and Veeam Backup for AWS.

Backup Appliances

A backup appliance is a Linux-based EC2 instance where Veeam Backup for AWS is installed.

If you have multiple backup appliances in AWS, you can add the appliances to Veeam Backup & Replication, and then use the Veeam Backup & Replication console as the central management console for Veeam Backup for AWS operations. For more information on the Veeam Backup & Replication console, see the [Veeam Backup & Replication User Guide](#).

Backup Appliance Software

The EC2 instance running Veeam Backup for AWS is deployed with the pre-installed set of software components:

- Ubuntu 22.04 LTS
- ASP.NET Core Runtime 8.0
- PostgreSQL 15
- nginx 1.18
- libpam-google-authenticator 20191231-2
- Veeam Backup for AWS installation packages

In case any software updates become available for the backup appliance, these updates can be installed using the Veeam Updater service as described in section [Updating Veeam Backup for AWS](#).

Backup Appliance Functionality

The backup appliance performs the following administrative activities:

- Manages architecture components.
- Coordinates snapshot creation, backup and recovery tasks.
- Controls backup policy scheduling.
- Generates daily reports and email notifications.
- Manages backup and snapshot retention tasks.

Backup Appliance Components

The backup appliance uses the following components:

- **Backup service** – coordinates data protection and disaster recovery operations.
- **Configuration database** – stores data on the existing backup policies, worker instance configurations, added IAM roles, sessions and so on, as well as information on the available and protected resources collected from AWS.
- **Configuration restore service** – allows users to back up and restore the configuration of the backup appliance.
- **Web UI** – provides a web interface that allows user to access to the Veeam Backup for AWS functionality.

- **Veeam Updater service** – allows Veeam Backup for AWS to check, view and install product and package updates.
- **Veeam FLR service** – allows users to restore individual files and folders of protected EC2 instances.
- **Self Backup service** – allows Veeam Backup for AWS to back up and restore the configuration database of the backup appliance.
- **REST API service** – allows users to perform operations with Veeam Backup for AWS entities using HTTP requests and standard HTTP methods. For more information, see the [Veeam Backup for AWS REST API Reference](#).

Backup Repositories

A backup repository is a folder in an Amazon S3 bucket where Veeam Backup for AWS stores EC2 and RDS image-level backups, additional copies of Amazon VPC backups, indexes of EFS file systems and configuration backups of standalone backup appliances.

To communicate with a backup repository, Veeam Backup for AWS uses **Veeam Data Mover** – the service that runs on a [worker instance](#) and that is responsible for data processing and transfer. When a backup policy addresses the backup repository, Veeam Data Mover establishes a connection with the repository to enable data transfer. To learn how Veeam Backup for AWS communicates with backup repositories, see [Managing Backup Repositories](#).

IMPORTANT

Backup files are stored in backup repositories in the native Veeam format and must be modified neither manually nor by 3rd party tools. Otherwise, Veeam Backup for AWS may fail to restore the backed-up data.

Encryption on Backup Repositories

For enhanced data security, Veeam Backup for AWS allows you to enable encryption at the repository level. Veeam Backup for AWS encrypts backup files stored in backup repositories the same way as Veeam Backup & Replication encrypts backup files stored in backup repositories. To learn what algorithms Veeam Backup & Replication uses to encrypt backup files, see the Veeam Backup & Replication User Guide, section [Encryption Standards](#). To learn how to enable encryption at the repository level, see [Adding Backup Repositories](#).

Veeam Backup for AWS also supports scenarios where data is backed up to S3 buckets with enabled Amazon S3 default encryption. You can add the S3 bucket to the backup infrastructure and use it as a target location for image-level backups. For information on Amazon S3 default encryption, see [AWS Documentation](#).

Worker Instances

To perform most data protection and disaster recovery operations (such as creating EC2 and RDS image-level backups, restoring backed-up data, EFS indexing or retention tasks), Veeam Backup for AWS uses worker instances. Worker instances are temporary Linux-based EC2 instances that are responsible for the interaction between the backup appliance, AWS services and other Veeam Backup for AWS components. Worker instances process backup workload and distribute backup traffic when transferring data to backup repositories.

Veeam Backup for AWS automatically deploys a worker instance in Amazon EC2 for the duration of a backup, restore or retention operation and removes it immediately as soon as the operation completes. For example, Veeam Backup for AWS deploys one worker instance per each AWS resource specified in a EC2 backup policy.

NOTE

The location of each deployed worker instance depends on the operation being performed and on the resource being processed. For more information, see sections [Worker Deployment Options](#) and [Worker Instance Locations](#).

Worker instances use the following services:

- **Veeam Data Mover** – the service that performs data processing tasks. During backup, Veeam Data Mover retrieves data of protected AWS resources and transfers it to backup repositories. During restore, Veeam Data Mover transfers backed-up data from backup repositories to the target location.
- **File-level recovery browser** – the web service that allows you to find and save files and folders of a backed-up EC2 instance to the local machine or to the original location. The file-level recovery browser is installed automatically on every worker instance that is deployed for file-level recovery.

Security Certificates for Worker Instances

During the file-level recovery process, Veeam Backup for AWS uses self-signed TLS certificates to establish secure communication between the web browser on the local machine and the file-level recovery browser on the worker instance. A self-signed certificate is generated automatically on the worker instance when the restore session starts.

Worker Instance Network Settings

To deploy worker instances, Veeam Backup for AWS uses either the default or the [most appropriate network settings](#) of AWS Regions. However, you can add specific worker configurations as described in section [Managing Worker Configurations](#).

Required Ports

The following network ports must be open to ensure proper operation of worker instances in Veeam Backup for AWS architecture:

From	To	Protocol	Port	Notes
Web browser (local machine)	Worker instances	TCP/HTTPS	443	Required to access the file-level recovery browser running on a worker instance during the file-level recovery process.

From	To	Protocol	Port	Notes
Worker instances	AWS services	TCP/HTTPS	443	Required to perform data protection and disaster recovery operations.
		TCP/NFS	2049	Required to perform EFS indexing.

Required AWS Services

To perform backup and restore operations, worker instances must have outbound internet access to the following AWS services:

- [AWS Systems Manager \(SSM\)](#), including access to the *ec2messages* and *ssmmessages* endpoints
- [Amazon Simple Queue Service \(SQS\)](#)
- [Amazon Simple Storage Service \(S3\)](#)
- [Amazon Elastic Block Store \(EBS\)](#)
- [Amazon Kinesis Data Streams](#)

If you want worker instances to operate in a private environment, you must enable the private network deployment functionality and configure VPC endpoints for all subnets to which the worker instances will be connected. Otherwise, the instances will not be able to access all the listed services. For more information, see [Private Network Deployment](#).

How To Configure Worker Instance Settings

You can configure the following worker instance settings:

1. [Choose whether you want to deploy worker instances in the backup or production accounts.](#)
2. [Specify groups of network settings that Veeam Backup for AWS will use to deploy worker instances in specific AWS Regions.](#)
3. [Specify instance types that Veeam Backup for AWS will use to deploy worker instances in specific AWS Regions.](#)
4. [Assign AWS tags to worker instances to help you differentiate the instances.](#)

Worker Deployment Options

Veeam Backup for AWS provides the following options for deploying worker instances:

- [Worker deployment in the backup account](#)
- [Worker deployment in production accounts](#)

Worker Deployment in Backup Account

The backup account is an AWS account in which Veeam Backup for AWS deploys worker instances to perform operations with resources belonging to either the same or any other AWS account. By default, worker instances are deployed in the backup account to perform most backup and restore operations with EC2 instances:

- [EC2 image-level backup](#)
- [Entire EC2 instance restore from image-level backups](#)
- [EC2 volume-level restore from image-level backups](#)
- [EC2 file-level recovery](#)
- [EC2 backup retention tasks](#)
- [RDS archived backup](#)

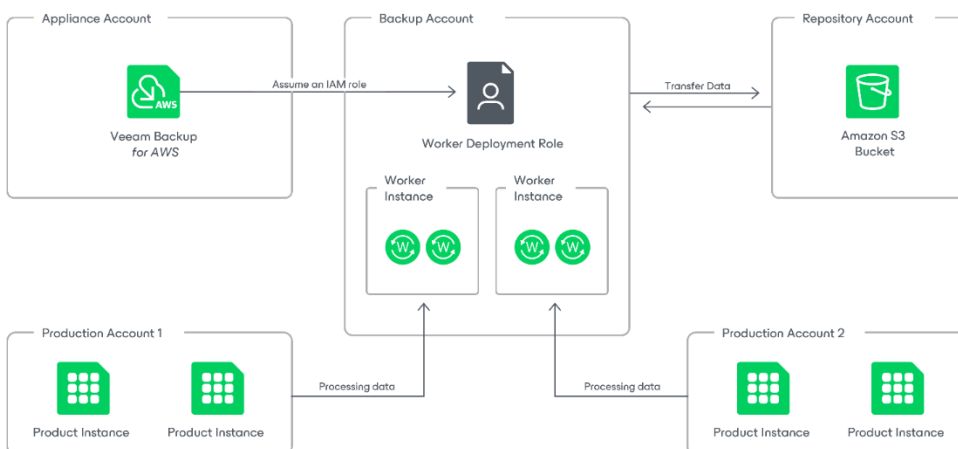
To deploy worker instances in the backup account, Veeam Backup for AWS employs a worker deployment role (service IAM role) that is then used to create temporary IAM roles to be attached to the deployed instances for communication with them. Out of the box, Veeam Backup for AWS uses the preconfigured *Default Backup Restore* role that has all the permissions required to perform data protection and disaster recovery operations. For more information on the *Default Backup Restore* role, see [Deploying Backup Appliance](#).

You can specify the worker deployment role in the worker instance settings as described in section [Managing Worker Instances](#). For more information on the IAM role permissions required to deploy worker instances in the backup account, see [Worker IAM Permissions](#).

How Worker Deployment in Backup Account Works

To perform a data protection or disaster recovery operation, Veeam Backup for AWS deploys worker instances in the following way:

1. Assumes a worker deployment role to deploy the worker instances.
2. Deploys in the backup account a worker instance for each AWS account to which the processed resources belongs, and attaches to this instance a temporary IAM role that will be used to communicate with it.
3. When the operation session completes, removes the worker instances and the temporary IAM role from AWS.



Worker Deployment in Production Accounts

Production accounts are AWS accounts in which Veeam Backup for AWS deploys worker instances to perform operations with processed AWS resources belonging to the same AWS accounts. By design, worker instances are deployed in production accounts to perform the following operations:

- [EFS file systems indexing](#)
- [RDS image-level backup](#)
- [RDS database restore from image-level backups](#)

Additionally, if you want to distribute workload across multiple AWS accounts and to manage resource costs for each account separately, you can instruct Veeam Backup for AWS to deploy worker instances in production accounts to perform the following operations:

- [EC2 image-level backup](#)
- [Entire EC2 instance restore from image-level backups](#)
- [EC2 volume-level restore from image-level backups](#)
- [EC2 file-level recovery from cloud-native snapshots](#)

To deploy worker instances in production accounts, Veeam Backup for AWS employs the following IAM roles:

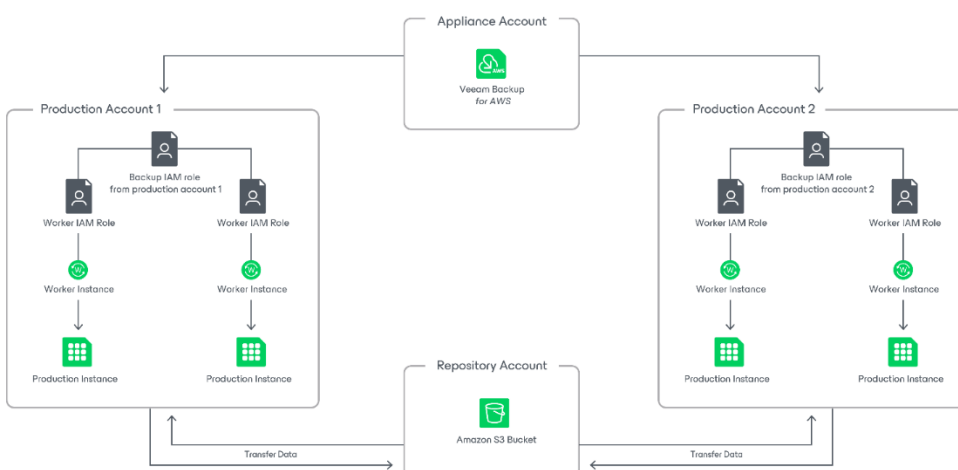
Role	Permissions	Settings
An IAM role that is used to perform an operation (that is, a backup or restore IAM role)	Depending on the operation, the role must be assigned additional permissions listed in either of the following sections: <ul style="list-style-type: none">• EFS Backup IAM Role Permissions• RDS Backup IAM Role Permissions• EC2 Backup IAM Role Permissions• EC2 Restore IAM Permissions	Depending on the operation, you must specify this IAM role in the backup policy or restore settings as described in either of the following sections: <ul style="list-style-type: none">• Creating EFS Backup Policies• Creating EC2 Backup Policies• Performing RDS Backup• Performing Entire EC2 Instance Restore• Performing Volume-Level Restore• Performing RDS Database Restore

Role	Permissions	Settings
An IAM role that is attached to the deployed worker instances and further used by Veeam Backup for AWS to communicate with the instances (that is, a worker IAM role)	Depending on the operation, the role must be assigned permissions listed in either of the following sections: <ul style="list-style-type: none"> Worker Deployment Role Permissions in Production Accounts FLR Worker IAM Role Permissions 	Depending on the operation, you must specify this IAM role when enabling worker deployment in production accounts in the backup policy or restore settings, as described in either of the following sections: <ul style="list-style-type: none"> Creating EFS Backup Policies Creating RDS Backup Policies Creating EC2 Backup Policies Performing RDS Database Restore Performing Entire EC2 Instance Restore Performing Volume-Level Restore Performing File-Level Recovery

How Worker Deployment in Production Accounts Works

To perform a data protection or disaster recovery operation, Veeam Backup for AWS deploys worker instances in the following way:

1. Assumes a backup or restore role to deploy the worker instances.
2. Deploys in each production account a worker instance for each AWS account to which the processed resources belongs, and attaches to this instance a worker IAM role that will be used to communicate with it.
3. When the operation session completes, removes the worker instance from AWS. Note that Veeam Backup for AWS does not remove the worker IAM role since it will be used for future backup and restore operations.



Worker Instance Locations

To minimize cross-region traffic charges and to speed up the data transfer, depending on the data protection and disaster recovery operation, Veeam Backup for AWS deploys worker instances in the following locations:

Operation	Worker Instance Location	Possibility to Deploy Worker Instances in Production Accounts	Default Worker Instance Type
Creating EC2 image-level backups	AWS Region in which a processed EC2 instance resides	Yes	<ul style="list-style-type: none"> c5.large – if the total EBS volume size is less than 1024 GB c5.2xlarge – if the total EBS volume size is 1024 GB - 16 TB c5.4xlarge – if the total EBS volume size is more than 16 TB
Restoring EC2 instances from image-level backups	AWS Region to which an EC2 instance is restored	Yes	
Restoring EC2 volumes from image-level backups	AWS Region to which the volumes of a processed EC2 instance are restored	Yes	
Performing health check for EC2 backups	AWS Region in which a backup repository with backed-up data resides	No	
Creating EC2 archived backups	AWS Region in which a standard backup repository with backed-up data resides	No	<ul style="list-style-type: none"> c5.2xlarge – if the total EBS volume size is less than 6 TB c5.4xlarge – if the total EBS volume size is more than 6 TB
Performing file-level recovery from image-level backups	AWS Region in which a backup repository with backed-up data resides	No	<ul style="list-style-type: none"> t3.medium
Performing file-level recovery from cloud-native snapshots and replicated snapshots	AWS Region in which a snapshot is located	<ul style="list-style-type: none"> No (if restoring to the original location) Yes (if restoring to a local machine) 	<ul style="list-style-type: none"> t3.medium

Operation	Worker Instance Location	Possibility to Deploy Worker Instances in Production Accounts	Default Worker Instance Type
Creating RDS image-level backups	AWS Region and VPC in which a processed PostgreSQL DB instance resides	Deploying worker instances in production accounts is the only available option.	<ul style="list-style-type: none"> • c5.large – if the total EBS volume size is less than 1024 GB • c5.2xlarge – if the total storage size is less than 6 TB • c5.4xlarge – if the total storage size is more than 6 TB
Restoring databases of PostgreSQL DB instances from image-level backups	AWS Region and VPC in which the restored DB instance databases will reside	Deploying worker instances in production accounts is the only available option.	
Performing health check for RDS backups	AWS Region in which a backup repository with backed-up data resides	No	
Creating RDS archived backups	AWS Region in which a standard backup repository with backed-up data resides	No	<ul style="list-style-type: none"> • c5.large – if the total EBS volume size is less than 1024 GB • c5.2xlarge – if the total EBS volume size is 1024 GB - 16 TB • c5.4xlarge – if the total EBS volume size is more than 16 TB
Performing EFS indexing	AWS Region, Availability Zone and VPC in which a file system has a mount target created	Deploying worker instances in production accounts is the only available option.	<ul style="list-style-type: none"> • t3.medium

Operation	Worker Instance Location	Possibility to Deploy Worker Instances in Production Accounts	Default Worker Instance Type
Applying retention policy settings to created restore points	AWS Region in which a backup repository with backed-up data resides	No	<ul style="list-style-type: none"> • c5.large – if the total size of backup files that must be deleted is 1-3 TB • c5.xlarge – if the total size of backup files that must be deleted is 3-6 TB • c5.2xlarge – if the total size of backup files that must be deleted is 6-13 TB • c5.4xlarge – if the total size of backup files that must be deleted is more than 13 TB

Worker instances are deployed based on worker configurations and profiles. For more information, see [Managing Worker Instances](#).

Additional Repositories and Tape Devices

Additional repositories and tape devices are any repositories where Veeam Backup & Replication keeps and stores copies of VM instance backups. For more information, see the Veeam Backup & Replication User Guide, sections [Backup Repository](#) and [Machines Backup to Tape](#).

Gateway Servers

The gateway server is an auxiliary backup infrastructure component that provides access from the backup server to the repositories. By default, the role of a gateway server is assigned to the backup server.

Gateway server caches data when you copy backups and restore application items, which helps you decrease the amount of traffic being sent over the network and reduce data transfer costs. For more information on caching data, see the Veeam Backup & Replication User Guide, section [Cache](#).

Protecting EC2 Instances

With Veeam Backup for AWS, you can perform the following operations to protect EC2 instances:

- Create cloud-native snapshots of EC2 instances and replicate these snapshots to any AWS Region within any AWS account.

A cloud-native snapshot of a EC2 instance includes point-in-time snapshots of EBS volumes attached to the processed instance. Snapshots of EBS volumes (also referred to as EBS snapshots) are taken using native [AWS capabilities](#).

- Create transactionally consistent backups using application-aware processing for Windows EC2 instances running VSS-aware applications.
- Create transactionally consistent snapshots using custom scripts to quiesce running applications for all processed EC2 instances.
- Create image-level backups of EC2 instances and keep them in Amazon Simple Storage Service (Amazon S3) for high availability, cost-effective and long-term storage.

An image-level backup captures the whole image of the processed EC2 instance (including instance configuration, OS data, application data and so on) at a specific point in time.

To protect EC2 instances, Veeam Backup for AWS runs [backup policies](#). A backup policy is a collection of settings that define the way backup operations are performed: what data to back up, which operations to perform, where to store backups, when to start the backup process, how to retain restore points, and so on.

Veeam Backup for AWS does not install agent software inside instances to back up EC2 instance data – it uses native AWS capabilities instead. During every backup session, Veeam Backup for AWS creates a cloud-native snapshot for each EC2 instance added to a backup policy. The cloud-native snapshot is further used to create a snapshot replica in another AWS Region or another AWS account and an image-level backup of the instance. For more information on how EC2 instance backup works, see [EC2 Backup](#).

Worker Deployment Considerations

Before you create EC2 backup policies, consider the following:

- By default, Veeam Backup for AWS deploys worker instances in the backup account and employs the worker deployment role (service IAM role). However, you can also instruct Veeam Backup for AWS to deploy worker instances in production accounts. For more information, see [Worker Deployment Options](#).
- To minimize cross-region traffic charges, Veeam Backup for AWS deploys worker instances in specific locations that depend on the data protection or disaster recovery operation. For more information, see [Worker Instance Locations](#).
- By default, Veeam Backup for AWS automatically chooses the default network settings of AWS Regions (if any) to deploy the worker instances. However, you can add worker configurations to define network settings for each region in which the worker instances will be deployed. For more information, see [Managing Worker Configurations](#).

How To Protect EC2 Instances

To create an EC2 backup policy, perform the following steps:

1. [Check limitations and prerequisites](#).
2. [Specify IAM roles to access AWS services and resources](#).

3. [Optional] Add backup repositories to store backed-up data.
4. [Optional] Configure worker instance settings to deploy workers while processing EC2 instance data.
5. [Optional] Configure global retention settings for obsolete snapshots and session records.
6. [Optional] Configure email notification settings for automated delivery of backup policy results and daily reports.
7. Complete the Add EC2 Policy wizard.

EC2 Backup

Veeam Backup for AWS performs EC2 backup in the following way:

1. Veeam Backup for AWS uses the [Amazon EC2 service](#) to create snapshots of EBS volumes that are attached to the processed EC2 instance.
2. EBS snapshots are assigned AWS tags upon creation. Keys and values of AWS tags contain encrypted metadata that helps Veeam Backup for AWS identify the related EBS snapshots and treat them as a single unit – a cloud-native snapshot.
3. If you enable snapshot replication for the backup policy, Veeam Backup for AWS copies cloud-native snapshots to the target AWS Region and AWS account specified in the backup policy settings.
4. If you enable image-level backup for the backup policy, Veeam Backup for AWS performs the following operations:
 - a. Deploys a worker instance in an AWS Region where the processed EC2 instance resides.

By default, Veeam Backup for AWS uses the default network settings of AWS Regions to deploy worker instances. However, you can add specific worker configurations. For more information on worker instances, see [Managing Worker Instances](#).

- b. Re-creates the EBS volumes from the cloud-native snapshot created at step 1 and attaches them to the worker instance. To increase backup performance, Veeam Backup for AWS can deploy worker instances with specific instance and volume types and the required number of copies of EBS volumes depending on the snapshot size.
- c. Reads data from the EBS volumes on the worker instance, transfers the data to a backup repository and stores it in the native Veeam format.

To reduce the amount of data read from EBS volumes, Veeam Backup for AWS uses the changed block tracking (CBT) mechanism: during incremental backup sessions, Veeam Backup for AWS compares the new cloud-native snapshot with the previous one and reads only those data blocks that have changed since the previous backup session. If CBT cannot be used, Veeam Backup for AWS reads all data from the re-created EBS volumes. For more information, see [Changed Block Tracking](#).

NOTE

By default, Veeam Backup for AWS compresses data saved to backup repositories. To learn how to encrypt data stored in backup repositories, see [Data Encryption](#).

- d. Removes the worker instance from Amazon EC2 when the backup session completes.
5. If you enable the [backup archiving mechanism](#), Veeam Backup for AWS performs the following operations:
 - a. Deploys a worker instance in an AWS Region where a backup repository storing backed-up data resides.
 - b. Retrieves data from the backup repository and transfers it to the archive backup repository.
 - c. Removes the worker instance from Amazon EC2 when the archive session completes.

Snapshot Chain

During every backup session, Veeam Backup for AWS creates a cloud-native snapshot for each instance added to the backup policy. The cloud-native snapshot itself is a collection of point-in-time snapshots that Veeam Backup for AWS takes using native AWS capabilities.

A sequence of cloud-native snapshots created during a set of backup sessions makes up a snapshot chain. Veeam Backup for AWS creates the snapshot chain in the following way:

1. During the first backup session, Veeam Backup for AWS creates a snapshot that contains all instance data and saves it in the AWS Region where the processed instance resides. This snapshot becomes a starting point in the snapshot chain.

The creation of the first snapshot may take significant time to complete, which depends on the number of volumes and their size, since Veeam Backup for AWS copies all the data of the instance volumes.


2. During subsequent backup sessions, Veeam Backup for AWS creates snapshots that contain only those data blocks that have changed since the previous backup session.

The creation of subsequent snapshots typically takes less time to complete, compared to the first snapshot in the chain. Note, however, that the completion time still depends on the amount of data being processed.

For more information on how incremental snapshots work, see [AWS Documentation](#).

Each cloud-native snapshot in the snapshot chain contains encrypted metadata. Metadata stores information about the protected instance and the backup policy that created the snapshot. Veeam Backup for AWS uses metadata to identify snapshots created by the Veeam backup service, to detect outdated snapshots, and to load the configuration of source instances during recovery operations, and so on.

Cloud-native snapshots act as independent restore points for backed-up instances. If you remove any snapshot, it will not break the snapshot chain – you will still be able to roll back instance data to any existing restore point.



The number of cloud-native snapshots kept in a snapshot chain is defined by retention policy settings. For more information, see [Snapshot Retention](#).

NOTE

Cloud-native snapshots created manually are not included into the snapshot chain. Therefore, these snapshots are not removed automatically according to retention policy settings. For information on how to remove them, see [Managing Backed-Up Data](#).

Snapshot Replica Chain

Snapshot replicas are copies of cloud-native snapshots that Veeam Backup for AWS creates during backup sessions. If you enable snapshot replication for a backup policy, Veeam Backup for AWS will make a copy of the initially created cloud-native snapshot and save it to the target AWS Region in the target AWS account specified in backup policy settings. Snapshot replicas created in the target AWS Region during a set of backup sessions make up a snapshot replica chain.


Veeam Backup for AWS creates and maintains the snapshot replica chain in the same way as the regular snapshot chain:

- The first snapshot replica of the processed instance becomes a starting point in the snapshot replica chain.
- Snapshot replicas created during subsequent backup sessions store only those data blocks that have changed since the previous backup session.

EC2 Snapshot Retention

For cloud-native snapshots and snapshot replicas, Veeam Backup for AWS retains the number of latest restore points defined in backup scheduling settings.

During every successful backup session, Veeam Backup for AWS creates a new restore point. If Veeam Backup for AWS detects that the number of restore points in the snapshot chain exceeds the retention limit, the earliest restore point is removed from the chain. However, some restore points can be retained longer than the period specified in the retention policy settings. For more information, see [CBT Impact on Snapshot Retention](#). For more information on the snapshot deletion process, see [AWS Documentation](#).



NOTE

Veeam Backup for AWS does not apply retention policy to cloud-native snapshots created manually. To learn how to remove these snapshots, see [Removing EC2 Backups and Snapshots](#).

Backup Chain


If you enable image-level backups for an EC2 backup policy, Veeam Backup for AWS creates a new backup in a backup repository during every backup session according to the backup policy schedule. A sequence of backups created during a set of backup sessions makes up a backup chain.

The backup chain includes backups of the following types:

- **Full** – a full backup stores a copy of the full EC2 image.
- **Incremental** – incremental backups store incremental changes of EC2 images.

To create a backup chain for an EC2 instance protected by a backup policy, Veeam Backup for AWS implements the forever forward incremental backup method:

1. During the first backup session, Veeam Backup for AWS copies the full EC2 image and creates a full backup in the backup repository. The full backup becomes a starting point in the backup chain.
2. During subsequent backup sessions, Veeam Backup for AWS copies only those data blocks that have changed since the previous backup session, and stores these data blocks to incremental backups in the backup repository. The content of each incremental backup depends on the content of the full backup and the preceding incremental backups in the backup chain.



Full and incremental backups act as restore points for backed-up EC2 instances that let you roll back instance data to the necessary state. To recover an EC2 instance to a specific point in time, the chain of backups created for the instance must contain a full backup and a set of incremental backups dependent on the full backup.

If some backup in the backup chain is missing, you will not be able to roll back to the necessary state. For this reason, you must not delete individual files from the backup repository manually. Instead, you must specify retention policy settings that will let you maintain the necessary number of backups in the backup repository. For more information, see [EC2 Backup Retention](#).

Changed Block Tracking

The changed block tracking (CBT) mechanism allows Veeam Backup for AWS to reduce the amount of data read from processed EBS volumes, and to increase the speed and efficiency of incremental backups:

- During a full backup session, Veeam Backup for AWS reads only written data blocks, while unallocated data blocks are filtered out.
- During an incremental backup session, Veeam Backup for AWS reads only those data blocks that have changed since the previous backup session.

To detect unallocated and changed data blocks, CBT relies on [EBS Direct APIs](#).

1. During the first (full) backup session, Veeam Backup for AWS [creates a cloud-native snapshot](#) of an EC2 instance. To do that, Veeam Backup for AWS sends API requests to access the content of the snapshot and to detect unallocated data blocks.
2. During subsequent sessions, new cloud-native snapshots are created. Veeam Backup for AWS sends API requests to access and to compare the content of the snapshot created during the previous backup session and the snapshot created during the current backup session. This allows Veeam Backup for AWS to detect data blocks that have changed since the previous backup session.

Limitations for Changed Block Tracking

Veeam Backup for AWS cannot use CBT for EC2 instances that reside in AWS Regions where EBS Direct APIs are not available.

If CBT cannot be used, Veeam Backup for AWS reads the whole content of processed EBS volumes and compares it with backed-up data that already exists in the backup repository. In this case, the completion time of incremental backups may occur to grow.

Archive Backup Chain


If you enable backup archiving for a backup policy, Veeam Backup for AWS creates a new backup in an archive backup repository during every archive session according to the backup policy schedule. A sequence of backups created during a set of archive sessions makes up an archive backup chain.

The archive backup chain includes backup files of the following types:

- **Full** – a full archive backup stores a copy of the full EC2 instance image.
- **Incremental** – incremental archive backups store incremental changes of the EC2 instance image.

To create an archive backup chain for a EC2 instance protected by a backup policy, Veeam Backup for AWS implements the forever forward incremental backup method:

1. During the first archive session, Veeam Backup for AWS detects backed-up data that is stored in the full backup and all incremental backups existing in the [standard backup chain](#), creates a full archive backup with all the data, and copies this backup to the archive backup repository. The full archive backup becomes a starting point in the archive chain.
2. During subsequent archive sessions, Veeam Backup for AWS checks the standard backup chain to detect data blocks that have changed since the previous archive session, creates incremental archive backups with only those changed blocks, and copies these backups to the archive backup repository. The content of each incremental archive backup depends on the content of the full archive backup and the preceding incremental archive backups in the archive backup chain.



Full and incremental archive backups act as restore points for backed-up EC2 instances that let you roll back instance data to the necessary state. To recover an EC2 instance to a specific point in time, the chain of backups created for the instance must contain a full archive backup and a set of incremental archive backups.

If some backup in the archive backup chain is missing, you will not be able to roll back to the necessary state. For this reason, you must not delete individual files from the archive backup repository manually. Instead, you must specify retention policy settings that will let you maintain the necessary number of backups in the archive backup repository. For more information, see [Retention Policy for Archived Backups](#).

EC2 Backup Retention


For image-level backups, Veeam Backup for AWS retains restore points for the number of days defined in backup scheduling settings as described in section [Creating EC2 Backup Policies](#).

To track and remove outdated restore points from a backup chain, Veeam Backup for AWS performs the following actions once a day:


1. Veeam Backup for AWS checks the configuration database to detect backup repositories that contain outdated restore points.
2. If an outdated restore point exists in a backup repository, Veeam Backup for AWS performs the following operations:
 - a. If the backup chain contains more than 20 backups and the total size of these backups exceeds 50 GB, deploys a worker instance in an AWS Region where the backup repository is located to process a retention task. Otherwise, Veeam Backup for AWS processes the task on the backup appliance.

By default, Veeam Backup for AWS uses the default network settings of AWS Regions to deploy worker instances. However, you can add specific worker configurations. For more information on worker instances, see [Managing Worker Instances](#).
 - b. Transforms the backup chain in the following way:


- i. Veeam Backup for AWS rebuilds the full backup to include in it data of the incremental backup that follows the full backup. To do that, Veeam Backup for AWS injects into the full backup data blocks from the earliest incremental backup in the chain. This way, a full backup 'moves' forward in the backup chain.



- ii. Veeam Backup for AWS removes the earliest incremental backup from the chain as redundant – this data has already been injected into the full backup.



3. Veeam Backup for AWS repeats step 2 for all other outdated restore points found in the backup chain until all the restore points are removed. As data from multiple restore points is injected into the rebuilt full backup, Veeam Backup for AWS ensures that the backup chain is not broken and that you will be able to recover your data when needed.



4. If the worker instance was deployed, Veeam Backup for AWS removes this worker instance from Amazon EC2 when the retention session completes.

NOTE

Consider the following:

- The retention task processes only 1 backup chain.
- Veeam Backup for AWS can process maximum 10 retention tasks at a time. If the number of retention tasks that must be processed on the backup appliance is more than the specified limit, the tasks exceeding this limit are queued.
- Each worker instance can process only one retention task at a time. Veeam Backup for AWS simultaneously can deploy maximum 10 worker instances that process retention tasks. If the number of retention tasks that must be processed on worker instances is more than the specified limit, the tasks exceeding this limit are queued.


CBT Impact on Snapshot Retention

If CBT is available, Veeam Backup for AWS does not remove the cloud-native snapshot used as a source for image-level backup from the snapshot chain until the next image-level backup session completes. Therefore, at some point you may discover that Veeam Backup for AWS ignores retention policy settings and keeps an additional restore point in the snapshot chain.

Consider the following example. You configured a backup policy to create cloud-native snapshots of your critical workloads 6 times a day (at 7:00 AM, 9:00 AM, 11:00 AM, 1:00 PM, 3:00 PM, and 5:00 PM) and to keep 2 daily snapshots in the snapshot chain. You also enabled creation of image-level backups 2 times a day (at 7:00 AM and 5:00 PM) and configured the retention policy settings to keep the backups in a backup repository for 7 days.

Veeam Backup for AWS will run the backup policy in the following way:

1. At 7:00 AM, the first backup session will create a cloud-native snapshot, and then will use this snapshot to create a full image-level backup.
2. From 9:00 AM to 3:00 PM, subsequent sessions will create only cloud-native snapshots.
 - a. After the backup session runs at 11:00 AM, the length of the snapshot chain (3 restore points) will exceed the retention limit (2 restore points). The earliest snapshot, however, will not be removed as it will be used to track changed data at 5:00 PM when the next image-level backup creation is scheduled.
 - b. After the backup session runs at 1:00 PM and 3:00 PM, Veeam Backup for AWS will remove the snapshots created at 9:00 AM and 11:00 AM. The length of the snapshot chain will remain 3 restore points.
3. At 5:00 PM, the backup session will create a new cloud-native snapshot. Veeam Backup for AWS will compare this snapshot with the one created at 7:00 AM to identify changed data blocks. After that, the backup session will create an incremental image-level backup based on the data obtained during the snapshot comparison.



4. After the snapshot comparison, Veeam Backup for AWS will apply the retention policy and remove from the chain the snapshot created at 7:00 AM (as it is no longer needed) and the snapshot created at 1:00 PM.



Retention Policy for Archived Backups


For archived backups, Veeam Backup for AWS retains restore points for the number of days defined in backup scheduling settings as described in section [EC2 Backup Policies](#).

To track and remove outdated restore points from an archive backup chain, Veeam Backup for AWS performs the following actions once a day:


1. Veeam Backup for AWS checks the configuration database to detect archive backup repositories that contain outdated restore points.
2. If an outdated restore point exists in a archive backup repository, Veeam Backup for AWS performs the following operations:
 - a. If the backup chain contains more than 20 backups and the total size of these backups exceeds 50 GB, deploys a worker instance in an AWS Region where the backup repository is located to process a retention task. Otherwise, Veeam Backup for AWS processes the task on the backup appliance.

By default, Veeam Backup for AWS uses the default network settings of AWS Regions to deploy worker instances. However, you can add specific worker configurations. For more information on worker instances, see [Managing Worker Instances](#).


- b. Transforms the archive backup chain in the following way:
 - i. Veeam Backup for AWS rebuilds the full archive backup to include there data of the incremental archive backup that follows the full archive backup. To do that, Veeam Backup for AWS injects into the full archive backup data blocks from the earliest incremental archive backup in the chain. This way, the full archive backup 'moves' forward in the archive backup chain.



- ii. Veeam Backup for AWS removes the earliest incremental archive backup from the chain as redundant – this data has already been injected into the full archive backup.



3. Veeam Backup for AWS repeats step 2 for all other outdated restore points found in the archive backup chain until all the restore points are removed. As data from multiple restore points is injected into the rebuilt full archive backup, Veeam Backup for AWS ensures that the archive backup chain is not broken and that you will be able to recover your data when needed.



4. If the worker instance was deployed, Veeam Backup for AWS removes this worker instance from Amazon EC2 when the retention session completes.

NOTE

Consider the following:

- The retention task processes only 1 backup chain.
- Veeam Backup for AWS can process maximum 10 retention tasks at a time. If the number of retention tasks that must be processed on the backup appliance is more than the specified limit, the tasks exceeding this limit are queued.
- Each worker instance can process only one retention task at a time. Veeam Backup for AWS simultaneously can deploy maximum 10 worker instances that process retention tasks. If the number of retention tasks that must be processed on worker instances is more than the specified limit, the tasks exceeding this limit are queued.

EC2 Restore

Veeam Backup for AWS offers the following restore options:

- [Instance restore](#) – restores an entire EC2 instance from a cloud-native snapshot, snapshot replica or an image-level backup. You can restore one or more EC2 instances at a time, to the original location or to a new location.
- [Volume restore](#) – restores EBS volumes attached to an EC2 instance from a cloud-native snapshot, snapshot replica or an image-level backup. You can restore EBS volumes to the original location or to a new location.
- [File-level recovery](#) – recovers individual files and folders of an EC2 instance from a cloud-native snapshot, snapshot replica or an image-level backup. You can download the necessary files and folders to a local machine, or restore the files and folders of the source EC2 instance to the original location.

You can restore EC2 instance data to the most recent state or to any available restore point.

EC2 Instance Restore

To restore EC2 instances from cloud-native snapshots, manual cloud-native snapshots and snapshot replicas, Veeam Backup for AWS uses native [AWS capabilities](#). To restore EC2 instances from image-level backups, Veeam Backup for AWS performs the following steps:

1. [This step applies only if you perform restore from an archived backup] Retrieves data from the archived restore point.
2. Deploys a worker instance in the AWS Region where the restored EC2 instance will reside.
3. Creates empty EBS volumes and attaches them to the worker instance.
The number of empty EBS volumes equals the number of EBS volumes attached to the backed-up EC2 instance.
4. Restores backed-up data to the empty EBS volumes on the worker instance.
5. Detaches EBS volumes with restored data from the worker instance.
6. Removes the worker instance from Amazon EC2.
7. Creates an EC2 instance in the specified location.
8. Attaches EBS volumes with restored data to the target EC2 instance.
9. [This step applies only if you perform restore to the original location] Powers off the source EC2 instance and removes it from Amazon EC2.

To learn how to restore an entire EC2 instance from a cloud-native snapshot, snapshot replica or an image-level backup, see [EC2 Restore](#).

Volume Restore

To restore EBS volumes from cloud-native snapshots, manual cloud-native snapshots and snapshot replicas, Veeam Backup for AWS uses native [AWS capabilities](#). To restore EBS volumes from image-level backups, Veeam Backup for AWS performs the following steps:

1. [This step applies only if you perform restore from an archived backup] Retrieves data from the archived restore point.

2. Deploys a worker instance in the AWS Region where the restored EBS volumes will reside.
3. Creates empty EBS volumes and attaches them to the worker instance.
The number of empty EBS volumes equals the number of volumes you selected to restore.
4. Restores backed-up data to the empty EBS volumes on the worker instance.
5. Detaches EBS volumes with restored data from the worker instance.
6. Removes the worker instance from Amazon EC2.

NOTE

Veeam Backup for AWS does not attach restored EBS volumes to any EC2 instances – the volumes are placed to the specified location as standalone EBS volumes.

To learn how to restore EBS volumes attached to an EC2 instance from a cloud-native snapshot, snapshot replica or an image-level backup, see [Performing Volume-Level Restore](#).

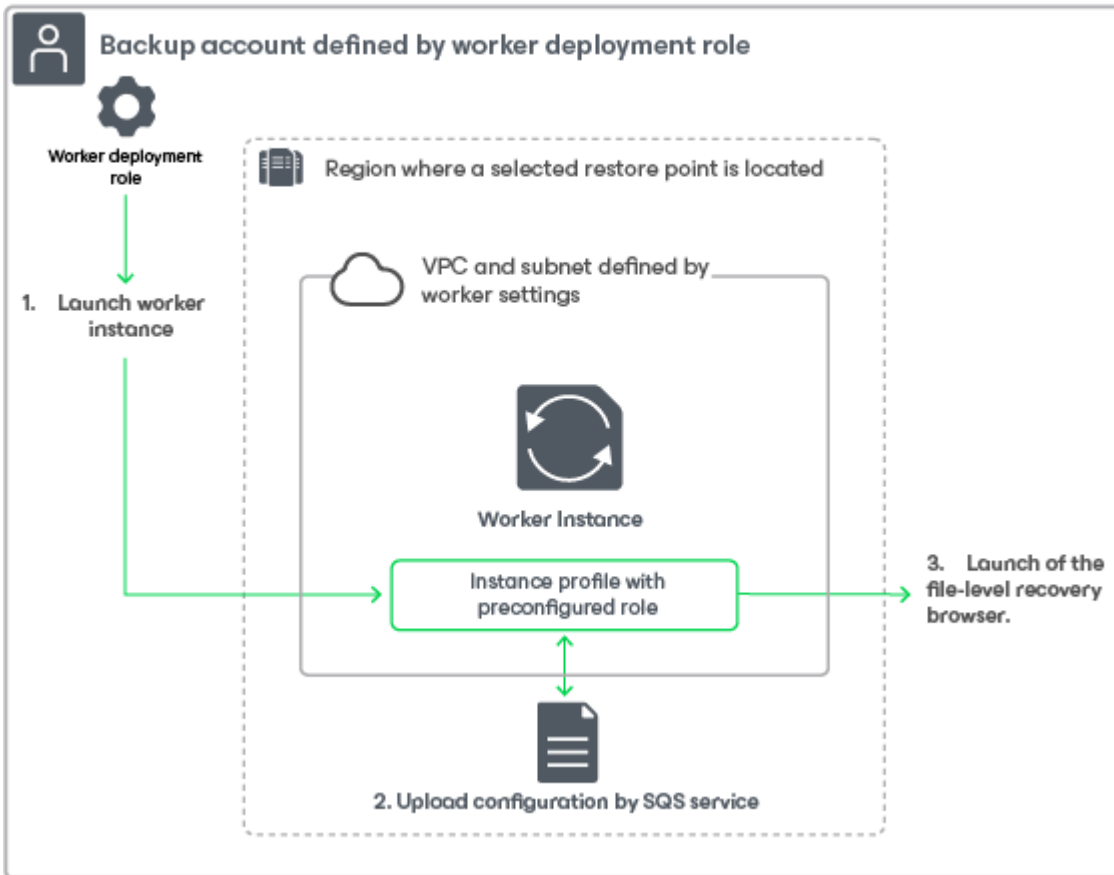
File-Level Recovery

To recover files and folders of a backed-up EC2 instance, Veeam Backup for AWS performs the following steps:

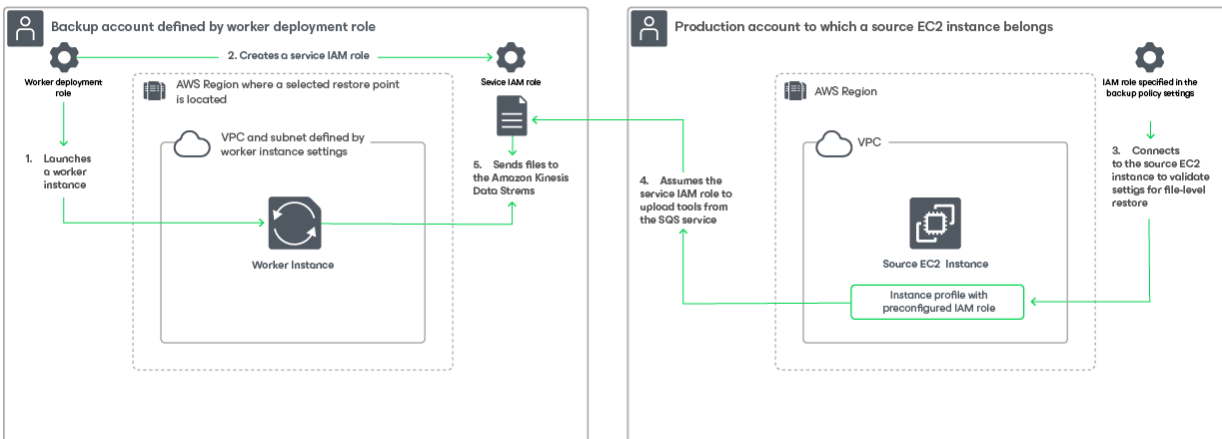
1. Deploys a worker instance in either of the following AWS Regions:
 - To restore files and folders from a cloud-native snapshot, manual cloud-native snapshots or a snapshot replicas, Veeam Backup for AWS deploys the worker instance in the AWS Region where the source EC2 snapshot or snapshot replica resides.
 - To restore files and folders from an image-level backup, Veeam Backup for AWS deploys the worker instance in the AWS Region where the backup repository with backed-up data resides.
2. Attaches and mounts EBS volumes of the EC2 instance to the worker instance.
[Applies to restore files and folders from an image-level backup] EBS volumes are not physically extracted from the backup – Veeam Backup for AWS emulates their presence on the worker instance. The source backup itself remains in the read-only state.
3. [This step applies only if you perform restore to the original location] Installs the Veeam restore tool on the source EC2 instance.
4. Launches the file-level recovery browser.
The file-level recovery browser displays the file system tree of the backed-up EC2 instance. In the browser, you select the necessary files and folders to restore.
5. Downloads the selected files and folders to the local machine.
6. [This step applies only if you perform restore to the original location] Restores the selected files and folders to the source EC2 instance, or downloads them to the local machine.
7. Unmounts and detaches EBS volumes of the backed-up EC2 instance from the worker instance.
8. [This step applies only if you perform restore to the original location] Removes the Veeam restore tool from the source EC2 instance.

- Removes the worker instance from Amazon EC2.

File-level recovery to a local machine



File-level recovery to the original location



To learn how to restore individual files and folders of an EC2 instance from a cloud-native snapshot or an image-level backup, see [Performing File-Level Recovery](#).

Protecting RDS Resources

With Veeam Backup for AWS, you can perform the following operations to protect RDS resources:

- Create cloud-native snapshots of RDS resources (DB instances and Amazon Aurora DB clusters) and replicate these snapshots to any AWS Region within any AWS account.

A cloud-native snapshot of a DB instance includes a storage volume snapshot of the instance. Snapshots of DB instances (also referred to as DB snapshots) are taken using native [AWS capabilities](#).

- Create image-level backups of PostgreSQL DB instances and keep them in Amazon Simple Storage Service (Amazon S3) for high availability, cost-effective and long-term storage.

An image-level backup captures the PostgreSQL databases of the processed DB instance.

To protect RDS resources, Veeam Backup for AWS runs backup policies. A backup policy is a collection of settings that define the way backup operations are performed: what data to back up, where to store backups, when to start the backup process, how to retain restore points, and so on.

Veeam Backup for AWS does not install agent software inside instances to back up RDS resource data – it uses native AWS capabilities instead. During every backup session, Veeam Backup for AWS creates a cloud-native snapshot for each RDS resource added to a backup policy. The cloud-native snapshot is further used to create a snapshot replica in another AWS Region or another AWS account and an image-level backup of the instance. For more information on how RDS resources backup works, see [RDS Backup](#).

Supported Applications

Veeam Backup for AWS supports backup of the following PostgreSQL versions:

- PostgreSQL 16
- PostgreSQL 15
- PostgreSQL 14
- PostgreSQL 13
- PostgreSQL 12

How To Protect RDS Resources

To create an RDS backup policy, perform the following steps:

1. [Check limitations and prerequisites](#).
2. [Specify IAM roles to access AWS services and resources](#).
3. [\[Optional\] Add backup repositories to store backed-up data](#).
4. [\[Optional\] Configure worker instance settings to deploy workers while processing DB instance data](#).
5. [\[Optional\] Configure global retention settings for obsolete snapshots and session records](#).
6. [\[Optional\] Configure email notification settings for automated delivery of backup policy results and daily reports](#).
7. [Complete the Add RDS Policy wizard](#).

RDS Backup

Veeam Backup for AWS performs RDS backup in the following way:

1. Veeam Backup for AWS creates a storage volume snapshot of the processed DB instance (that is, a DB snapshot) or of the processed Aurora DB cluster (that is, a DB cluster snapshot).

The snapshot is assigned AWS tags upon creation. Keys and values of AWS tags contain encrypted metadata that helps Veeam Backup for AWS identify the related snapshot. For the Aurora DB cluster metadata saved in AWS tags also contains information on every DB instance launched in the cluster.

2. If you enable snapshot replication for the backup policy, Veeam Backup for AWS copies the snapshot to the target AWS Region and AWS account specified in the backup policy settings.
3. If you enable image-level backup for the backup policy, Veeam Backup for AWS performs the following operations:
 - a. Deploys a worker instance in an AWS Region in which the processed DB instance resides in an AWS account to which the instance belongs – that is, the production AWS account.

By default, Veeam Backup for AWS selects the most appropriate network settings of AWS Regions in production accounts. However, you can add specific worker configurations. For more information on worker instances, see [Managing Worker Configurations](#).
 - b. Creates 2 security groups that are associated with the source DB instance and the worker instance to allow direct network traffic between them. The security group associated with the source instance allows inbound traffic through opened on the instance port only from the worker instance, whereas the security group associated with the worker instance allows outbound traffic through opened on the instance port only to the source instance.
 - c. Uses PostgreSQL capabilities to dump out PostgreSQL databases.
 - d. Uses the worker instance to retrieve dumps, triggers, stored procedures and transfers the retrieved data to the target backup repository and stores the data in the native Veeam format.
 - e. Removes the worker instance and 2 created security groups from AWS when the backup session completes.
5. If you enable the [backup archiving mechanism](#), Veeam Backup for AWS performs the following operations:
 - a. Deploys a worker instance in an AWS Region where a backup repository storing backed-up data resides in the [backup account](#).
 - b. Retrieves data from the backup repository and transfers it to the archive backup repository.
 - c. Removes the worker instance from Amazon EC2 when the archive session completes.

Snapshot Chain

During every backup session, Veeam Backup for AWS creates a cloud-native snapshot for each instance added to the backup policy. The cloud-native snapshot is taken using native AWS capabilities.

A sequence of cloud-native snapshots created during a set of backup sessions makes up a snapshot chain. Veeam Backup for AWS creates the snapshot chain in the following way:

1. During the first backup session, Veeam Backup for AWS creates a snapshot that contains all instance data and saves it in the AWS Region where the processed instance resides. This snapshot becomes a starting point in the snapshot chain.

The creation of the first snapshot may take significant time to complete since Veeam Backup for AWS copies the whole image of the instance.


2. During subsequent backup sessions, Veeam Backup for AWS creates snapshots that contain only those data blocks that have changed since the previous backup session.

The creation of subsequent snapshots typically takes less time to complete, compared to the first snapshot in the chain. Note, however, that the completion time still depends on the amount of data being processed.

For more information on how incremental snapshots work, see [AWS Documentation](#).

Each cloud-native snapshot in the snapshot chain contains encrypted metadata. Metadata stores information about the protected instance and the backup policy that created the snapshot. Veeam Backup for AWS uses metadata to identify snapshots created by the Veeam backup service, to detect outdated snapshots, and to load the configuration of source instances during recovery operations, and so on.

Cloud-native snapshots act as independent restore points for backed-up instances. If you remove any snapshot, it will not break the snapshot chain – you will still be able to roll back instance data to any existing restore point.



The number of cloud-native snapshots kept in a snapshot chain is defined by retention policy settings. For more information, see [RDS Snapshot Retention](#).

NOTE

Cloud-native snapshots created manually are not included into the snapshot chain. Therefore, these snapshots are not removed automatically according to retention policy settings. For information on how to remove them, see [Managing Backed-Up Data](#).

Snapshot Replica Chain

Snapshot replicas are copies of cloud-native snapshots that Veeam Backup for AWS creates during backup sessions. If you enable snapshot replication for a backup policy, Veeam Backup for AWS will make a copy of the initially created cloud-native snapshot and save it to the target AWS Region in the target AWS account specified in backup policy settings. Snapshot replicas created in the target AWS Region during a set of backup sessions make up a snapshot replica chain.

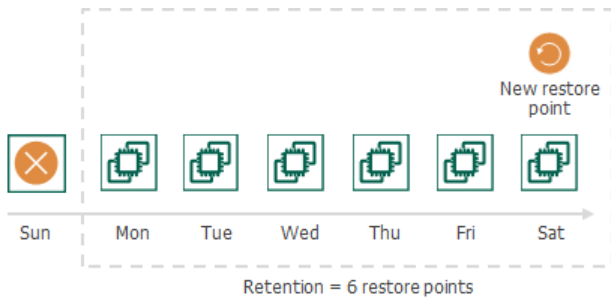
Veeam Backup for AWS creates and maintains the snapshot replica chain in the same way as the regular snapshot chain:

- The first snapshot replica of the processed instance becomes a starting point in the snapshot replica chain.
- Snapshot replicas created during subsequent backup sessions store only those data blocks that have changed since the previous backup session.

RDS Snapshot Retention

For cloud-native snapshots and snapshot replicas, Veeam Backup for AWS retains the number of latest restore points defined in backup scheduling settings.

During every successful backup session, Veeam Backup for AWS creates a new restore point. If Veeam Backup for AWS detects that the number of restore points in the snapshot chain exceeds the retention limit, the earliest restore point is removed from the chain. For more information on the snapshot deletion process, see [AWS Documentation](#).



NOTE

Veeam Backup for AWS does not apply retention policy to cloud-native snapshots created manually. To learn how to remove them, see [Managing Backed-Up Data](#).

Backup Chain

The forever forward incremental backup method is not implemented for DB instances – during every backup session Veeam Backup for AWS creates a full backup in the regular backup chain.

Each RDS backup in the backup chain contains encrypted metadata that stores information about the protected DB instance, the backup policy that created the backup, as well as the date, time and configured retention settings. Veeam Backup for AWS uses metadata to identify outdated backups, to retrieve information on the source instance configuration during recovery operations, and so on.

RDS backups act as independent restore points for backed-up DB instances. If you remove any backup, it will not break the backup chain – you will still be able to roll back data to any existing restore point.

The period of time during which RDS backups are kept in the backup chain is defined by retention policy settings. For more information, see [RDS Backup Retention](#).

Archive Backup Chain


If you enable backup archiving for a backup policy, Veeam Backup for AWS creates a new backup in an archive backup repository during every archive session according to the backup policy schedule. A sequence of backups created during a set of archive sessions makes up an archive backup chain.

The archive backup chain includes backup files of the following types:

- **Full** – a full archive backup stores a copy of the full DB instance image.
- **Incremental** – incremental archive backups store incremental changes of the DB instance image.

To create an archive backup chain for a DB instance protected by a backup policy, Veeam Backup for AWS implements the forever forward incremental backup method:

1. During the first archive session, Veeam Backup for AWS detects backed-up data that is stored in the full backup and all incremental backups existing in the [standard backup chain](#), creates a full archive backup with all the data, and copies this backup to the archive backup repository. The full archive backup becomes a starting point in the archive chain.
2. During subsequent archive sessions, Veeam Backup for AWS checks the standard backup chain to detect data blocks that have changed since the previous archive session, creates incremental archive backups with only those changed blocks, and copies these backups to the archive backup repository. The content of each incremental archive backup depends on the content of the full archive backup and the preceding incremental archive backups in the archive backup chain.



Full and incremental archive backups act as restore points for backed-up DB instances that let you roll back instance data to the necessary state. To recover a DB instance to a specific point in time, the chain of backups created for the instance must contain a full archive backup and a set of incremental archive backups.

If some backup in the archive backup chain is missing, you will not be able to roll back to the necessary state. For this reason, you must not delete individual files from the archive backup repository manually. Instead, you must specify retention policy settings that will let you maintain the necessary number of backups in the archive backup repository. For more information, see [Retention Policy for Archived Backups](#).

RDS Backup Retention

The forever forward incremental backup method is not implemented for DB instances – during every backup session Veeam Backup for AWS creates a full backup in the regular backup chain. If Veeam Backup for AWS detects an outdated restore point in a backup repository, it removes this restore point from the backup chain.

Retention Policy for Archived Backups

For archived backups, Veeam Backup for AWS retains restore points for the number of days defined in [backup scheduling settings](#).


To track and remove outdated restore points from an archive backup chain, Veeam Backup for AWS performs the following actions once a day:

1. Veeam Backup for AWS checks the configuration database to detect archive backup repositories that contain outdated restore points.
2. If an outdated restore point exists in a archive backup repository, Veeam Backup for AWS performs the following operations:
 - a. If the total size of backups that must be deleted is more than 50 GB, deploys a worker instance in an AWS Region where the backup repository is located to process a retention task. Otherwise, Veeam Backup for AWS processes the task on the backup appliance.


By default, Veeam Backup for AWS uses the default network settings of AWS Regions to deploy worker instances. However, you can add specific worker configurations. For more information on worker instances, see [Managing Worker Instances](#).

b. Transforms the archive backup chain in the following way:

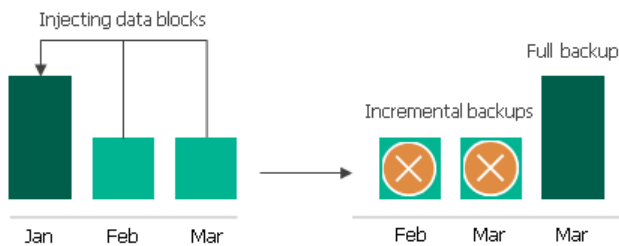
- i. Veeam Backup for AWS rebuilds the full archive backup to include the data of the incremental archive backup that follows the full archive backup. To do that, Veeam Backup for AWS injects into the full archive backup data blocks from the earliest incremental archive backup in the chain. This way, the full archive backup 'moves' forward in the archive backup chain.



- ii. Veeam Backup for AWS removes the earliest incremental archive backup from the chain as redundant – this data has already been injected into the full archive backup.



3. Veeam Backup for AWS repeats step 2 for all other outdated restore points found in the archive backup chain until all the restore points are removed. As data from multiple restore points is injected into the rebuilt full archive backup, Veeam Backup for AWS ensures that the archive backup chain is not broken and that you will be able to recover your data when needed.



4. If the worker instance was deployed, Veeam Backup for AWS removes this worker instance from Amazon EC2 when the retention session completes.

NOTE

Consider the following:

- The retention task processes only 1 backup chain.
- Veeam Backup for AWS can process maximum 10 retention tasks at a time. If the number of retention tasks that must be processed on the backup appliance is more than the specified limit, the tasks exceeding this limit are queued.
- Each worker instance can process only one retention task at a time. Veeam Backup for AWS simultaneously can deploy maximum 10 worker instances that process retention tasks. If the number of retention tasks that must be processed on worker instances is more than the specified limit, the tasks exceeding this limit are queued.

RDS Restore

Veeam Backup for AWS offers the following restore operations:

- [RDS instance restore](#) – restores an entire DB instance or an Aurora DB cluster from a cloud-native snapshot or snapshot replica.
- [Database restore](#) – restores specific databases of a PostgreSQL DB instance from an image-level backup.

You can restore EC2 instance data to the most recent state or to any available restore point.

RDS Instance Restore

To restore a DB instance from a snapshot, Veeam Backup for AWS performs the following steps using native [AWS capabilities](#):

1. Creates a DB instance in the specified location.
2. Modifies the configuration setting values of the created DB instance.
3. Restores backed-up databases to the restored DB instance.

To restore an Aurora DB cluster from a snapshot, Veeam Backup for AWS performs the following steps using native [AWS capabilities](#):

1. Creates an Aurora DB cluster in the specified location.
2. Restores backed-up databases to the created Aurora DB cluster.
3. Modifies the configuration setting values of the created Aurora DB cluster.
4. In the created Aurora DB cluster, creates all backed-up DB instances (when restoring to the original location) or creates the primary DB instance (when restoring to a new location).
5. Modifies the configuration setting values of each created DB instance.

To learn how to restore a DB instance or an Aurora DB cluster from a cloud-native snapshot or snapshot replica, see [RDS Restore](#).

Database Restore

To restore a database of a PostgreSQL DB instance from an image-level backup, Veeam Backup for AWS performs the following steps:

1. [This step applies only if you perform restore from an archived backup] Retrieves data from the archived restore point.
2. Deploys a worker instance in an AWS Region in which DB instance that will host the restored databases resides in an AWS account to which the instance belongs – that is, the production AWS account.

By default, Veeam Backup for AWS selects the most appropriate network settings of AWS Regions in production accounts. However, you can add specific worker configurations. For more information on worker instances, see [Managing Worker Configurations](#).

2. Creates 2 security groups that are associated with the target DB instance and worker instance to allow direct network traffic between the resources. The security group associated with the target instance allows inbound traffic through opened on the instance port only from the worker instance, whereas the security group associated with the worker instance allows outbound traffic through opened on the instance port only to the target instance.
3. Uses the worker instance to retrieve dumps, triggers and stored procedures from a backup file stored in the target backup repository.
4. Uses PostgreSQL capabilities to restore the PostgreSQL databases to the specified DB instance.
5. Removes the worker instance and 2 created security groups from AWS.

To learn how to restore databases of a DB instance from an image-level backup, see [RDS Restore](#).

NOTE

Due to [AWS technical limitations](#), the next run of the backup policy protecting the source DB instance may take more time to complete in case the restore operation completes successfully, depending on the specified target AWS Region and target DB instance. For more information, see [Performing RDS Database Restore](#).

Protecting DynamoDB Tables

With Veeam Backup for AWS, you can perform the following operations to protect DynamoDB tables:

- Create cloud-native backups of DynamoDB tables and store them in any backup vault in the source AWS Region.

An Amazon DynamoDB backup captures the whole image of the DynamoDB table at a specific point of time. DynamoDB backups are taken using native [AWS capabilities](#).

- Create backup copies of DynamoDB tables and store them in any AWS Region within the same AWS account.

By default, DynamoDB backups are stored only in the AWS Region where the processed resources reside. For enhanced data safety, you can instruct Veeam Backup for AWS to create copies of these backups and store them in any other AWS Region within the same AWS account.

To protect DynamoDB tables, Veeam Backup for AWS runs backup policies. A backup policy is a collection of settings that define the way backup operations are performed: what data to back up, where to store backups, when to start the backup process, how to retain restore points, and so on.

Veeam Backup for AWS does not install agent software inside instances to back up DynamoDB tables – it uses native AWS capabilities instead. During every backup session, Veeam Backup for AWS creates a cloud-native backup for each table added to a backup policy. The cloud-native backup is further used to create a backup copy in another AWS Region. For more information on how DynamoDB table backup works, see [DynamoDB Backup](#).

Supported DynamoDB Table Properties

Veeam Backup for AWS allows you to back up and restore the following table properties:

Property	Description	Ability to Change Property During Restore to New Location
Table name	Name of a DynamoDB table.	Yes
Partition key	First attribute of the primary key.	No
Sort key	Second attribute of the primary key.	No
Global secondary index (GSI) and local secondary index (LSI)	Additional indexes that provide efficient access to the table data.	No
Table class	Defines how often the table data is accessed.	Yes
Capacity mode	Defines how read/write operations are charged and managed.	Yes
Provisioned read/write capacity units	Read/write throughput for the table and its indexes.	Yes

Property	Description	Ability to Change Property During Restore to New Location
Maximum throughput capacity	Maximum read/write throughput for the on-demand table.	No
Tags	Table identifiers.	No
Deletion protection	Defines whether the table is protected against accidental deletion.	Yes
Server-side encryption	Defines the key used for data-at-rest encryption.	Yes
Point-in-time recovery (PITR)	Defines whether the table data can be restored to any point in time during the last 35 days.	Yes
DynamoDB Time to Live (TTL)	Attribute name with a timestamp that determines when the table items are no longer needed.	No

IMPORTANT

Veeam Backup for AWS does not support the following:

- CloudWatch alarms
- Resource-based policies
- DynamoDB global table feature
- Adjusted provisioned throughput capacity provided by Amazon DynamoDB auto scaling
- Item-level modifications captured by Amazon Kinesis Data Streams
- Time-ordered sequences of item-level modifications captured by Amazon DynamoDB Streams
- Restore of the configured CloudWatch Contributor Insights diagnostic tool

How To Protect DynamoDB Tables

To create a DynamoDB policy, perform the following steps:

1. [Check limitations and prerequisites.](#)
2. [Specify IAM roles to access AWS services and resources.](#)
3. [\[Optional\] Configure global retention settings for obsolete session records.](#)
4. [\[Optional\] Configure email notification settings for automated delivery of backup policy results and daily reports.](#)
5. [Complete the Add DynamoDB Policy wizard.](#)

DynamoDB Backup

Veeam Backup for AWS performs DynamoDB backup in the following way:

1. Veeam Backup for AWS uses the [AWS Backup service](#) to create a cloud-native backup of the DynamoDB table, and saves this backup to the specified backup vault in the same AWS Region in which the source table resides.

The backup is assigned AWS tags upon creation. Keys and values of AWS tags contain encrypted metadata that helps Veeam Backup for AWS identify the related table backup.

2. If you configure the DynamoDB backup policy to copy backup files to another AWS Region, Veeam Backup for AWS copies the created backup to the target AWS Region in the same AWS account.

Backup Chain

During every backup session, Veeam Backup for AWS creates a new cloud-native backup for each DynamoDB table added to the backup policy. To create the backup, Veeam Backup for AWS uses the [AWS Backup service](#). A sequence of cloud-native backups created during a set of backup sessions makes up a backup chain.

DynamoDB backups



Each DynamoDB backup in the backup chain contains encrypted metadata. Metadata stores information about the protected table, the backup policy that created the backup, and the date, time and applied retention settings. Veeam Backup for AWS uses metadata to identify outdated backups, to load the configuration of source tables during recovery operations, and so on.

NOTES

- Due to [AWS Backup service limitations](#), during every backup session, Veeam Backup for AWS creates a full backup in the regular backup chain.
- DynamoDB backups created manually are not included into the DynamoDB backup chain. Therefore, these backups are not removed automatically according to retention policy settings. To learn how to remove them, see [Removing DynamoDB Backups Created Manually](#).

DynamoDB backups act as independent restore points for backed-up tables. If you remove any backup, it will not break the DynamoDB backup chain – you will still be able to roll back table data to any existing restore point. The period of time during which DynamoDB backups are kept in the DynamoDB backup chain is defined by retention policy settings. For more information, see [DynamoDB Backup Retention](#).

DynamoDB Backup Copy Chain


If you enable backup copying for a backup policy, Veeam Backup for AWS will make a copy of the initially created full DynamoDB backup and save it to the target AWS Region specified in the backup policy settings. In the target AWS Region, backup copies created during a set of backup sessions make up a backup copy chain.

Veeam Backup for AWS creates and maintains a DynamoDB backup copy chain in the same way as a regular DynamoDB backup chain – during every backup copy session Veeam Backup for AWS creates a full backup in the backup copy chain.

DynamoDB Backup Retention

For DynamoDB backups, Veeam Backup for AWS retains restore points for the period of time specified in [backup scheduling settings](#).

During every successful backup session, Veeam Backup for AWS creates a restore point and saves the date, time and applied retention settings in the restore point metadata. If Veeam Backup for AWS detects that the period of time for which the restore point was stored exceeds the period specified in the retention settings, it automatically removes the restore point from the DynamoDB chain. You can also remove unnecessary DynamoDB backups manually as described in section [Removing DynamoDB Backups](#).



NOTE

Veeam Backup for AWS does not apply retention policy to DynamoDB backups created manually. To learn how to remove them, see [Removing DynamoDB Backups Created Manually](#).

DynamoDB Restore

IMPORTANT

You can restore a DynamoDB table only to the same AWS account to which the source table belongs.

To restore a DynamoDB table from a backup, Veeam Backup for AWS performs the following steps using native [AWS capabilities](#):

1. Creates a table in the specified location.
2. Restores backed-up data (items and attributes) to the restored table.
3. Modifies the configuration setting values of the created DynamoDB table.

To learn how to restore a DynamoDB table from a DynamoDB backup or a backup copy, see [DynamoDB Restore](#).

Protecting Redshift Clusters

Veeam Backup for AWS allows you to create cloud-native backups of Redshift clusters and store them in any backup vault in the source AWS Region. An Amazon Redshift backup captures the whole image of the Redshift cluster at a specific point of time. Redshift backups are taken using native [AWS capabilities](#).

To protect Redshift clusters, Veeam Backup for AWS runs backup policies. A backup policy is a collection of settings that define the way backup operations are performed: what data to back up, where to store backups, when to start the backup process, how to retain restore points, and so on.

Veeam Backup for AWS does not install agent software inside instances to back up Redshift clusters – it uses native AWS capabilities instead. During every backup session, Veeam Backup for AWS creates a cloud-native backup for each cluster added to a backup policy. For more information on how Redshift clusters backup works, see [Redshift Backup](#).

Supported Redshift Cluster Properties

Veeam Backup for AWS allows you to back up and restore the following cluster properties:

Property	Description	Ability to Change Property During Restore to New Location
Cluster ID	Unique identifier of the Redshift cluster.	Yes
Node type	Type of the nodes for the Redshift cluster.	Yes
Number of nodes	Total number of compute nodes for the Redshift cluster.	Yes
Associated IAM roles	List of IAM roles associated with the Redshift cluster.	Yes
Default role ARN	ARN of the default IAM role associated with the Redshift cluster.	Yes
Admin user password	Defines whether the admin password is managed by AWS Secrets Manager.	No
Secret Manager KMS key ID	ID of the custom KMS key used to encrypt the secret.	Yes
VPC	VPC to which the Redshift cluster is connected.	Yes

Property	Description	Ability to Change Property During Restore to New Location
Subnet group	Subnet group in which the Redshift cluster is launched.	Yes
Availability Zone	Availability Zone where the Redshift cluster resides.	Yes
Public access	Defines whether the Redshift cluster is accessible outside the VPC to which the cluster is connected.	Yes
Enhanced VPC routing	Defines whether network traffic is routed between the Redshift cluster and the data repositories through a Redshift-managed VPC endpoint.	No
Port	Port used to access the Redshift cluster.	Yes
Parameter group	Parameter group associated with the Redshift cluster.	Yes
Cluster relocation	Defines whether the Redshift cluster can be moved to another Availability Zone.	No
Encryption key	AWS KMS key that is used to encrypt the Redshift cluster data.	Yes
Tags	User-defined labels assigned to the Redshift cluster.	No
Backup and maintenance	Defines whether daily automatic backup is scheduled and whether a weekly maintenance window is configured for the Redshift cluster.	No

IMPORTANT

Veeam Backup for AWS does not support the following:

- Multi-AZ deployment mode of the Redshift clusters
- HSM encryption scheme of the Redshift cluster
- Elastic IP assigned to the network interface of the Redshift cluster

How To Protect Redshift Clusters

To create a Redshift policy, perform the following steps:

1. [Check limitations and prerequisites.](#)
2. [Specify IAM roles to access AWS services and resources.](#)
3. [\[Optional\] Configure global retention settings for obsolete session records.](#)
4. [\[Optional\] Configure email notification settings for automated delivery of backup policy results and daily reports.](#)
5. [Complete the Add Redshift Policy wizard.](#)


Redshift Backup

Veeam Backup for AWS performs Redshift backup in the following way:

1. Veeam Backup for AWS uses the [AWS Backup service](#) to create a cloud-native backup of the Redshift cluster, and saves this backup to the specified backup vault in the same AWS Region in which the source cluster resides.
2. The backup is assigned AWS tags upon creation. Keys and values of AWS tags contain encrypted metadata that helps Veeam Backup for AWS identify the related cluster backup.

Backup Chain

During every backup session, Veeam Backup for AWS creates a new cloud-native backup for each Redshift cluster added to the backup policy. To create the backup, Veeam Backup for AWS uses the [AWS Backup service](#). A sequence of cloud-native backups created during a set of backup sessions makes up a backup chain.



Each Redshift backup in the backup chain contains encrypted metadata. Metadata stores information about the protected cluster, the backup policy that created the backup, and the date, time and applied retention settings. Veeam Backup for AWS uses metadata to identify outdated backups, to load the configuration of source cluster during recovery operations, and so on.

NOTES

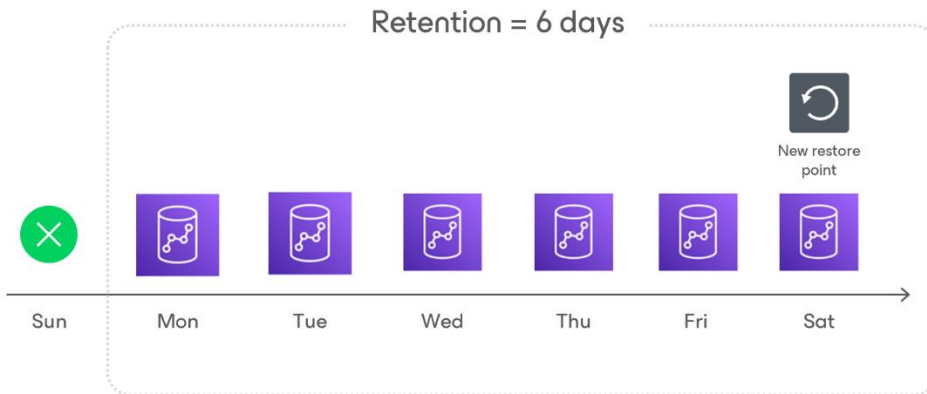
- Due to [AWS Backup service limitations](#), during every backup session, Veeam Backup for AWS creates a full backup in the regular backup chain.
- Redshift backups created manually are not included into the Redshift backup chain. Therefore, these backups are not removed automatically according to retention policy settings. To learn how to remove them, see [Removing Redshift Backups Created Manually](#).

Redshift backups act as independent restore points for backed-up clusters. If you remove any backup, it will not break the Redshift backup chain – you will still be able to roll back cluster data to any existing restore point. The period of time during which Redshift backups are kept in the Redshift backup chain is defined by retention policy settings. For more information, see [Redshift Backup Retention](#).

Redshift Backup Retention

For Redshift backups, Veeam Backup for AWS retains restore points for the period of time specified in [backup scheduling settings](#).

During every successful backup session, Veeam Backup for AWS creates a restore point and saves the date, time and applied retention settings in the restore point metadata. If Veeam Backup for AWS detects that the period of time for which the restore point was stored exceeds the period specified in the retention settings, it automatically removes the restore point from the Redshift chain. You can also remove unnecessary Redshift backups manually as described in section [Removing Redshift Backups](#).



NOTE

Veeam Backup for AWS does not apply retention policy to Redshift backups created manually. For learn how to remove them, see [Removing Redshift Backups Created Manually](#).

Redshift Restore

IMPORTANT

You can restore a Redshift cluster only to the same AWS account to which the source cluster belongs and the same AWS Region where the source cluster resides.

To restore a Redshift cluster from a backup, Veeam Backup for AWS performs the following steps using native [AWS capabilities](#):

1. Creates a cluster in the specified location.
2. Modifies the configuration setting values of the created Redshift cluster.
3. Restores backed-up databases to the restored Redshift clusters.

To learn how to restore a Redshift cluster from a Redshift backup, see [Redshift Restore](#).

Protecting EFS File Systems

With Veeam Backup for AWS, you can perform the following operations to protect EFS file systems:

- Create cloud-native backups of EFS file systems and store them in any backup vault in the source AWS Region.

An Amazon EFS file system backup captures the whole image of the EFS file system (including file system configuration, files, directories and so on) at a specific point of time. EFS backups are taken using native [AWS capabilities](#).
- Create backup copies of EFS file systems and store them in any AWS Region within the same AWS account.

By default, EFS backups are stored only in the AWS Region where the processed resources reside. For enhanced data safety, you can instruct Veeam Backup for AWS to create copies of these backups and store them in any other AWS Region within the same AWS account. You can also combine the backup copy functionality with various [data recovery options](#) to migrate file system data between AWS Regions.

To protect EFS file systems, Veeam Backup for AWS runs backup policies. A backup policy is a collection of settings that define the way backup operations are performed: what data to back up, where to store backups, when to start the backup process, how to retain restore points, and so on.

Veeam Backup for AWS does not install agent software inside instances to back up EFS file systems – it uses native AWS capabilities instead. During every backup session, Veeam Backup for AWS creates a cloud-native backup for each file system added to a backup policy. The cloud-native backup is further used to create a backup copy in another AWS Region.

EFS Indexing

Veeam Backup for AWS allows you to perform indexing of the processed EFS file systems – that is, to perform EFS file-level recovery operations without specifying the exact paths to the necessary files and to restore files using different restore points during one restore session. While performing EFS indexing of a file system, Veeam Backup for AWS creates a catalog of all files and directories (an index) and saves the index to a backup repository. This index is further used to reproduce the file system structure and to enable browsing and searching for specific files within an EFS backup. For more information on how EFS file systems backup works, see [EFS Backup](#).

To allow Veeam Backup for AWS to perform indexing of the processed EFS file systems, this functionality must be enabled in the [backup policy settings](#).

How To Protect EFS File Systems

To create an EFS backup policy, perform the following steps:

1. [Check limitations and prerequisites](#).
2. [Specify IAM roles to access AWS services and resources](#).
3. [\[Optional\] Add backup repositories to store backed-up data](#).
4. [\[Optional\] Configure worker instance settings to deploy workers while performing indexing of the processed EFS file systems](#).
5. [\[Optional\] Configure email notification settings for automated delivery of backup policy results and daily reports](#).
6. [Complete the Add EFS Policy wizard](#).

EFS Backup

Veeam Backup for AWS performs EFS backup in the following way:

1. Veeam Backup for AWS uses the [AWS Backup service](#) to create a cloud-native backup of the file system, and saves this backup to the specified backup vault in the same AWS Region in which the source file system resides.

The backup is assigned AWS tags upon creation. Keys and values of AWS tags contain encrypted metadata that helps Veeam Backup for AWS identify the related EFS file system backup.

2. If you configure the EFS backup policy to copy backup files to another AWS Region, Veeam Backup for AWS copies the created backup to the target AWS Region in the same AWS account.
3. If you enable EFS indexing in the backup policy settings, Veeam Backup for AWS performs the following operations:

- a. Deploys a worker instance in an AWS Region in which the processed file system resides in an AWS account to which the file system belongs – that is, the production AWS account.

By default, Veeam Backup for AWS selects the most appropriate network settings of AWS Regions in production accounts (for example, selects a VPC specified as a mount target for the processed file system). However, you can add specific worker configurations. For more information on worker instances, see [Managing Worker Configurations](#).

- b. Mounts the source file system on the worker instance.
- c. Reads data from the file system using the worker instance, creates a catalog of files and folders (index) of the system, transfers the index to a backup repository and stores it in the native Veeam format.
- d. The EFS index is associated with the cloud-native backup created at step 1 and the backup copy created at step 2. However, if the indexing session does not complete by the time a new backup session starts, a new indexing session is not launched and Veeam Backup for AWS associates the created EFS index with 2 cloud-native backups and backup copies created by 2 backup sessions.

NOTE

By default, Veeam Backup for AWS compresses data saved to backup repositories. To learn how to encrypt data stored in backup repositories, see [Data Encryption](#).

4. When the indexing session completes, removes the worker instance from Amazon EC2.

Backup Chain

During every backup session, Veeam Backup for AWS creates a cloud-native backup for each EFS file system added to the backup policy. To create the backup, Veeam Backup for AWS uses the [AWS Backup service](#).


A sequence of cloud-native backups created during a set of backup sessions makes up a backup chain. Veeam Backup for AWS creates the backup chain in the following way:

1. During the first backup session, Veeam Backup for AWS creates a backup that contains all EFS file system data and saves it in the selected backup vault of the AWS Region where the processed file system resides. This backup becomes a starting point in the backup chain.

The creation of the first backup may take significant time to complete since Veeam Backup for AWS copies the whole image of the EFS file system.

2. During subsequent backup sessions, Veeam Backup for AWS creates backups that contain only those data blocks (files and directories) that have changed since the previous backup session.

The creation of subsequent backups typically takes less time to complete, compared to the first backup in the chain. Note, however, that the completion time still depends on the amount of processed data.



Each EFS backup in the backup chain contains encrypted metadata. Metadata stores information about the protected file system, the backup policy that created the backup, and the date, time and applied retention settings. Veeam Backup for AWS uses metadata to identify outdated backups, to load the configuration of source file systems during recovery operations, and so on.

EFS backups act as independent restore points for backed-up file systems. If you remove any backup, it will not break the EFS backup chain – you will still be able to roll back file system data to any existing restore point. The period of time during which EFS backups are kept in the EFS backup chain is defined by retention policy settings. For more information, see [EFS Backup Retention](#).

NOTE

EFS backups created manually are not included into the EFS backup chain. Therefore, these backups are not removed automatically according to retention policy settings. For information on how to remove them, see [Removing EFS Backups Created Manually](#).

EFS Backup Copy Chain

If you enable backup copying for a backup policy, Veeam Backup for AWS will make a copy of the initially created EFS backup and save it to the target AWS Region specified in the backup policy settings. In the target AWS Region, backup copies created during a set of backup sessions make up a backup copy chain.

Veeam Backup for AWS creates and maintains an EFS backup copy chain in the same way as a regular EFS backup chain:

- The first created backup copy of the processed file system becomes a starting point in the backup copy chain.
- Backup copies created during subsequent backup sessions store only those data blocks that have changed since the previous backup session.

EFS Indexing Chain

If you enable EFS indexing for a backup policy, Veeam Backup for AWS during each indexing session creates and index of the processed file system and associates the index with one or multiple restore points as described in section [EFS Backup](#). In the target backup repository, EFS indexes created during a set of indexing sessions make up an indexing chain.

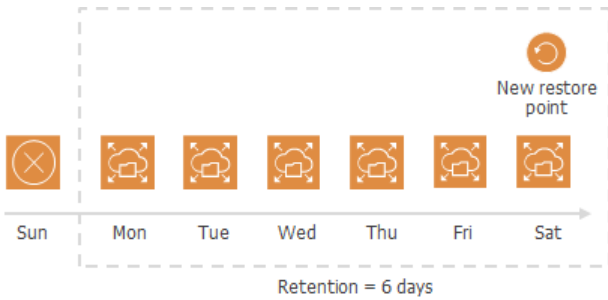
EFS indexes always contain full file catalogs of the processed file system. Therefore, if you delete any index from the backup repository, the index chain will not be corrupted but you may not be able to restore file and folders to a restore point associated with the deleted index using the file-level recovery browser. To learn how to perform file-level recovery, see [Performing File-Level Restore](#).

The period of time during which EFS indexes are kept in the indexing chain is defined by time stamps that were saved in the index metadata when creating the indexes. For more information, see [EFS Backup Retention](#).

EFS Backup Retention

For EFS file system backups, Veeam Backup for AWS retains restore points for the period of time specified in [backup scheduling settings](#).

During every successful backup session, Veeam Backup for AWS creates a restore point and saves the date, time and applied retention settings in the restore point metadata. If Veeam Backup for AWS detects that the period of time for which the restore point was stored exceeds the period specified in the retention settings, it automatically removes the restore point from the EFS backup chain. You can also remove unnecessary EFS backups manually as described in section [Removing EFS Backups](#).



NOTE

Veeam Backup for AWS does not apply retention policy to EFS backups created manually. To learn how to remove them, see [Removing EFS Backups Created Manually](#).

EFS Indexing Retention

When creating an index, Veeam Backup for AWS writes to the index metadata a time stamp when the index must be deleted. The time stamp is defined by the retention specified in the backup policy settings for the first restore point with which the index is associated. If you change retention settings for the backup policy, time stamps of earlier created indexes will not change. However, even if the index must be deleted according to the time stamp, Veeam Backup for AWS will not delete the index until all associated restore points are removed from the Veeam Backup for AWS configuration database.

EFS Restore

Veeam Backup for AWS offers the following restore options:

- File system restore – restores an entire Amazon EFS file system from an EFS backup or a backup copy. You can restore one or more Amazon EFS file systems at a time, to the original location or to a new location.
- File-level recovery – recovers individual files and folders stored in a file system from an EFS backup or backup copy. You can restore files and folders to the original file system or to another file system.

You can restore EFS file system data to the most recent state or to any available restore point.

IMPORTANT

You can restore an EFS file system only to the same AWS account to which the source file system belongs.

How File System Restore Works

To restore an EFS file system from a backup, Veeam Backup for AWS performs the following steps using native [AWS capabilities](#):

1. Creates a file system in the specified location.
2. Modifies the configuration setting values of the created EFS file system.
3. Restores backed-up files and folders to the restored file system.

To learn how to restore an entire Amazon EFS file system from an EFS backup or a backup copy, see [EFS Restore](#).

How EFS File-Level Recovery Works

To recover files and folders of a backed-up file system using specific file paths, Veeam Backup for AWS sends an API request to AWS to restore the specified files to the selected file system.

To recover files and folders of a backed-up file system using specific file paths, Veeam Backup for AWS performs the following steps:

1. On the backup appliance, restores the EFS index associated with the specified restore point.
2. Launches the file-level recovery browser.

The file-level recovery browser displays the file system tree of the backed-up EFS file system. In the browser, you select the necessary files and folders to restore.

3. Creates a new EFS directory `aws-backup-restore_<datetime>` in the root directory of the selected file system and restores the specified backed-up files and folders to the created directory.

To learn how to restore individual files and folders stored in a file system from an EFS backup or backup copy, see [EFS Restore](#).

Protecting FSx File Systems

With Veeam Backup for AWS, you can perform the following operations to protect FSx file systems:

- Create cloud-native backups of FSx file systems and store them in any backup vault in the source AWS Region.

An Amazon FSx file system backup captures the whole image of the FSx file system (including file system configuration, files, directories and so on) at a specific point of time. FSx backups are taken using native [AWS capabilities](#).

- Create backup copies of FSx file systems and store them in AWS Regions within the same AWS account.

By default, FSx backups are stored only in the AWS Region where the processed resources reside. For enhanced data safety, you can instruct Veeam Backup for AWS to create copies of these backups and store them in any other [default AWS Region](#) (that is, an AWS Region activated by default for your AWS account) within the same AWS account.

To protect FSx file systems, Veeam Backup for AWS runs backup policies. A backup policy is a collection of settings that define the way backup operations are performed: what data to back up, where to store backups, when to start the backup process, how to retain restore points, and so on.

Veeam Backup for AWS does not install agent software inside instances to back up FSx file systems – it uses native AWS capabilities instead. During every backup session, Veeam Backup for AWS creates a cloud-native backup for each file system added to a backup policy. The cloud-native backup is further used to create a backup copy in another AWS Region.

Supported FSx File System Properties

Veeam Backup for AWS allows you to back up and restore the following file system properties:

Property	Description	Ability to Change Property During Restore to New Location
File system name	Name of an FSx file system.	Yes
Deployment type	Deployment type of the FSx file system.	No
Storage type	Storage type of the FSx file system.	No
Storage capacity	Storage capacity of the FSx file system.	No
Throughput capacity [Applies only to FSx for Windows File Server and FSx for OpenZFS file systems]	Sustained speed at which file servers hosting the FSx file system can process data.	No

Property	Description	Ability to Change Property During Restore to New Location
Throughput per unit of storage [Applies only to FSx for Lustre file systems with the Persistent deployment type]	Read/write throughput for each 1 TiB of provisioned storage, in MB/s/TiB.	No
IOPS	Maximum number of input/output operations per second that the FSx file system can process.	No
Provisioned Metadata IOPS [Applies only to FSx for Lustre file systems]	Maximum rate of metadata operations supported by the FSx file system.	No
Volume properties [Applies only to FSx for OpenZFS file systems]	Properties of the FSx file system configured to manage its root volume storage.	No
Provisioned SSD IOPS [Applies only to FSx for Windows File Server file systems]	Additional IOPS provisioned to the FSx file system.	No
Data compression type [Applies only to FSx for Lustre file systems]	Defines whether the FSx file system data is compressed.	No
Windows authentication [Applies only to FSx for Windows File Server file systems]	Type of an Microsoft Active Directory to which the FSx file system is joined.	No
Active Directory domain name [Applies only to FSx for Windows File Server file systems]	Domain name of the Microsoft AD to which the FSx file system is joined.	Yes
DNS server IP addresses [Applies only to FSx for Windows File Server file systems]	IPv4 addresses of DNS servers configured for the domain.	Yes
Service account user name [Applies only to FSx for Windows File Server file systems]	Name of a service account that has access to the FSx file system.	Yes

Property	Description	Ability to Change Property During Restore to New Location
Organizational unit [Applies only to FSx for Windows File Server file systems]	Path name of an organizational unit in which the FSx file system is connected.	Yes
System administrators group [Applies only to FSx for Windows File Server file systems]	Name of an Active Directory group that has privileges to manage the FSx file system.	Yes
DNS Aliases [Applies to FSx for Windows File Server file systems]	Additional names that are used to access FSx file system data.	No
Root squash configuration [Applies only to FSx for Lustre file systems]	Additional layer of access control configured for the FSx file system.	No
Tags	File systems identifiers.	No
VPC	VPC to which the FSx file system is connected.	Yes
Subnet	Subnet in which the elastic network interface of the FSx file system resides.	Yes
Preferred and standby subnet [Applies only to FSx for Windows File Server and FSx for OpenZFS file systems file systems]	Subnets in which the network interfaces of the primary and standby file servers reside.	Yes
Security groups	Security groups associated with the FSx file system.	Yes
Route tables [Applies only to FSx for OpenZFS file systems file systems]	Route tables associated with the subnet of the VPC to which the FSx file system is connected.	Yes
Encryption key	AWS KMS key that is used to encrypt the FSx file system data.	Yes

Property	Description	Ability to Change Property During Restore to New Location
Backup and maintenance	Defines whether daily automatic backup is scheduled and whether a weekly maintenance window is configured for the FSx file system.	No

IMPORTANT

Veeam Backup for AWS does not support restore of the following items:

- [For FSx for OpenZFS file systems] Copy tags to snapshots from the root volume and copy tags to backups settings
- [For FSx for OpenZFS file systems with the Multi-AZ deployment type] Configured endpoint IP address ranges
- [For FSx for Windows File Server file systems] Copy tags to backups setting and file access auditing configurations
- [For FSx for Windows File Server file systems] Associated DNS aliases when restoring to a new location or to the original location but with different settings
- [For Amazon FSx for Lustre file systems] Copy tags to backups setting, CloudWatch logging configurations, and exceptions to root squash settings

Required Ports

If you plan to protect FSx file systems, make sure that security groups associated with these file systems allow access to the following ports:

File System Type	Protocol	Ports	Notes
Amazon FSx for Windows File Server	UDP	53, 88, 123, 389, 464	Requires inbound and outbound access.
	TCP	53, 88, 135, 389, 445, 464, 636, 3268, 3269, 5985, 9389, 49152-65535	
Amazon FSx for Lustre	TCP	988, 1018-1023	Requires inbound access. For more information on file system access control, see AWS Documentation .
Amazon FSx for OpenZFS	TCP, UDP	111, 2049, 20001-20003	Requires inbound access.

To learn how to authorize access to security groups, see [AWS Documentation](#).

How To Protect FSx File Systems

To create an FSx backup policy, perform the following steps:

1. [Check limitations and prerequisites.](#)
2. [Specify IAM roles to access AWS services and resources.](#)
3. [\[Optional\] Configure email notification settings for automated delivery of backup policy results and daily reports.](#)
4. [Complete the Add FSx Policy wizard.](#)

FSx Backup

Veeam Backup for AWS performs FSx backup in the following way:

1. Veeam Backup for AWS uses the [AWS Backup service](#) to create a cloud-native backup of the file system, and saves this backup to the specified backup vault in the same AWS Region in which the source file system resides.

The backup is assigned AWS tags upon creation. Keys and values of AWS tags contain encrypted metadata that helps Veeam Backup for AWS identify the related FSx file system backup.

2. If you configure the FSx backup policy to copy backup files to another AWS Region, Veeam Backup for AWS copies the created backup to the target AWS Region in the same AWS account.

Backup Chain

During every backup session, Veeam Backup for AWS creates a cloud-native backup for each FSx file system added to the backup policy. To create the backup, Veeam Backup for AWS uses the [AWS Backup service](#).


A sequence of cloud-native backups created during a set of backup sessions makes up a backup chain. Veeam Backup for AWS creates the backup chain in the following way:

1. During the first backup session, Veeam Backup for AWS creates a backup that contains all FSx file system data and saves it in the selected backup vault of the AWS Region where the processed file system resides. This backup becomes a starting point in the backup chain.

The creation of the first backup may take significant time to complete since Veeam Backup for AWS copies the whole image of the FSx file system.

2. During subsequent backup sessions, Veeam Backup for AWS creates backups that contain only those data blocks (files and directories) that have changed since the previous backup session.

The creation of subsequent backups typically takes less time to complete, compared to the first backup in the chain. Note, however, that the completion time still depends on the amount of processed data.



Each FSx backup in the backup chain contains encrypted metadata. Metadata stores information about the protected file system, the backup policy that created the backup, and the date, time and applied retention settings. Veeam Backup for AWS uses metadata to identify outdated backups, to load the configuration of source file systems during recovery operations, and so on.

FSx backups act as independent restore points for backed-up file systems. If you remove any backup, it will not break the FSx backup chain – you will still be able to roll back file system data to any existing restore point. The period of time during which FSx backups are kept in the FSx backup chain is defined by retention policy settings. For more information, see [FSx Backup Retention](#).

NOTE

FSx backups created manually are not included into the FSx backup chain. Therefore, these backups are not removed automatically according to retention policy settings. For information on how to remove them, see [Removing FSx Backups Created Manually](#).

FSx Backup Copy Chain

If you enable backup copying for a backup policy, Veeam Backup for AWS will make a copy of the initially created FSx backup and save it to the target AWS Region specified in the backup policy settings. In the target AWS Region, backup copies created during a set of backup sessions make up a backup copy chain.

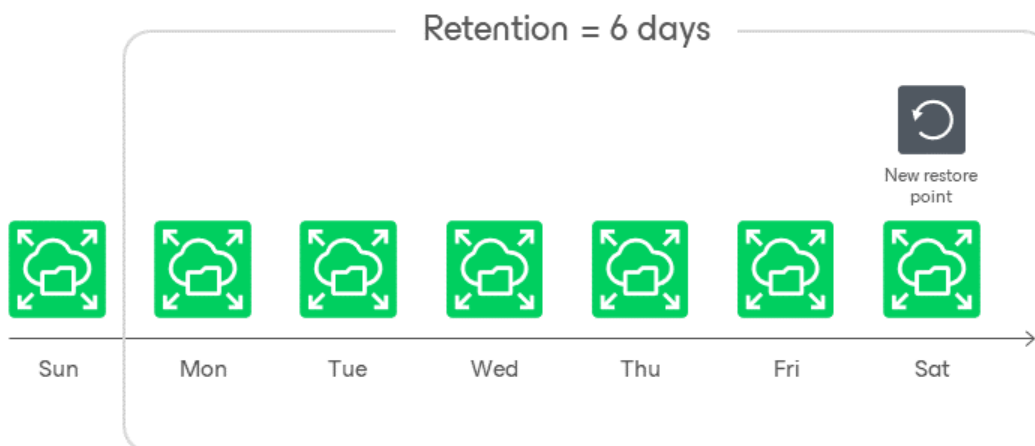
Veeam Backup for AWS creates and maintains a FSx backup copy chain in the same way as a regular FSx backup chain:

- The first created backup copy of the processed file system becomes a starting point in the backup copy chain.
- Backup copies created during subsequent backup sessions store only those data blocks that have changed since the previous backup session.

FSx Backup Retention

For FSx file system backups, Veeam Backup for AWS retains restore points for the period of time specified in [backup scheduling settings](#).

During every successful backup session, Veeam Backup for AWS creates a restore point and saves the date, time and applied retention settings in the restore point metadata. If Veeam Backup for AWS detects that the period of time for which the restore point was stored exceeds the period specified in the retention settings, it automatically removes the restore point from the FSx backup chain. You can also remove unnecessary FSx backups manually as described in section [Removing FSx Backups](#).



NOTE

Veeam Backup for AWS does not apply retention policy to FSx backups created manually. To learn how to remove them, see [Removing FSx Backups Created Manually](#).

FSx Restore

IMPORTANT

You can restore an FSx file system only to the same AWS account to which the source file system belongs.

To restore an FSx file system from a backup, Veeam Backup for AWS performs the following steps using native AWS capabilities:

1. Creates a file system in the specified location.
2. Modifies the configuration setting values of the created FSx file system.
3. Restores backed-up files and folders to the restored file system.

To learn how to restore an Amazon FSx file system from an FSx backup or a backup copy, see [FSx Restore](#).

Protecting VPC Configurations

To protect Amazon VPC configurations, Veeam Backup for AWS retrieves configuration data through API and saves this data to the configuration database. You can also instruct Veeam Backup for AWS to store copies of VPC configuration backups in a backup repository. For more information on how VPC configuration backup works, see [VPC Configuration Backup](#).

How To Protect VPC Configurations

To configure the VPC configuration backup policy settings, perform the following steps:

1. [Check limitations and prerequisites](#).
2. [Specify IAM role or add custom IAM roles to access AWS services and resources](#).
3. [Add backup repositories to save additional VPC configuration backup copies](#).
4. [\[Optional\] Configure global retention settings for obsolete session records](#).
5. [\[Optional\] Configure email notification settings for automated delivery of backup policy results and daily reports](#).
6. [Complete the VPC Configuration Backup wizard](#).

VPC Configuration Backup

Veeam Backup for AWS performs VPC configuration backup in the following way:

1. Sends API requests to AWS to retrieve the VPC configuration data, and saves this data in the configuration database.


To back up VPC configurations of AWS Regions added to backup policies, Veeam Backup for AWS uses permissions of IAM roles specified in the backup policy settings. The VPC configuration data is collected for the selected AWS Regions in the AWS accounts to which the specified IAM roles belong.

2. Veeam Backup for AWS creates a configuration record for each pair of the AWS account and an AWS Region whose VPC configuration data is being backed up. Every time the VPC Configuration Backup policy runs, Veeam Backup for AWS updates the record to create a new restore point for the VPC configurations. For more information, see [VPC Configuration Backup Chain](#).
3. If you [enable additional backup copy](#) for the VPC Configuration Backup policy, Veeam Backup for AWS launches the Veeam Data Mover service on the backup appliance to copy restore points to the target backup repository, creating an individual folder for each AWS account whose VPC configuration data is protected by the policy.

Backup Chain

During every backup session, Veeam Backup for AWS creates a restore point with backed-up VPC configuration data for each AWS Region protected by the VPC Configuration Backup policy. The restore point contains encrypted metadata that includes information on the date and time when the policy ran, AWS Regions whose VPC configuration settings were backed up by the policy, and AWS accounts whose IAM roles were used to collect VPC configuration settings for each AWS Region.

A sequence of restore points created during a set of backup sessions makes up a VPC configuration backup chain for each configuration record.



You cannot delete specific restore points created for a configuration record – these points are removed automatically according to the specified [retention policy settings](#). However, you can manually remove a configuration record with all restore points created for it, as described in section [Removing VPC Configuration Backups](#).

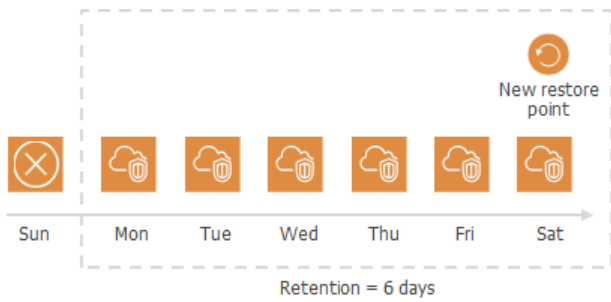
VPC Configuration Backup Retention

For VPC configuration backups, Veeam Backup for AWS retains restore points for the period of time specified in [backup retention settings](#).

During every successful backup session, Veeam Backup for AWS creates a restore point and saves the date, time and the applied retention settings in the restore point metadata. If Veeam Backup for AWS detects that the period of time for which the restore point was stored exceeds the period specified in the retention settings, it automatically removes the restore point from the VPC configuration backup chain. You can also remove unnecessary VPC configuration backups manually as described in section [Removing VPC Configuration Backups](#).

NOTE

Veeam Backup for AWS applies the retention settings configured for the [VPC Configuration Backup policy](#) both to VPC configuration backups stored in the Veeam Backup for AWS database and to VPC configuration backups stored in the backup repository selected for the policy. For VPC configuration backups stored in backup repositories that are not specified in the VPC Configuration Backup policy settings, Veeam Backup for AWS applies retention settings saved in the backup metadata.



Exporting VPC Configuration

You can export backed-up VPC configuration data to an AWS CloudFormation template in the JSON format using one of the following options:

- [Perform the entire VPC configuration export.](#)
- [Perform the selected VPC configuration items export.](#)

VPC Configuration Restore

Veeam Backup for AWS offers the following disaster recovery operations:

- [VPC configuration restore](#) – restores an entire VPC configuration from a VPC configuration backup. You can restore the VPC configuration to the original location or to a new location.
- [Selected items restore](#) – restores the selected VPC configuration items from a VPC configuration backup. You can restore specific VPC configuration items only to the original location.

You can restore the VPC configuration data to the most recent state or to any available restore point.

IMPORTANT

When restoring VPC route tables, consider that routes that had the `blackhole` state when a restore point was created will not be restored and a restore session will complete with warning. In this case, it is recommended to check the restored target route table configurations in the AWS Management Console to ensure that all traffic flows correctly. To learn how to configure routes in route tables, see [AWS Documentation](#).

Entire VPC Configuration Restore

To restore the entire VPC configuration from a backup, Veeam Backup for AWS performs the following steps:

1. Retrieves the backed-up VPC configuration from the Veeam Backup for AWS database.
2. Validates the restore operation: sends API requests to AWS to verify that AWS service quotas are not exceeded and there are no subnet CIDR block conflicts.
3. Retrieves information on existing items and their settings in the current Amazon VPC configuration.
4. Restores the backed-up VPC configuration:
 - a. Creates the missing VPC configuration items.
 - b. Modifies settings of the existing items that do not match the backed-up settings.

To learn how to restore an entire VPC configuration from a VPC configuration backup, see [Performing Entire Configuration Restore](#).

Selected Items Restore

To restore specific items of the VPC configuration from a backup, Veeam Backup for AWS performs the following steps:

1. Retrieves from the Veeam Backup for AWS database the backed-up VPC configuration data on items added to a [restore list](#).
2. Validates the restore operation: sends API request to AWS to verify that AWS service quotas are not exceeded and there are no subnet CIDR block conflicts.
3. Retrieves information on existing items and their settings in the current Amazon VPC configuration.

4. Validates the restore list: sends API requests to AWS to check whether any of the selected VPC configuration items depend on other items that are missing from the current VPC configuration.

In case any VPC configuration items on which the selected items depend are missing, Veeam Backup for AWS allows the user to add the missing items to the restore list.

5. Restores the selected items of the backed-up VPC configuration:
 - Creates the missing VPC configuration items.
 - Modifies settings of the existing items that do not match the backed-up settings.

IMPORTANT

Consider the following:

- VPC peering connections will have the *Pending Acceptance* status after restoring. To accept the restored VPC peering connections, use the AWS Management Console. For more information, see [AWS Documentation](#).
- If restore of any selected item fails, Veeam Backup for AWS will stop the restore operation and initiate a rollback. During the rollback, Veeam Backup for AWS will delete all newly created items, but will retain all changes made to the existing VPC configuration items.

To learn how to restore restores the selected VPC configuration items from a VPC configuration backup, see [Performing Selected Items Restore](#).

Retention Policies

Cloud-native snapshots, snapshot replicas and image-level backups are not kept forever. They are removed according to retention policy specified in the backup schedule settings while creating a backup policy.

Depending on the data protection scenario, retention policy can be specified:

- **In restore points** – for cloud-native snapshots and snapshot replicas.

The snapshot chain can contain only the allowed number of restore points. If the number of allowed restore points is exceeded, Veeam Backup for AWS removes the earliest restore point from the snapshot chain. For more information, see [EC2 Backup Retention](#) and [RDS Backup Retention](#).

- **In days/months/years** – for backups and archives.

Restore points in the backup chain (either standard or archive) can be stored for the allowed period of time. If a restore point is older than the specified limit, Veeam Backup for AWS removes it from the backup chain. For more information, see sections [EC2 Backup Retention](#), [RDS Backup Retention](#), [Redshift Backup Retention](#), [DynamoDB Backup Retention](#), [EFS Backup Retention](#), [FSx Backup Retention](#) and [VPC Configuration Backup Retention](#).

NOTE

When configuring policy scheduling, consider that Veeam Backup for AWS runs retention sessions at 4:00 AM by default, according to the time zone set on the backup appliance. If you schedule backup policies to execute at 4:00 AM, the backup policies and retention tasks will be queued.

You can also specify global retention settings for obsolete snapshots and replicas. For more information, see [Configuring Global Retention Settings](#).

Immutability

NOTE

This section applies only to EC2 and RDS image-level backups stored in backup repositories.

Veeam Backup for AWS allows you to protect data stored in backup repositories from deletion by making the data temporarily immutable. To do that, Veeam Backup for AWS uses [Amazon S3 Object Lock](#) – once imposed, S3 Object Lock prevents objects from being deleted or overwritten for a specific immutability period. The immutability period is set based on the retention policy configured in the backup policy settings.

Block Generation

If you choose a repository with immutability settings enabled as the target location for image-level backups, Veeam Backup for AWS creates an immutable backup chain in the repository instead of a regular backup chain. Immutable backup chains are built the same way as the chains of standard and archived EC2 and RDS backups, which means that each immutability chain is composed of a set of backups produced during a sequence of backup sessions, and that the same retention policies apply to these chains. The only difference is that objects in immutable backup chains can be neither removed nor modified until the immutability period is over. Therefore, every time Veeam Backup for AWS creates a new incremental backup containing modified data blocks, the retention period of the dependent unchanged data blocks (in the preceding incremental and full backups) is supposed to be extended. This can cause a substantial increase in I/O operations and associated costs incurred by Amazon S3.

To reduce the number of requests to the repository, thus to save traffic and to reduce transaction costs, Veeam Backup for AWS leverages the Block Generation mechanism. A generation is a period of up to 25 days that extends the retention period configured for backups composing the immutable backup chain. This means that the retention period is not explicitly extended for each dependent data block every time Veeam Backup for AWS creates a new incremental backup in the chain within one generation (during these 25 days).

NOTE

Veeam Backup for AWS initiates a dedicated generation for each type of the backup schedule configured in the [EC2 backup policy settings](#) or in the [RDS backup policy settings](#).

How Block Generation Works

Block Generation works in the following way:

1. During the first backup session, Veeam Backup for AWS creates a full backup in a backup repository and adds 25 days to its retention period. The full backup becomes a starting point in the first generation of the immutable backup chain.
2. During subsequent backup sessions, Veeam Backup for AWS copies only those data blocks that have changed since the previous backup session, and stores these data blocks to incremental backups in the backup repository. The content of each incremental backup depends on the content of the full backup and the preceding incremental backups in the immutable backup chain. Veeam Backup for AWS adds $<25 - N>$ days to the retention period of these backups, where N is the number of days since the first backup in the generation was created.

As a result, all backups within one generation will have the same retention date, and will not be removed by the retention policy before this date.

3. On the 26th day a new block generation period is initiated. Veeam Backup for AWS creates a new incremental backup and adds 25 days to its retention period. This backup becomes a starting point in the second generation of the immutable backup chain. The new generation is automatically applied to all dependent data blocks from the preceding backups.
4. Veeam Backup for AWS repeats step 2 for the second generation.
5. Veeam Backup for AWS continues keeping dependent data blocks immutable by applying new generations to these blocks, thus continuously extending their retention period.

IMPORTANT

As soon as a block generation is initiated, the immutability period of data blocks in the generation cannot be reduced. Even if you change the retention period configured for image-level backups in the backup policy settings, this will not affect the expiration date of the restore points that have been already created.

Block Generation Example

Consider the following example. You want a backup policy to create image-level backups of your critical workloads once a day starting from March 1, and to keep the backed-up data immutable for 5 days. In this case, you do the following:

1. In the policy target settings, you set the **Enable backups** toggle to *On*, and select a backup repository with immutability enabled as the target location for the created backups.
2. In the daily scheduling settings, you select an hour when backups will be created (for example, *7:00 AM*), and specify the number of days for which Veeam Backup for AWS will retain the created backups (*5 days*).

According to the specified scheduling settings, Veeam Backup for AWS will create image-level backups in the following way:

1. On March 1, a backup session will start at 7:00 AM to create the full backup in the immutable backup chain. Veeam Backup for AWS will add 25 days to the retention period specified in the backup policy settings. Thus, the retention period of the backup will be prolonged to 30 days, and the immutability expiration date will become March 31.
2. On March 2, Veeam Backup for AWS will create a new incremental backup at 7:00 AM and add 24 days to the retention period specified in the backup policy settings. Thus, the retention period of the incremental backup will be prolonged to 29 days, and the retention date will become March 31.
3. On March 3-25, Veeam Backup for AWS will continue creating incremental backups and extending their retention period so that the retention date will still remain March 31.
4. On March 26, Veeam Backup for AWS will create a new backup at 7:00 AM. During the backup session, Veeam Backup for AWS will initiate a new block generation period, and apply the new generation to the newly created backup and all dependent data blocks. The retention period of this backup will be prolonged to 30 days, and the immutability expiration date will become April 25.

Then, all data blocks of the preceding backups whose retention period has not been extended will be removed by a retention session due to the immutability period expiration.

How To Create Immutable Backups

To protect backups created with Veeam Backup for AWS from deletion by making them temporarily immutable, perform the following steps:

1. [Check limitations and prerequisites.](#)
2. [Add a backup repository with immutability enabled.](#)

3. Create an EC2 backup policy and specify the repository with immutability enabled as the target location for image-level backups.
4. Create a RDS backup policy and specify the repository with immutability enabled as the target location for image-level backups.

Private Network Deployment

The private deployment feature allows you to increase the security of your environment by retaining network traffic within a private network.

With Veeam Backup for AWS, you can deploy [backup appliances](#) and [worker instances](#) in a private environment. However, it is not necessary that all these components are connected to a private network – you can configure the backup infrastructure the way that suits your security concerns best.

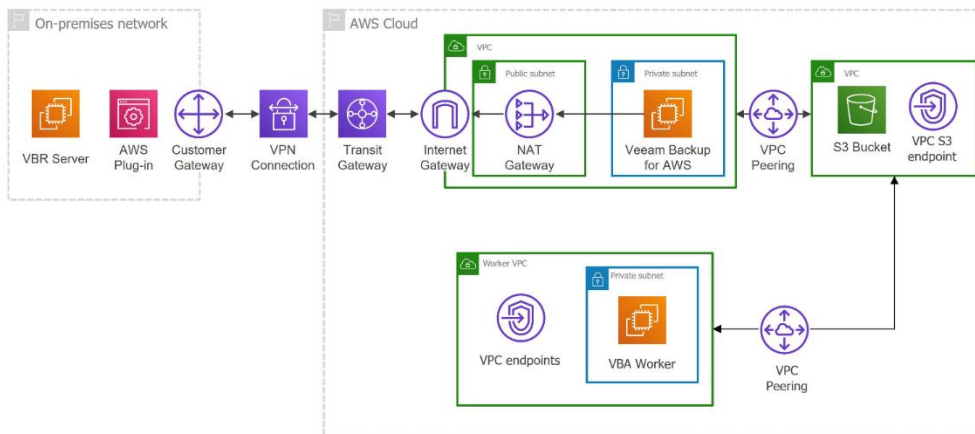
In This Section

- [Backup Appliances in Private Environment](#)
- [Worker Instances in Private Environment](#)

Backup Appliances in Private Environment

Starting from Veeam Backup for AWS version 7.0, you can deploy backup appliances in private networks to increase the security of your environment. When a backup appliance is deployed in a private environment, it is not assigned any public IPv4 address, and you will have to perform a number of additional configuration actions to allow private network access.

When deploying a backup appliance [from the Veeam Backup & Replication console](#), the only option is to connect it to an existing VPC. In this case, you must allow communication between the Veeam Backup & Replication server and the backup appliance. One possible solution is to establish an AWS Site-to-Site VPN (Site-to-Site VPN) connection between the VPC of the appliance and your on-premises network, as described in [Configuring Access to Backup Appliances in AWS](#).



Note that you must take into account the backup appliance requirements listed below.

Requirements for Backup Appliances

For a backup appliance to be able to operate in a private environment, the following requirements must be met:

- To download information on available product updates, the backup appliance requires the following outbound internet access:

From	To	Protocol	Port
Backup appliance	Veeam Update Repository (repository.veeam.com)	HTTPS	443
	Ubuntu Security Repository and OS Update repository (security.ubuntu.com , archive.ubuntu.com)	HTTP	80
	Microsoft Package Repository (packages.microsoft.com , dotnetcli.blob.core.windows.net)	HTTPS	443

From	To	Protocol	Port
	PostgreSQL Apt Repository (apt.postgresql.org)	HTTP	80
	PostgreSQL Website* (postgresql.org)	HTTPS	443

*Required to download the repository key <https://www.postgresql.org/media/keys/ACCC4CF8.asc>.

- To perform data protection and disaster recovery operations, the backup appliance must have outbound internet access to the [AWS services](#).
- If you want to receive daily reports and email notifications on backup policy results, outbound internet access must be allowed from the backup appliance to the email service through port **443** over the HTTPS protocol or through the SMTP port specified in the email server settings (port **25** by default).
- If you want to enable single sign-on (SSO) authentication to log in to different software systems with the same credentials using the identity provider service, outbound internet access must be allowed from the user workstation to the identity provider through port **443** over the HTTPS protocol.
- If you want to access the Web UI component from a user workstation, inbound internet access must be allowed from the user workstation to the appliance through port **443** over the HTTPS protocol.
- If the backup appliance is managed by a Veeam Backup & Replication server, inbound internet access must be allowed from the server to the appliance through port **443** over the HTTPS protocol.

Configuring Access to Backup Appliances in AWS

NOTE

This section provides instructions on steps performed in a third-party application. Keep in mind that the instructions may become outdated. For up-to-date instructions, see [AWS Documentation](#).

To allow a Veeam Backup & Replication server to communicate with a backup appliance operating in a [private environment](#), you can establish an AWS Site-to-Site VPN (Site-to-Site VPN) connection between the VPC of the appliance and your on-premises network:

1. [Create a customer gateway](#).
2. [Create a virtual private target gateway and attach the gateway to the VPC](#).
3. [Enable route propagation](#).
4. [Allow inbound traffic to the backup appliance](#).
5. [Create a VPN connection](#).

Step 1. Create Customer Gateway

A customer gateway device is a physical device or software application in your on-premises network. A customer gateway is a resource in AWS representing the customer gateway device in the on-premises network. For more information, see [AWS Documentation](#).

To provide information on a customer gateway device to AWS, create a customer gateway:

1. Log in to the **AWS Management Console** using credentials of an AWS account in which you want to create the Site-to-Site VPN connection.
2. Navigate to **All Services > Networking & Content Delivery** and click **VPC**.
3. In the **VPC** console, navigate to **Virtual Private Network > Customer Gateways** and click **Create Customer Gateway**.
4. Complete the **Create customer gateway** wizard:
 - a. At the **Details** step of the wizard, do the following:
 - i. [Optional] In the **Name tag** field, specify a name for the gateway.
 - ii. In the **BGP ASN** field, specify a Border Gateway Protocol (BGP) Autonomous System Number (ASN) for the gateway.
 - iii. In the **IP address** field, specify a static, internet-routable IP address for the gateway.
 - iv. From the **Certificate ARN** drop-down list, specify the Amazon Resource Name of a private certificate that will be used to connect to the gateway.
 - v. [Optional] In the **Device** field, specify a name for the customer gateway device.
 - b. Click **Create customer gateway**.

Step 2. Create Virtual Private Target Gateway

To establish a VPN connection between the VPC of the backup appliance and your on-premises network, create a virtual private target gateway on the AWS side and attach the gateway to the VPC:

1. In the **VPC** console, navigate to **Virtual Private Network > Virtual Private Gateways** and click **Create Virtual Private Gateway**.
2. Complete the **Create virtual private gateway** wizard:
 - a. At the **Details** step, do the following:
 - i. [Optional] In the **Name tag** field, specify a name for the virtual private target gateway.
 - ii. In the **Autonomous System Number (ASN)** section, choose whether you want to keep the default ASN or specify a custom one. This ASN must not match the BGP ASN that you have specified for the customer gateway at [step 1](#).

For custom ASNs, the following limitations apply. For a 16-bit ASN, its value must be between 64512 and 65534; for a 32-bit ASN, its value must be between 4200000000 and 4294967294.

Note that after you create the VPN connection, you will not be able to change the ASN for it.
 - b. Click **Create virtual private gateway**.
3. To attach the created virtual private gateway to the VPC, select the gateway in the **Virtual private gateways** list and click **Actions > Attach to VPC**.
4. Complete the **Attach to VPC** wizard:
 - a. At the **Details** step, select the VPC from the list of available VPCs.
 - b. Click **Attach to VPC**.

Step 3. Configure Routing

To allow the backup appliance to access the customer gateway and to automatically propagate Site-to-Site VPN routes, enable route propagation in the route table associated with a subnet of the appliance VPC:

1. In the **VPC** console, navigate to **Virtual Private Cloud > Route Tables**.
2. In the **Route tables** list, choose the necessary route table and click **Actions > Edit Route Propagation**.
3. In the **Edit route propagation** wizard, select the **Enable** check box and click **Save**.

Step 4. Update Security Group

To allow inbound traffic to the backup appliance from the on-premises network, update the security group for the appliance VPC:

1. In the **VPC** console, navigate to **Security > Security Groups**.
2. In the **Security Group** list, choose the default security group and click **Actions > Edit Inbound Rules**.
3. In the **Edit inbound rules** wizard, click **Add rule**, add a new inbound rule for the SSH, RDP and ICMP protocols, and click **Save rules**.

To learn how to add security group rules, see [AWS Documentation](#).

Step 5. Create VPN Connection

To enable access to your on-premises network, create a VPN connection between the created virtual private gateway and the customer gateway:

1. In the **VPC** console, navigate to **Virtual private network > Site-to-Site VPN Connections** and click **Create VPN Connection**.
2. Complete the **Create VPN connection** wizard:
 - a. At the **Details** step of the wizard, do the following:
 - i. [Optional] In the **Name tag** field, specify a name for the VPN connection.
 - ii. In the **Autonomous System Number (ASN)** section, select the **Virtual private gateway** option and specify the ID of the virtual private gateway that you have created at [step 2](#).
 - iii. In the **Customer gateway** section, select the **Existing** option and specify the ID of the customer gateway.
 - iv. [Applies only if the customer device does not support Border Gateway Protocol] In the **Routing options** section, select the **Static** option and specify the IP prefixes of the appliance VPC.
 - b. Click **Create VPN connection**.

TIP

When you create a VPN connection, AWS generates a sample configuration file that can be further used to configure a customer gateway device. To download the file, do the following:

1. In the **VPC** console, navigate to **Virtual Private Network > Site-to-Site VPN Connections**.
2. From the **VPN connections** drop-down list, select the created connection and click **Download configuration**.
3. In the **Download configuration** window, select the vendor, class and operating system of the customer gateway device, and the IKE version that is used for the VPN connection. Then, click **Download**.

To learn how to configure a customer gateway device, see [AWS Documentation](#).

Worker Instances in Private Environment

Veeam Backup for AWS automatically deploys worker instances in Amazon EC2 for the duration of backup, restore and retention processes and removes it immediately after the processes complete. Veeam Backup for AWS deploys one worker instance per each processed AWS resource. To minimize cross-region traffic charges, depending on the data protection or disaster recovery operation, Veeam Backup for AWS deploys the worker instance in [specific locations](#).

IMPORTANT

If you want worker instances to operate in a private network, consider that worker instances must have outbound access to specific [AWS services](#).

By default, Veeam Backup for AWS uses public access to communicate with worker instances. However, you can instruct Veeam Backup for AWS to deploy worker instances without public IPv4 addresses, and then configure worker settings to allow private network access. One possible solution is to enable the private network deployment functionality in the Veeam Backup for AWS Web UI, create interface endpoints and ensure connectivity between your resources:

1. Set the **Private network deployment** toggle to *On* as described in section [Configuring Private Network Deployment](#).

NOTE

If you enable the private network deployment functionality, worker instances will communicate with the Amazon S3 service through a private S3 endpoint specified in [repository settings](#) – but only to perform data protection and recovery tasks, as well as retention tasks. To access the service while restoring the backup appliance configuration, exporting VPC configuration, creating and editing backup repositories, Veeam Backup for AWS will still use the public `s3.<region>.amazonaws.com` endpoint.

2. To allow worker instances to access AWS services, create specific VPC interface endpoints for all subnets to which the worker instances will be connected.

For the list of VPC interface endpoints required for backup and restore operations, see [Configuring Private Networks](#).

3. To enable route traffic between different VPCs, create peering connections between those VPCs.
4. To enable private traffic between different VPCs, add routes to the route tables associated with the subnets of those VPCs.

The actions you perform depend on specific use cases. For more information, see [Example 1. Creating EC2 Backups](#) and [Example 2. Archiving EC2 Backups](#).

Requirements for Private Network Deployment

If you enable the private network deployment functionality, consider the following:

- The backup appliance and worker instances must be able to communicate with the Amazon S3 service through an S3 interface endpoint. That is why security groups associated with the endpoint network interface must allow inbound HTTPS traffic from both the backup appliance and the worker instances through port **443**.

- Specific VPC interface endpoints must be created for subnets to which worker instances will be connected to access [AWS services](#), and the security group associated with the endpoint network interfaces must allow local inbound traffic through port **443**.

For the list of VPC interface endpoints required for specific backup and restore operations, see [Configuring Private Networks](#).

- Security groups associated with the worker instances must allow outbound HTTPS traffic to all endpoints through port **443**.

IMPORTANT

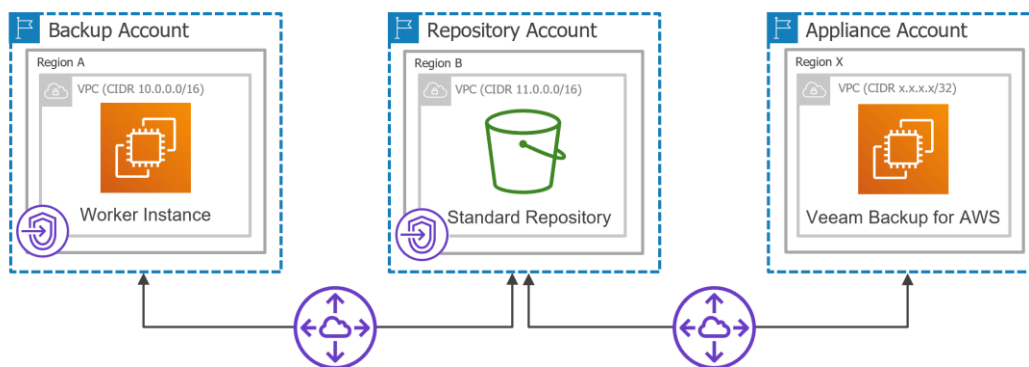
S3 gateway endpoints are not supported when using the private network deployment functionality.

Example 1. Creating EC2 Backups

Consider the following example. You need to backup an EC2 instance that belongs to a production account located in region A by deploying a worker instance in a backup account located in the same region and store its image-level backups in a backup repository that belongs to a repository account located in region B. The backup appliance belongs to an appliance account located in region X.

NOTE

To perform EC2 backup, Veeam Backup for AWS by default deploys worker instances in the backup account in the same AWS Region where source EC2 instances reside. However, you can instruct Veeam Backup for AWS to deploy worker instances in a production account. For more information, see [Worker Deployment Options](#).



In this case, you can perform the following steps in the AWS Management Console.

1. Establish a connection between the backup account and the repository account. To do that:
 - a. Create interface VPC endpoints required for worker instances to access AWS services.
 - i. Create a VPC to which the worker instances will be connected (for example, *10.0.0.0/16*) and a private subnet in the backup account. Note that the **Enable DNS name** check box must be enabled in the VPC settings.
 - ii. Create interface VPC endpoints for the private subnet to which the worker instances will be connected. These endpoints will be used to access the *ssm*, *sqs*, *ebs*, *ec2messages* services.

- iii. Create security groups associated with the endpoint network interfaces to allow local inbound HTTPS traffic (port **443**).

It is recommended to specify the full IPv4 address range of the VPC in the security group settings to make the created interface endpoints available for all resources in the VPC. If a security group restricts inbound HTTPS traffic from the resources, you will not be able to send traffic through the endpoint network interfaces.

- b. Configure an S3 interface endpoint required for the worker instances to access the Amazon S3 service.
 - i. Create a VPC (for example, *11.0.0.0/16*) and a private subnet in the repository account. Make sure that the CIDR block of the repository VPC differs from the CIDR block of the worker instance VPC to avoid subnet CIDR block conflicts.
 - ii. Create an S3 interface VPC endpoint for the private subnet to which the worker instances will be connected. The endpoint will be used to access the Amazon S3 service.
 - iii. Create a security group associated with the endpoint network interface to allow inbound HTTPS traffic (port **443**) from both the backup appliance and the worker instances.

Security Group	From	Protocol	Port	Notes
Group associated with the VPC interface endpoints	Worker instance VPC (10.0.0.0/16)	TCP	443	Allows local inbound HTTPS traffic
Group associated with the S3 interface endpoint	Appliance VPC (x.x.x.x/32)	TCP	443	Allows inbound HTTPS traffic from the backup appliance
	Worker instance VPC (10.0.0.0/16)	TCP	443	Allows inbound HTTPS traffic from the worker instances

- c. Configure the following peering connection settings.
 - i. Create a VPC peering connection between the worker instance VPC and repository VPC, and accept the peering request to enable route traffic between those VPCs.
 - ii. Add routes to the route tables associated with the subnets of the worker instance VPC and repository VPC to enable private traffic between those VPCs.

Destination	Target
Worker Instance VPC	
Worker instance VPC (10.0.0.0/16)	<i>Local</i>
Repository VPC (11.0.0.0/16)	<i>pcx-xxxx</i>

Destination	Target
Repository VPC	
Repository VPC (11.0.0.0/16)	<i>Local</i>
Worker instance VPC (10.0.0.0/16)	<i>pcx-xxxx</i>

2. Establish a connection between the repository account and the appliance account. To do that:
 - a. Create a VPC peering connection between the appliance VPC and repository VPC, and accept the peering request to enable route traffic between those VPCs.
 - b. Add routes to the route tables of the repository VPC and appliance VPC to enable private traffic between those VPCs.

Destination	Target
Repository VPC	
Repository VPC (11.0.0.0/16)	<i>Local</i>
Worker instance VPC (10.0.0.0/16)	<i>pcx-xxxx</i>
Appliance VPC (x.x.x.x/32)	<i>pcx-yyyy</i>
Appliance VPC	
Appliance VPC (x.x.x.x/32)	<i>Local</i>
Repository VPC (11.0.0.0/16)	<i>pcx-yyyy</i>

For detailed instructions on how to create interface endpoints, set up VPC peering connections and add routing, see [Configuring Private Networks](#).

TIP


If you have multiple AWS accounts and want to deploy worker instances in [production accounts](#), you can create a single resource share in one AWS account for all subnets to which the worker instances will be connected. The resource share can be further used to share these subnets with other AWS accounts in any organization. For more information, see [Configuring Private Networks for Production Accounts](#).

Example 2. Archiving EC2 Backups

Consider the following example. You need to copy backed-up data stored in a standard backup repository that belongs to a repository account located in region B to an archive backup repository that belongs to a repository account located in region C.

NOTE

To archive EC2 backups, Veeam Backup for AWS deploys worker instances in the [backup account](#) in the same AWS Region in which the standard backup repository with backed-up data resides.



In this case, you can perform the following steps in AWS Management Console.

1. Establish a connection between the backup account and the standard repository account. To do that:
 - a. Create interface VPC endpoints required for worker instances to access AWS services:
 - i. Create a VPC to which the worker instances will be connected (for example, *12.0.0.0/16*) and a private subnet in the backup account.
 - ii. Create interface VPC endpoints for the private subnet to which the worker instances will be connected. These endpoints will be used to access the *ssm*, *sqs* and *ec2messages* services.
 - iii. Create security groups associated with the endpoint network interfaces to allow local inbound HTTPS traffic (port **443**).

It is recommended to specify the full IPv4 address range of the VPC in the security group settings to make the created interface endpoints available for all resources in the VPC. If a security group restricts inbound HTTPS traffic from the resources, you will not be able to send traffic through the endpoint network interfaces.
 - b. Make sure that you have already configured the S3 interface endpoint in the standard repository account required for the worker instances to access the Amazon S3 service, as described in [Example 1](#).

IMPORTANT

The backup appliance and worker instances must be able to communicate with the Amazon S3 service through the created S3 interface endpoint. That is why security groups associated with the endpoint network interface must allow inbound HTTPS traffic from both the backup appliance and the worker instances through port **443**.

- c. Configure the following peering connection settings.

- i. Create a VPC peering connection between the worker instance VPC and standard repository VPC, and accept the peering request to enable route traffic between those VPCs.
- ii. Add routes to the route tables associated with the subnets of the worker instance VPC and standard repository VPC to enable private traffic between those VPCs.

Destination	Target
Worker Instance VPC	
Worker instance VPC (12.0.0.0/16)	<i>Local</i>
Standard repository VPC (11.0.0.0/16)	<i>pcx-zzzz</i>
Standard Repository VPC	
Standard repository VPC (11.0.0.0/16)	<i>Local</i>
Worker instance VPC (12.0.0.0/16)	<i>pcx-zzzz</i>

2. Establish a connection between the backup account and archive repository account. To do that:
 - a. Create interface VPC endpoints required for the worker instances to access AWS services to archive backups:
 - i. Create a VPC (for example, *13.0.0.0/16*) and a private subnet in the archive repository account.
 - ii. Create an S3 interface VPC endpoint for the private subnet to which the worker instances will be connected. The endpoint will be used to access the Amazon S3 service.
 - iii. Create a security group associated with the endpoint network interface to allow inbound HTTPS traffic (port **443**) from the worker instances, standard backup repository and backup appliance.
 - b. Configure the following peering connection settings.
 - i. Create a VPC peering connection between the worker instance VPC and archive repository VPC, and accept the peering request to enable route traffic between those VPCs.
 - ii. Add routes to the route tables associated with the worker instance VPC and archive repository VPC to enable private traffic between those VPCs.

Destination	Target
Worker Instance VPC	
Worker instance VPC (12.0.0.0/16)	<i>Local</i>
Standard repository VPC (11.0.0.0/16)	<i>pcx-zzzz</i>

Destination	Target
Archive repository VPC (13.0.0.0/16)	<i>pcx-xxxx</i>
Archive Repository VPC	
Archive repository VPC (13.0.0.0/16)	<i>Local</i>
Worker instance VPC (12.0.0.0/16)	<i>pcx-xxxx</i>

3. Establish a connection between the standard repository account and archive repository account. To do that:
 - a. Create a VPC peering connection between the standard repository VPC and archive repository VPC, and accept the peering request to enable route traffic between those VPCs.
 - b. Add routes to the route tables associated with the standard repository VPC and archive repository VPC to enable private traffic between those VPCs.

Destination	Target
Standard Repository VPC	
Standard repository VPC (11.0.0.0/16)	<i>Local</i>
Worker instance VPC (12.0.0.0/16)	<i>pcx-zzzz</i>
Archive repository VPC (13.0.0.0/16)	<i>pcx-xxxx</i>
Archive Repository VPC	
Archive repository VPC (13.0.0.0/16)	<i>Local</i>
Worker instance VPC (12.0.0.0/16)	<i>pcx-xxxx</i>
Standard repository VPC (11.0.0.0/16)	<i>pcx-xxxx</i>

- Update the security groups associated with the endpoint network interfaces to allow inbound HTTPS traffic (port **443**) from the backup appliance, the worker instances, the standard backup repository and the archive backup repository.

Security Groups Associated with S3 Interface Endpoint	From	Protocol	Port	Notes
Standard repository VPC	Backup appliance VPC (x.x.x.x/32)	TCP	443	Allows inbound HTTPS traffic from the backup appliance
	Worker instance VPC (12.0.0.0/16)	TCP	443	Allows inbound HTTPS traffic from the worker instances
	Archive repository VPC (13.0.0.0/16)	TCP	443	Allows inbound HTTPS traffic from the archive backup repository
Archive repository VPC	Backup appliance VPC (x.x.x.x/32)	TCP	443	Allows inbound HTTPS traffic from the backup appliance
	Worker instance VPC (12.0.0.0/16)	TCP	443	Allows inbound HTTPS traffic from the worker instances
	Standard repository VPC (11.0.0.0/16)	TCP	443	Allows inbound HTTPS traffic from the standard backup repository

For detailed instructions on how to create interface endpoints, set up VPC peering connections and add routing, see [Configuring Private Networks](#).

Configuring Private Networks

If you want worker instances to operate in a private environment – that is, to allow Veeam Backup for AWS to deploy worker instances with disabled auto-assignment of Public IPv4 addresses – you must configure specific endpoints for services used by the backup appliance to perform backup and restore operations:

Operation	Worker Instance Location	Possibility to Deploy Worker Instances in Production Accounts	Interface Endpoints	S3 Interface Endpoints
Creating EC2 image-level backups	AWS Region in which a processed EC2 instance resides	Yes	<ul style="list-style-type: none"> com.amazonaws.<region>.ec2 messages com.amazonaws.<region>.ssm com.amazonaws.<region>.sqs com.amazonaws.<region>.ebs 	<ul style="list-style-type: none"> com.amazonaws.<region>.s3
Restoring EC2 instances from image-level backups	AWS Region to which an EC2 instance is restored	Yes	<ul style="list-style-type: none"> com.amazonaws.<region>.ec2 messages com.amazonaws.<region>.ssm com.amazonaws.<region>.sqs 	<ul style="list-style-type: none"> com.amazonaws.<region>.s3
Restoring EC2 volumes from image-level backups	AWS Region to which the volumes of a processed EC2 instance are restored	Yes	<ul style="list-style-type: none"> com.amazonaws.<region>.ec2 messages com.amazonaws.<region>.ssm com.amazonaws.<region>.sqs 	<ul style="list-style-type: none"> com.amazonaws.<region>.s3
Performing health check for EC2 backups	AWS Region in which a backup repository with backed-up data resides	No	<ul style="list-style-type: none"> com.amazonaws.<region>.ec2 messages com.amazonaws.<region>.ssm com.amazonaws.<region>.sqs 	<ul style="list-style-type: none"> com.amazonaws.<region>.s3

Operation	Worker Instance Location	Possibility to Deploy Worker Instances in Production Accounts	Interface Endpoints	S3 Interface Endpoints
Creating EC2 archived backups	AWS Region in which a standard backup repository with backed-up data resides	No	<ul style="list-style-type: none"> com.amazonaws.<region>.ec2messages com.amazonaws.<region>.ssm com.amazonaws.<region>.sqs 	<ul style="list-style-type: none"> com.amazonaws.<region>.s3
Creating RDS image-level backups	AWS Region in which a processed DB instance resides	Yes	<ul style="list-style-type: none"> com.amazonaws.<region>.ssmmessages com.amazonaws.<region>.ssm com.amazonaws.<region>.sqs 	<ul style="list-style-type: none"> com.amazonaws.<region>.s3
Restoring PostgreSQL DB instances from image-level backups	AWS Region to which a PostgreSQL DB instance is restored	Yes	<ul style="list-style-type: none"> com.amazonaws.<region>.ssmmessages com.amazonaws.<region>.ssm com.amazonaws.<region>.sqs 	<ul style="list-style-type: none"> com.amazonaws.<region>.s3
Performing health check for RDS backups	AWS Region in which a backup repository with backed-up data resides	No	<ul style="list-style-type: none"> com.amazonaws.<region>.ssmmessages com.amazonaws.<region>.ssm com.amazonaws.<region>.sqs 	<ul style="list-style-type: none"> com.amazonaws.<region>.s3

Operation	Worker Instance Location	Possibility to Deploy Worker Instances in Production Accounts	Interface Endpoints	S3 Interface Endpoints
Creating RDS archived backups	AWS Region in which a standard backup repository with backed-up data resides	No	<ul style="list-style-type: none"> com.amazonaws.<region>.ssmmessages com.amazonaws.<region>.ssm com.amazonaws.<region>.sqs 	<ul style="list-style-type: none"> com.amazonaws.<region>.s3
Applying retention policy settings to created restore points	AWS Region in which a backup repository with backed-up data resides	No	<ul style="list-style-type: none"> com.amazonaws.<region>.ec2messages com.amazonaws.<region>.ssm com.amazonaws.<region>.sqs 	<ul style="list-style-type: none"> com.amazonaws.<region>.s3
Performing file-level recovery from image-level backups	AWS Region in which a backup repository with backed-up data resides	No	<ul style="list-style-type: none"> com.amazonaws.<region>.ec2messages com.amazonaws.<region>.ssm com.amazonaws.<region>.sqs [Applies only if you restore to the original location] com.amazonaws.<region>.kinesis-streams 	<ul style="list-style-type: none"> com.amazonaws.<region>.s3

Operation	Worker Instance Location	Possibility to Deploy Worker Instances in Production Accounts	Interface Endpoints	S3 Interface Endpoints
Performing file-level recovery from cloud-native snapshots and replicated snapshots	AWS Region in which a snapshot is located	<ul style="list-style-type: none"> No (if restoring to the original location) Yes (if restoring to a local machine) 	<ul style="list-style-type: none"> com.amazonaws.<region>.ec2messages com.amazonaws.<region>.ssm com.amazonaws.<region>.sqs [Applies only if you restore to the original location] com.amazonaws.<region>.kinesis-streams 	<ul style="list-style-type: none"> com.amazonaws.<region>.s3
Performing EFS indexing	Availability Zone in which a file system has a mount target created	Yes	<ul style="list-style-type: none"> com.amazonaws.<region>.ssmmessages com.amazonaws.<region>.ssm com.amazonaws.<region>.sqs 	<ul style="list-style-type: none"> com.amazonaws.<region>.s3

To create these endpoints, use the specified endpoint names, where `<region>` is the name of an AWS Region in which worker instances will be deployed.

How to Configure Private Networks

NOTE

This section provides instructions on steps performed in a third-party application. Keep in mind that the instructions may become outdated. For up-to-date instructions, see [AWS Documentation](#).

To configure private networks, use either of the following options:

- [Configuring private networks to deploy worker instances in the backup account.](#)
- [Configuring private networks to deploy worker instances in production accounts.](#)


NOTE

Following the provided instructions is not the only way to configure connectivity between your VPCs. Keep in mind that there exists a number of other possible workarounds.

Configuring Private Networks for Backup Account

For Veeam Backup for AWS to be able to deploy worker instances in a private environment in the [backup account](#), perform the following steps:

1. [Create VPC interface and S3 interface endpoints for subnets to which worker instances will be connected.](#)
2. [Create a peering connection between VPCs.](#)
3. [Add routes to the route tables associated with the subnets of the VPCs.](#)



Step 1. Create Interface Endpoints

To allow Veeam Backup for AWS to create EC2 and RDS image-level backups, to perform restore operations and to save EFS indexes to backup repositories, you must configure specific VPC interface endpoints for all subnets to which worker instances deployed for these operations will be connected. For the list of VPC interface endpoints required for backup and restore operations, see [Configuring Private Networks](#).

To deploy worker instances, Veeam Backup for AWS uses either the default or the [most appropriate network settings](#) of AWS Regions where the processed resources reside. However, you can add specific worker configurations as described in section [Managing Worker Configurations](#).

Creating Interface Endpoints

To create an interface VPC endpoint, do the following:

1. Log in to the **AWS Management Console** using credentials of an AWS account in which you want to create the endpoint.
2. Navigate to **Services > Networking & Content Delivery** and click **VPC**.
3. In the **VPC** console, navigate to **Virtual Private Cloud > Endpoints** and click **Create Endpoint**.
4. Complete the **Create endpoint** wizard:
 - a. At the **Endpoint settings** step, do the following:
 - i. [Optional] In the **Name tag** field, specify a name for the endpoint.
 - ii. In the **Service category** section, select the **AWS services** option.
 - b. At the **Services** step, enter *Interface* in the search field, and choose a service for which you want to create the endpoint.
 - c. At the **VPC** step, do the following:
 - i. From the **VPC** drop-down list, choose a VPC to which the deployed worker instances will be connected. Make sure that the **Enable DNS hostnames** check box is selected for the VPC.
 - ii. In the **Additional settings** section, select the **Enable DNS name** check box.
 - d. At the **Subnets** step, choose a subnet for each Availability Zone where the worker instances will be deployed, and specify the IP address type. Make sure that the **Auto-assign public IPv4 address** check box is not selected for the subnet.
 - e. At the **Security groups** step, choose security groups that will be associated with the endpoint network interface.

Ensure that each security group allows communication between the associated endpoint network interface and the resources in your VPC communicating with the selected service. If a security group restricts inbound HTTPS traffic (port **443**) from the resources in the VPC, you will not be able to send traffic through the endpoint network interface.
 - f. At the **Policy** step, select the **Full access** option to allow full access to the service. Alternatively, select the **Custom** option, and attach a VPC endpoint policy that will control permissions required to access available resources over the VPC endpoint.
 - g. Click **Create Endpoint**.

For more information on interface VPC endpoints, see [AWS Documentation](#).

Creating S3 Interface Endpoints

To create an S3 interface VPC endpoint, do the following:

1. In the **VPC** console, navigate to **Virtual Private Cloud > Endpoints** and click **Create Endpoint**.
2. Complete the **Create endpoint** wizard:
 - a. At the **Endpoint settings** step, do the following:
 - i. [Optional] In the **Name tag** field, specify a name for the endpoint.
 - ii. In the **Service category** section, select the **AWS services** option.
 - b. At the **Services** step, enter *S3* in the search field and choose the `com.amazonaws.<region>.s3` service with the *Interface* type, where `<region>` is the name of an AWS Region in which a backup repository is located.
 - c. At the **VPC** step, choose a VPC to which the deployed worker instances will be connected.
 - d. At the **Subnets** step, choose a subnet for each Availability Zone where the worker instances will be deployed, and specify the IP address type.
 - e. At the **Security groups** step, choose security groups that will be associated with the endpoint network interface.
 - h. At the **Policy** step, select the **Full access** option to allow full access to the service. Alternatively, select the **Custom** option, and attach a VPC endpoint policy that will control permissions required to access available resources over the VPC endpoint.
 - f. Click **Create Endpoint**.

IMPORTANT

The backup appliance and worker instances must be able to communicate with the Amazon S3 service through the created S3 interface endpoint. That is why security groups associated with the endpoint network interface must allow inbound HTTPS traffic from both the backup appliance and the worker instances through port **443**.

For more information on interface endpoints for Amazon S3, see [AWS Documentation](#).

Step 2. Create VPC Peering Connection

If you have created interface endpoints and S3 interface endpoints in subnets of two different VPCs, you must create a peering connection between the acceptor and requester VPC to enable route traffic between those VPCs using private IP addresses.

To create a VPC peering connection, do the following:

1. In the **VPC** console, navigate to **Virtual Private Cloud > Peering connections** and click **Create peering connection**.
2. Complete the **Create peering connection** wizard:
 - a. At the **Peering connection settings** step, do the following:
 - i. [Optional] In the **Name** field, specify a name for the connection.
 - ii. In the **Select a local VPC to peer with** section, choose the requester VPC.
 - iii. In the **Select another VPC to peer with** section, choose an AWS account and AWS Region in which you want to create the connection, and specify the ID of the acceptor VPC.
 - iv. In the **Tags** section, specify AWS tags that will be assigned to the connection.
 - b. Click **Create Peering Connection**.
3. To enable route traffic between the requester and acceptor VPC, select the created peering connection in the **Peering connections** list and click **Actions > Accept request**.

Step 3. Configure Routing

If you have created a peering connection between two different VPCs, you must add routes to the route tables associated with the subnets of the acceptor and requester VPC to enable private traffic between those VPCs.

To add a route to a route table, do the following:

1. In the **VPC** console, navigate to **Virtual Private Cloud > Route tables**, choose the route table and click **Actions > Edit routes**.
2. Complete the **Edit routes** wizard:
 - a. Click **Add routes**.
 - b. In the **Destination** field, specify the range of IPv4 addresses to which the network traffic in the peering connection must be directed.
The IPv4 address range must be specified in the CIDR notation (for example, `12.23.34.0/24`).
 - c. In the **Target** field, select the **Peering Connection** option and specify the ID of the peering connection.
To obtain the ID, you can look it up on the **Peering connections** page in the **VPC** console.
 - d. Click **Save changes**.

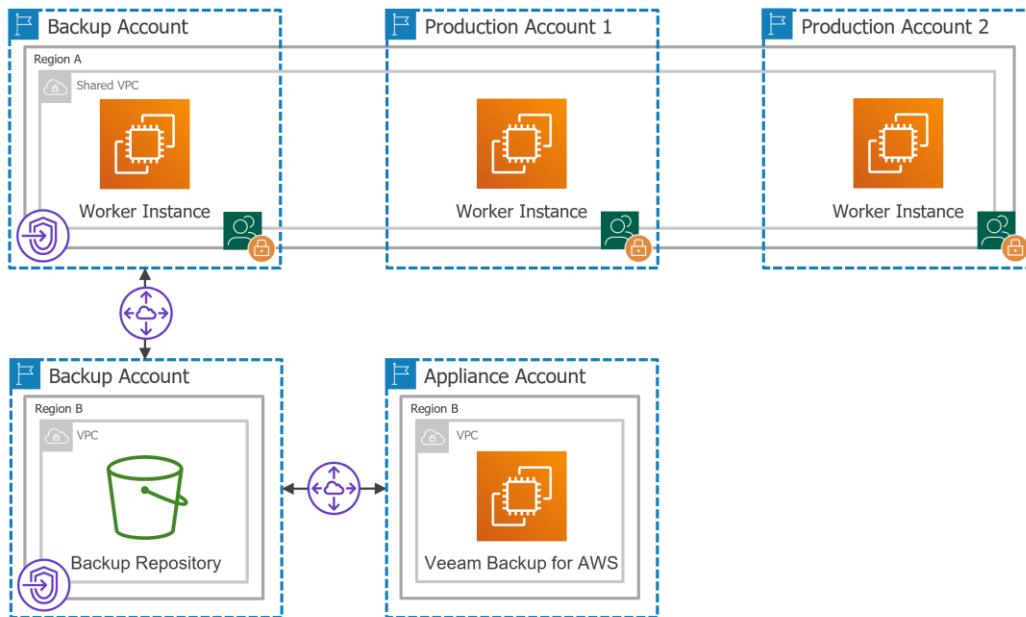
Configuring Private Networks for Production Accounts

If you have multiple AWS accounts and want to deploy worker instances in [production accounts](#), the estimated cost of VPC endpoints per account may occur to be significantly high. To reduce the cost, you can create a single resource share in one AWS account for all subnets to which the worker instances will be connected, and share the resource with other AWS accounts belonging to the same organization.

For Veeam Backup for AWS to be able to deploy worker instances in a private environment in production accounts, perform the following steps:

1. [Create VPC interface and S3 interface endpoints for subnets to which the worker instances will be connected](#).
2. [Create a peering connection between VPCs](#).
3. [Add routes to the route tables associated with the subnets of the VPCs](#).
4. [Create a resource share to share the subnets with other AWS accounts](#).

5. In each production account, create security groups that will be associated with worker instances connected to the shared subnets.



Step 1. Create Interface Endpoints

To allow Veeam Backup for AWS to create image-level backups of EC2 instances, to perform restore operations and to save EFS indexes to backup repositories, you must configure specific VPC interface endpoints for all subnets to which worker instances deployed for these operations will be connected. For the list of VPC interface endpoints required for backup and restore operations, see [Configuring Private Networks](#).

To deploy worker instances, Veeam Backup for AWS uses either the default or the [most appropriate network settings](#) of AWS Regions where the processed resources reside. However, you can add specific worker configurations as described in section [Managing Worker Configurations](#).

Creating Interface Endpoints

To create an interface VPC endpoint, do the following:

1. Log in to the **AWS Management Console** using credentials of an AWS account in which you want to create the endpoint.
2. Navigate to **Services > Networking & Content Delivery** and click **VPC**.
3. In the **VPC** console, navigate to **Virtual Private Cloud > Endpoints** and click **Create Endpoint**.
4. Complete the **Create endpoint** wizard:
 - a. At the **Endpoint settings** step, do the following:
 - i. [Optional] In the **Name tag** field, specify a name for the endpoint.
 - ii. In the **Service category** section, select the **AWS services** option.
 - b. At the **Services** step, enter *Interface* in the search field and choose a service for which you want to create a VPC endpoint.
 - c. At the **VPC** step, do the following:
 - i. From the **VPC** drop-down list, choose a VPC to which the deployed worker instances will be connected. Make sure that the **Enable DNS hostnames** check box is selected for the VPC.
 - ii. In the **Additional settings** section, select the **Enable DNS name** check box.
 - d. At the **Subnets** step, choose a subnet for each Availability Zone where the worker instances will be deployed, and specify the IP address type. Make sure that the **Auto-assign public IPv4 address** check box is not selected for the subnet.
 - e. At the **Security groups** step, choose security groups that will be associated with the endpoint network interfaces.

Ensure that each security group allows communication between the associated endpoint network interface and resources in your VPC communicating with the selected service. If a security group restricts inbound HTTPS traffic (port **443**) from the resources in the VPC, you will not be able to send traffic through the endpoint network interface.
 - f. At the **Policy** step, select the **Full access** option to allow full access to the service. Alternatively, select the **Custom** option, and attach a VPC endpoint policy that will control permissions required to access available resources over the VPC endpoint.
 - g. Click **Create Endpoint**.

For more information on interface VPC endpoints, see [AWS Documentation](#).

Creating S3 Interface Endpoints

To create an S3 interface VPC endpoint, do the following:

1. In the **VPC** console, navigate to **Virtual Private Cloud > Endpoints** and click **Create Endpoint**.
2. Complete the **Create endpoint** wizard:
 - a. At the **Endpoint settings** step, do the following:
 - i. [Optional] In the **Name tag** field, specify a name for the endpoint.
 - ii. In the **Service category** section, select the **AWS services** option.
 - b. At the **Services** step, enter *S3* in the search field and choose the `com.amazonaws.<region>.s3` service with the *Interface* type, where `<region>` is the name of an AWS Region in which a backup repository is located.
 - c. At the **VPC** step, choose a VPC to which the deployed worker instances will be connected.
 - d. At the **Subnets** step, choose a subnet for each Availability Zone where the worker instances will be deployed, and specify the IP address type.
 - e. At the **Security groups** step, choose security groups that will be associated with the endpoint network interface.
 - h. At the **Policy** step, select the **Full access** option to allow full access to the service. Alternatively, select the **Custom** option, and attach a VPC endpoint policy that will control permissions required to access available resources over the VPC endpoint.
 - f. Click **Create Endpoint**.

IMPORTANT

The backup appliance and worker instances must be able to communicate with the Amazon S3 service through the created S3 interface endpoint. That is why security groups associated with the endpoint network interface must allow inbound HTTPS traffic from both the backup appliance and the worker instances through port **443**.

For more information on interface endpoints for Amazon S3, see [AWS Documentation](#).

Step 2. Create VPC Peering Connection

If you have created interface endpoints and S3 interface endpoints in subnets of two different VPCs, you must create a peering connection between the acceptor and requester VPC to enable route traffic between those VPCs using private IP addresses.

To create a VPC peering connection, do the following:

1. In the **VPC** console, navigate to **Virtual Private Cloud > Peering connections** and click **Create peering connection**.
2. Complete the **Create peering connection** wizard:
 - a. At the **Peering connection settings** step, do the following:
 - i. [Optional] In the **Name** field, specify a name for the connection.
 - ii. In the **Select a local VPC to peer with** section, choose the requester VPC.
 - iii. In the **Select another VPC to peer with** section, choose an AWS account and AWS Region in which you want to create the connection, and specify the ID of the acceptor VPC.
 - iv. In the **Tags** section, specify AWS tags that will be assigned to the connection.
 - b. Click **Create Peering Connection**.
3. To enable route traffic between the requester and acceptor VPC, select the created peering connection in the **Peering connections** list and click **Actions > Accept request**.

Step 3. Configure Routing

If you have created a peering connection between two different VPCs, you must add routes to the route tables associated with the subnets of the acceptor and requester VPC to enable private traffic between those VPCs.

To add a route to a route table, do the following:

1. In the **VPC** console, navigate to **Virtual Private Cloud > Route tables**, choose the route table and click **Actions > Edit routes**.
2. Complete the **Edit routes** wizard:
 - a. Click **Add routes**.
 - b. In the **Destination** field, specify the range of IPv4 addresses to which the network traffic in the peering connection must be directed.
The IPv4 address range must be specified in the CIDR notation (for example, `12.23.34.0/24`).
 - c. In the **Target** field, select the **Peering Connection** option and specify the ID of the peering connection.
To obtain the ID, you can look it up on the **Peering connections** page in the **VPC** console.
 - d. Click **Save changes**.

Step 4. Create Resource Share

If you have multiple AWS accounts and want to deploy worker instances in [production accounts](#), you can create a single resource share in one AWS account for all subnets to which the worker instances will be connected. The resource share can be further used to share these subnets with other AWS accounts belonging to the same organization. For information, see [AWS Documentation](#).

To create a resource share, do the following:

1. Navigate to **Services > Security, Identity & Compliance** and click **Resource Access Manager**.
2. In the **Resource Access Manager** console, use the Region selector to choose an AWS Region in which the resource share will be created.
3. Navigate to **Shared by me > Resource shares** and click **Create resource share**.
4. Complete the **Create resource share** wizard:
 - a. At the **Specify resource share details** step, configure the following settings:
 - i. In the **Resource share** field, specify a name for the resource share.
 - ii. In the **Resources** section, enter *Subnets* in the search field and choose subnets that you want to share.
 - iii. In the **Tags** section, specify AWS tags that will be assigned to the resource share.
 - b. At the **Associate managed permissions** step, keep the default managed permissions associated with the specified subnets.
 - c. At the **Grant access to principal** step, use the **Principals** section to choose whether you want to share the subnets within your organization only. Then, select the AWS account option and specify the IDs of AWS accounts with which you want to share the subnets.

To obtain the IDs, you can either look them up in the AWS Management Console, or send a query to the AWS Command Line Interface (AWS CLI).
 - d. At the **Review and create** step, review the configured settings and click **Create resource share**.

Step 5. Create Security Groups

Security groups associated with shared subnets are not automatically propagated to other AWS accounts during resource sharing. That is why if you have created a single resource share in one AWS account for all subnets to which worker instances will be connected, you must create security groups in each production account – these groups will be associated with worker instances connected to the shared subnets.

To create a security group, do the following:

1. Log in to the **AWS Management Console** using credentials of an AWS account in which you want to create the security group.
2. Navigate to **Services > Networking & Content Delivery** and click **VPC**.
3. In the **VPC** console, navigate to **Security > Security Groups** and click **Create security group**.
4. Complete the **Create security group** wizard:
 - a. At the **Basic details** step, do the following:
 - i. In the **Security group name** and **Description** field, specify a name and description for the security group.
 - ii. In the **VPC** field, specify the ID of the VPC in which you want to create the security group.
To obtain the ID, you can look it up on the **Your VPCs** page in the VPC console.
 - b. At the **Inbound rules** step, do not specify any inbound rules.
 - c. At the **Outbound rules** step, specify rules to allow outbound HTTPS traffic to all VPC endpoints used by worker instances that will be connected to the shared subnets through port **443**.
 - d. At the **Tags** step, specify AWS tags that will be assigned to the security group.
 - e. Click **Create security group**.

IMPORTANT

After you create a security group, you must either add a new worker configuration or edit the network settings of an existing one to specify the created security group for each production account in which worker instances will be deployed. To learn how to do that, see [Adding Configurations for Production Accounts](#).

Data Encryption

By default, Amazon S3 Buckets are encrypted by default with Amazon S3 managed keys (SSE-S3). For more information on S3 encryption, see [AWS Documentation](#).

For enhanced data security, Veeam Backup for AWS allows you to encrypt backed-up data in backup repositories using Veeam encryption mechanisms. Additionally, Veeam Backup for AWS supports native AWS KMS encryption of EC2 and RDS instance volumes, and cloud-native snapshots, as well as encryption of EFS and FSx file systems and DynamoDB tables. To encrypt data, Veeam Backup for AWS uses the 256-bit Advanced Encryption Standard (AES). For more information about AES, see [Advanced Encryption Standard \(AES\)](#).

NOTE

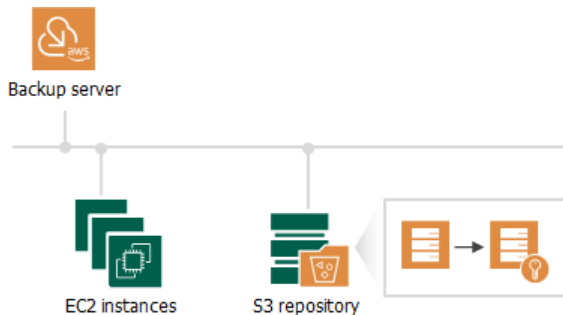
Sensitive customer data (credentials of user accounts required to connect to virtual servers and other systems, cloud credentials, and so on) is stored in the configuration database in the encrypted format.

Backup Repository Encryption

Veeam Backup for AWS allows you to enable encryption at the repository level. Veeam Backup for AWS encrypts backup files stored in backup repositories the same way as Veeam Backup & Replication encrypts backup files stored in backup repositories. To learn what algorithms Veeam Backup & Replication uses to encrypt backup files, see the Veeam Backup & Replication User Guide, section [Encryption Standards](#).

To enable encryption for a backup repository added to Veeam Backup for AWS, configure the repository settings as described in section [Adding Backup Repositories](#) and choose whether you want to encrypt data using a password or using a KMS encryption key. After you create a backup policy and specify the backup repository as a target location for EC2 image-level backups, RDS image-level backups, EFS indexing or VPC configuration backup copies, as described in sections [Creating EC2 Backup Policies](#), [Creating RDS Backup Policies](#), [Creating EFS Backup Policies](#) and [Editing VPC Configuration Backup Policy](#), Veeam Backup for AWS performs the following steps:

1. Based on the provided password or KMS key, generates an encryption key to protect backed-up data stored in the backup repository, and stores the key in the configuration database on the backup appliance.
2. Uses the generated key to encrypt backed-up data transferred to the backup repository when running the backup policy.



AWS KMS Encryption

NOTE

Veeam Backup for AWS does not use automatic AWS KMS key rotation for KMS keys, as well as AWS Secrets Manager for storing secrets.

Veeam Backup for AWS allows you to back up, replicate and restore data of EC2 and RDS instance volumes encrypted with [AWS KMS keys](#), as well as back up and restore EFS and FSx file systems, DynamoDB tables and Redshift clusters encrypted with AWS KMS keys. Additionally, you can encrypt unencrypted data and change KMS keys used to encrypt data when performing the following operations:

- [Creating EC2 instance snapshot replicas.](#)
- [Creating RDS instance snapshot replicas.](#)
- [Creating cloud-native snapshots of EC2 instances manually.](#)
- [Creating cloud-native snapshots of RDS instances manually.](#)
- [Restoring entire EC2 instances to a new location.](#)
- [Restoring entire RDS instances to a new location.](#)
- [Restoring EC2 instance volumes to a new location.](#)
- [Restoring entire EFS file systems to a new location.](#)
- [Restoring FSx file systems to a new location.](#)
- [Restoring DynamoDB tables to a new location.](#)
- [Restoring Redshift clusters to the original location.](#)

Depending on the operation performed for an encrypted RDS instance or an EC2 instance that has encrypted EBS volumes, the IAM role that Veeam Backup for AWS uses for the operation requires permissions to access various KMS keys:

- [Creating cloud-native snapshots](#)
- [Creating snapshot replicas](#)
- [Restoring from cloud-native snapshots](#)
- [Creating image-level backups](#)
- [Restoring from image-level backups](#)

IMPORTANT

If you back up, replicate or restore data of an unencrypted RDS instance or EC2 instance, and if you want to encrypt the backed-up or restored data, you must grant to the IAM role that Veeam Backup for AWS uses to perform the operation permissions to access only the KMS key with which you want to encrypt the data. To learn how to grant to an IAM role permissions to use a KMS key, see [this Veeam KB article](#).

Creating Cloud-Native Snapshots

The process of creating cloud-native snapshots of an EC2 instance with encrypted EBS volumes and an encrypted RDS instance does not differ from the same process for an instance with unencrypted volumes. The IAM role used to create cloud-native snapshots does not require any additional permissions – Veeam Backup for AWS encrypts these snapshots with the same KMS keys with which the source instance or volume is encrypted.

Creating Snapshot Replicas

The process of creating a snapshot replica of an encrypted RDS instance and an EC2 instance with encrypted EBS volumes differs depending on whether you create snapshot replicas within the same AWS account to which the instance belongs or not:

- [Creating the snapshot replica in the same AWS account to which the instance belongs.](#)
- [Creating the snapshot replica in an AWS account other than the AWS account to which the instance belongs.](#)

Creating Snapshot Replica in Same AWS Account

To create a snapshot replica in the same AWS account to which the encrypted EC2 or RDS instance belongs, Veeam Backup for AWS performs the following steps:

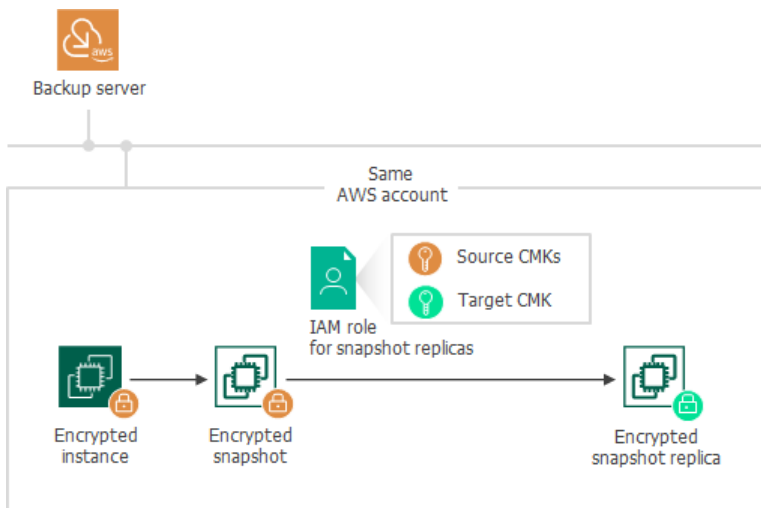
1. Takes an encrypted cloud-native snapshot of the instance.
2. Copies the created snapshot to the target AWS Region.

To copy the encrypted snapshot, Veeam Backup for AWS uses an IAM role specified at the **Targets** step of the **Add Policy** wizard, as described in sections [Creating EC2 Backup Policies](#) and [Creating RDS Backup Policies](#). The IAM role must have permissions to access the following KMS keys:

- KMS keys with which data of the source instance is encrypted (source KMS keys).
- A KMS key with which you want to encrypt instance data in the snapshot replica (target KMS key).

IMPORTANT

If you do not specify a target KMS key in the backup policy settings, Veeam Backup for AWS will not create a snapshot replica for the encrypted instance, and the backup session will complete with warnings.



Creating Snapshot Replica in Another AWS Account

The process of creating a snapshot replica differs depending on the AWS resource for which you want to create a snapshot replica:

- [Creating the snapshot replica in an AWS account other than the AWS account to which the EC2 instance belongs.](#)
- [Creating the snapshot replica in an AWS account other than the AWS account to which the RDS instance belongs.](#)

Creating Snapshot Replica of EC2 Instance

To create a snapshot replica in an AWS account other than the AWS account to which the EC2 instance with encrypted EBS volumes belongs, Veeam Backup for AWS performs the following steps:

1. Takes an encrypted cloud-native snapshot of the EC2 instance.
2. Shares the created snapshot with the target AWS account.

To share the encrypted snapshot, Veeam Backup for AWS uses an IAM role specified at the **Sources** step of the **Add Policy** wizard, as described in section [Creating EC2 Backup Policies](#). The IAM role must have permissions to access the KMS keys with which EBS volumes of the EC2 instance are encrypted (source KMS keys).

IMPORTANT

If EBS volumes of the EC2 instance are encrypted with the [default key for EBS encryption \(aws/ebs alias\)](#), Veeam Backup for AWS will not be able to share the snapshot with another AWS account and the replication process will fail to complete successfully. For more information, see [this Veeam KB article](#).

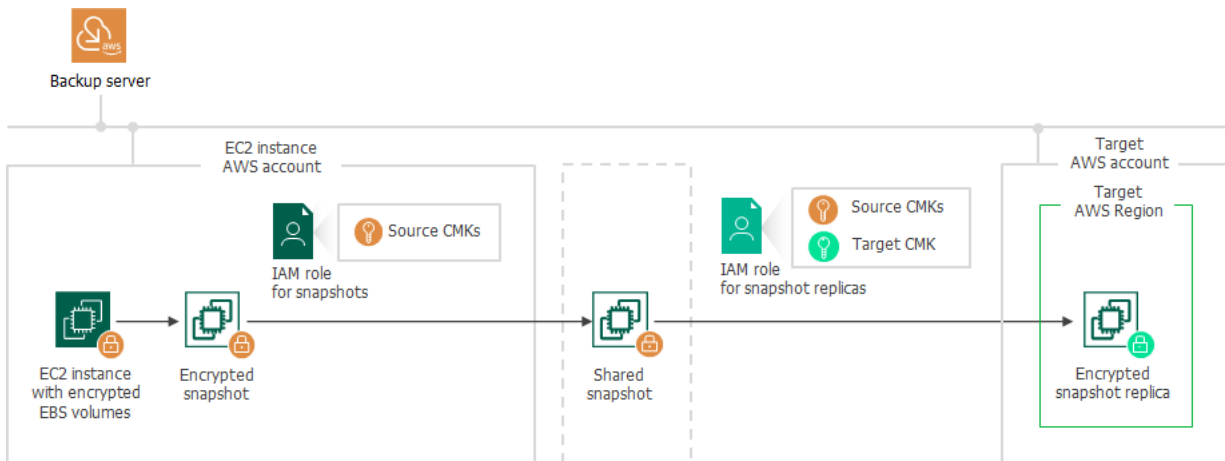
3. Copies the shared snapshot to the target AWS Region in the target AWS account.

To copy the shared encrypted snapshot, Veeam Backup for AWS uses an IAM role specified at the **Targets** step of the **Add Policy** wizard, as described in section [Creating EC2 Backup Policies](#). The IAM role must have permissions to access the following KMS keys:

- The KMS keys with which EBS volumes of the EC2 instance are encrypted (source KMS keys).
- A KMS key with which you want to encrypt EBS volume data in the snapshot replica (target KMS key).

IMPORTANT

Note that if you do not specify a target KMS key in the backup policy settings, Veeam Backup for AWS will not create a snapshot replica for the encrypted instance, and the backup session will complete with warnings.



Creating Snapshot Replica of RDS Instance

To create a snapshot replica in an AWS account other than the AWS account to which the encrypted RDS instance belongs, Veeam Backup for AWS performs the following steps:

1. Takes an encrypted cloud-native snapshot of the RDS instance.
2. Shares the created snapshot with the target AWS account.

To share the encrypted snapshot, Veeam Backup for AWS uses an IAM role specified at the **Sources** step of the **Add Policy** wizard, as described in section [Creating RDS Backup Policies](#). The IAM role must have permissions to access the KMS key with which the RDS instance is encrypted (source KMS key).

IMPORTANT

If the RDS instance is encrypted with the [default encryption key \(aws/rds alias\)](#), Veeam Backup for AWS will not be able to share the snapshot with another AWS account and the replication process will fail to complete successfully. For more information, see [this Veeam KB article](#).

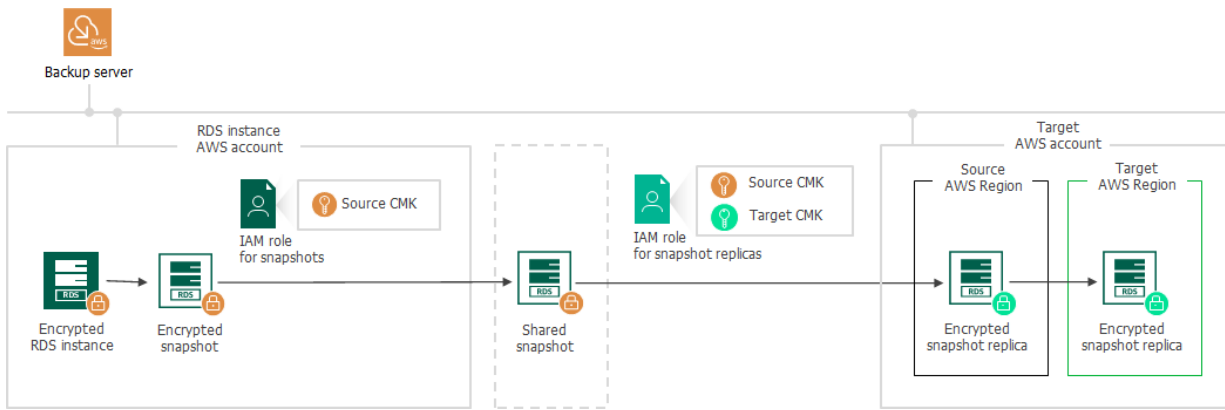
3. In the target AWS account, copies the shared encrypted snapshot to the same AWS Region to which the RDS instance belongs in the source AWS account. Then, if the target AWS Region differs from the source AWS Region, copies the shared snapshot to the target AWS Region.

To copy the shared encrypted snapshot, Veeam Backup for AWS uses an IAM role specified at the **Targets** step of the **Add Policy** wizard, as described in section [Creating RDS Backup Policies](#). The IAM role must have permissions to access the following KMS keys:

- The KMS key with which the RDS instance is encrypted (source KMS key).
- A KMS key with which you want to encrypt RDS instance data in the snapshot replica (target KMS key).

IMPORTANT

If you do not specify a target KMS key in the backup policy settings, Veeam Backup for AWS will not create a snapshot replica for the encrypted instance, and the backup session will complete with warnings.



Restoring From Snapshots and Replicas

The process of restoring an RDS or EC2 instance from an encrypted cloud-native snapshot differs depending on whether you perform restore to the original location where the cloud-native snapshot was stored or to a new location:

- [Restoring the instance to the original location where the snapshot resides.](#)
- [Restoring the instance to a new location.](#)

NOTE

Consider the following:

- An AWS account to which the cloud-native snapshot belongs is also referred to as the source AWS account.
- An AWS account to which you restore the instance is also referred to as the target AWS account.

Restoring to Original Location

To restore an EC2 or RDS instance to the location where the snapshot resides, Veeam Backup for AWS uses an IAM role specified for the restore operation, as described in sections [Performing Entire EC2 Instance Restore](#) and [Performing RDS Instance Restore](#). The IAM role must have permissions to access the following KMS keys:

- KMS keys with which the cloud-native snapshot is encrypted.
- A KMS key with which you want to encrypt data of the restored instance.

Restoring to New Location

The process of restoring to a new location differs depending on the AWS resource you want to restore and the specific use case:

- [Restoring the EC2 instance to another AWS Region in the same AWS account.](#)
- [Restoring the EC2 instance in another AWS account to the same AWS Region.](#)
- [Restoring the EC2 instance in another AWS account to another AWS Region.](#)
- [Restoring the RDS instance to another AWS Region in the same AWS account.](#)
- [Restoring the RDS instance in another AWS account to the same AWS Region.](#)

- [Restoring the RDS instance in another AWS account to another AWS Region.](#)


Restoring EC2 instance to Another AWS Region in Same AWS Account

To restore an EC2 instance to another AWS Region in the same AWS account to which the cloud-native snapshot belongs, Veeam Backup for AWS performs the following steps:

1. Copies the encrypted cloud-native snapshot to the target AWS Region.
2. Creates an EC2 instance in the target AWS Region.
3. Creates encrypted EBS volumes from the copied encrypted snapshot and attaches them to the created EC2 instance.

To copy the encrypted snapshot, and to create and encrypt EBS volumes, Veeam Backup for AWS uses an IAM role specified for the restore operation, as described in section [Performing Entire EC2 Instance Restore](#). The IAM role must have permissions to access the following KMS keys:

- KMS keys with which the cloud-native snapshot is encrypted (source KMS keys).
- A KMS key with which you want to encrypt EBS volumes of the restored EC2 instance (target KMS key).



Restoring EC2 Instance to Same AWS Region but in Another AWS Account

To restore an EC2 instance in another AWS account to the same AWS Region where the cloud-native snapshot resides, Veeam Backup for AWS performs the following steps:

1. Shares the encrypted cloud-native snapshot with the target AWS account.

To share the encrypted snapshot, Veeam Backup for AWS uses an IAM role specified in the backup policy settings [for creating cloud-native snapshots](#) (if you perform restore from a snapshot) or [for copying and storing snapshot replicas](#) (if you perform restore from a snapshot replica). The IAM role must have permissions to access the KMS key with which the cloud-native snapshot is encrypted (source KMS keys).


IMPORTANT

Due to AWS limitations, cloud-native snapshots encrypted with the [default key for EBS encryption \(aws/ebs alias\)](#) cannot be shared with other AWS accounts. Thus, if the cloud-native snapshot is encrypted with the default key for EBS encryption, Veeam Backup for AWS will not be able to share the snapshot and the restore process will fail to complete successfully. For more information, see [this Veeam KB article](#).

2. Creates an EC2 instance in the target AWS account in the same AWS Region where the snapshot resides in the source AWS account.
3. Creates encrypted EBS volumes from the shared encrypted snapshot and attaches them to the created EC2 instance.

To create and encrypt EBS volumes, Veeam Backup for AWS uses an IAM role specified for the restore operation, as described in section [Performing Entire EC2 Instance Restore](#). The IAM role must have permissions to access the following KMS keys:

- The KMS keys with which the cloud-native snapshot is encrypted (source KMS keys).
- A KMS key with which you want to encrypt EBS volumes of the restored EC2 instance (target KMS key).



Restoring EC2 Instance to Another AWS Region in Another AWS Account

To restore an EC2 instance to another AWS Region in an AWS account other than the AWS account to which the cloud-native snapshot belongs, Veeam Backup for AWS performs the following steps:

1. Shares the encrypted cloud-native snapshot with the target AWS account.

To share the encrypted snapshot, Veeam Backup for AWS uses an IAM role specified in the backup policy settings [for creating cloud-native snapshots](#) (if you perform restore from a snapshot) or [for copying and storing snapshot replicas](#) (if you perform restore from a snapshot replica). The IAM role must have permissions to access the following KMS keys:

- KMS keys with which the cloud-native snapshot is encrypted (source KMS keys).
- A KMS key with which you want to encrypt EBS volumes of the restored EC2 instance (target KMS key).

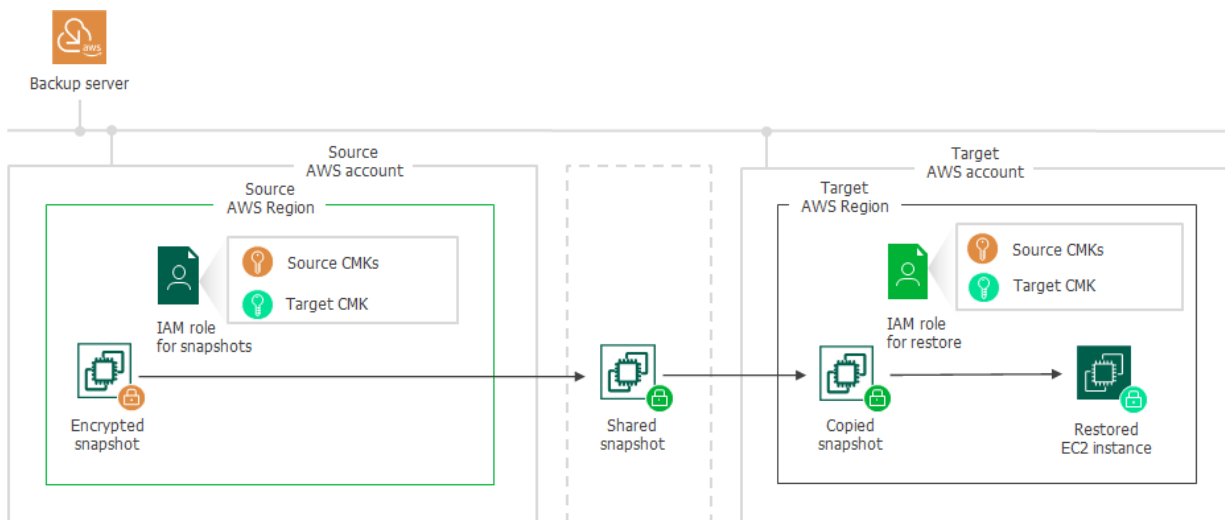
IMPORTANT

Due to AWS limitations, cloud-native snapshots encrypted with the [default key for EBS encryption \(aws/ebs alias\)](#) cannot be shared with other AWS accounts. Thus, if the cloud-native snapshot is encrypted with the default key for EBS encryption, Veeam Backup for AWS will not be able to share the snapshot and the restore process will fail to complete successfully. For more information, see [this Veeam KB article](#).

2. Copies the shared snapshot to the target AWS Region in the target AWS account.
3. Creates an EC2 instance in the target AWS Region in the target AWS account.
4. Creates encrypted EBS volumes from the shared encrypted snapshot and attaches them to the created EC2 instance.

To copy the snapshot, create and encrypt EBS volumes, Veeam Backup for AWS uses an IAM role specified for the restore operation, as described in section [Performing Entire EC2 Instance Restore](#). The IAM role must have permissions to access the following KMS keys:

- The KMS keys with which the cloud-native snapshot is encrypted (source KMS keys).
- The KMS key with which you want to encrypt EBS volumes of the restored EC2 instance (target KMS key).



Restoring RDS Instance to Another AWS Region in Same AWS Account


To restore an RDS instance to a another AWS Region in the same AWS account to which the cloud-native snapshot belongs, Veeam Backup for AWS performs the following steps:

1. Copies the encrypted cloud-native snapshot to the target AWS Region.
2. Creates an RDS instance from the copied encrypted snapshot in the target AWS Region.

To copy the encrypted snapshot, and to create the RDS instance, Veeam Backup for AWS uses an IAM role specified for the restore operation, as described in section [Performing RDS Instance Restore](#). The IAM role must have permissions to access the following KMS keys:

- A KMS key with which the cloud-native snapshot is encrypted (source KMS key).

- A KMS key with which you want to encrypt the restored RDS instance (target KMS key).



Restoring RDS Instance to Same AWS Region but in Another AWS Account

To restore an RDS instance in another AWS account to the same AWS Region where the cloud-native snapshot resides, Veeam Backup for AWS performs the following steps:

1. Shares the encrypted cloud-native snapshot with the target AWS account.

To share the encrypted snapshot, Veeam Backup for AWS uses an IAM role specified in the backup policy settings [for creating cloud-native snapshots](#) (if you restore from a snapshot) or [for copying and storing snapshot replicas](#) (if you restore from a snapshot replica). The IAM role must have permissions to access the KMS key with which the cloud-native snapshot is encrypted (source KMS key).

IMPORTANT

Due to AWS limitations, cloud-native snapshots encrypted with the [default encryption key \(aws/rds alias\)](#) cannot be shared with other AWS accounts. Thus, if the cloud-native snapshot is encrypted with the default encryption key, Veeam Backup for AWS will not be able to share the snapshot and the restore process will fail to complete successfully. For more information, see [this Veeam KB article](#).


2. In the target AWS account, copies the shared snapshot to the same AWS Region where the snapshot resides in the source AWS account, and re-encrypts the snapshot with the KMS keys that you specified to encrypt the restored RDS instance.

To copy the shared encrypted snapshot and to re-encrypt it, Veeam Backup for AWS uses an IAM role specified for the restore operation, as described in section [Performing RDS Instance Restore](#). The IAM role must have permissions to access the following KMS keys:

- The KMS key with which the cloud-native snapshot is encrypted (source KMS key).
- A KMS key with which you want to encrypt the restored RDS instance (target KMS key).

- Creates an encrypted RDS instance from the copied encrypted snapshot in the target AWS account in the same AWS Region where the snapshot resides in the source AWS account.

To create and encrypt the RDS instance, Veeam Backup for AWS uses an IAM role specified for the restore operation, as described in section [Performing RDS Instance Restore](#). The IAM role must have permissions to access the KMS key with which you want to encrypt the restored RDS instance (target KMS key).



Restoring RDS Instance to Another AWS Region in Another AWS Account

To restore an RDS instance to another AWS Region in an AWS account other than the AWS account to which the cloud-native snapshot belongs, Veeam Backup for AWS performs the following steps:

- Shares the encrypted cloud-native snapshot with the target AWS account.

To share the encrypted snapshot, Veeam Backup for AWS uses an IAM role specified in the backup policy settings [for creating cloud-native snapshots](#) (if you restore from a snapshot) or [for copying and storing snapshot replicas](#) (if you restore from a snapshot replica). The IAM role must have permissions to access the following KMS keys:

- A KMS key with which the cloud-native snapshot is encrypted (source KMS key).
- A KMS key with which you want to encrypt the restored RDS instance (target KMS key).

IMPORTANT

Due to AWS limitations, cloud-native snapshots encrypted with the [default encryption key \(aws/rds alias\)](#) cannot be shared with other AWS accounts. Thus, if the cloud-native snapshot is encrypted with the default encryption key, Veeam Backup for AWS will not be able to share the snapshot and the restore process will fail to complete successfully. For more information, see [this Veeam KB article](#).

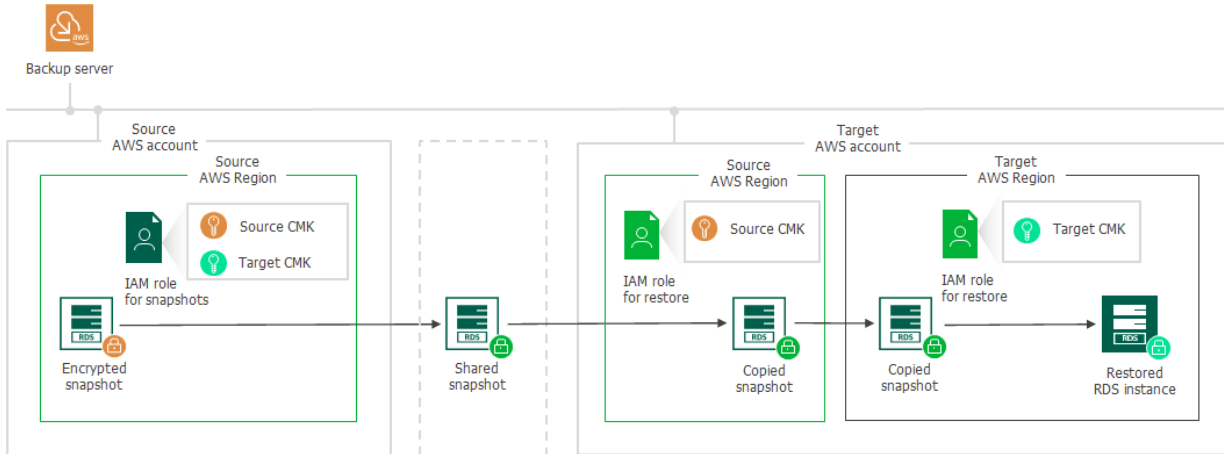
- In the target AWS account, copies the shared snapshot to the same AWS Region where the snapshot resides in the source AWS account.

To copy the shared encrypted snapshot, Veeam Backup for AWS uses an IAM role specified for the restore operation, as described in section [Performing RDS Instance Restore](#). The IAM role must have permissions to access the KMS key with which the cloud-native snapshot is encrypted (source KMS key).

- Copies the copied encrypted snapshot to the target AWS Region in the target AWS account and re-encrypts the snapshot with the KMS key specified to encrypt the restored RDS Instance.

- Creates an encrypted RDS instance in the target AWS Region in the target AWS account.

To copy and re-encrypt the snapshot, create and encrypt the RDS instance, Veeam Backup for AWS uses an IAM role specified for the restore operation, as described in section [Performing RDS Instance Restore](#). The IAM role must have permissions to access the KMS key with which you want to encrypt the restored RDS instance (target KMS key).



Creating Image-Level Backups

The process of creating an image-level backup of an EC2 instance with encrypted EBS volumes differs depending on whether a worker instance processing EBS volume data is deployed in the same AWS account or not:

- Creating the image-level backup in the same AWS account where the worker instance is deployed.
- Creating the image-level backup in an AWS account other than the AWS account where the worker instance is deployed.

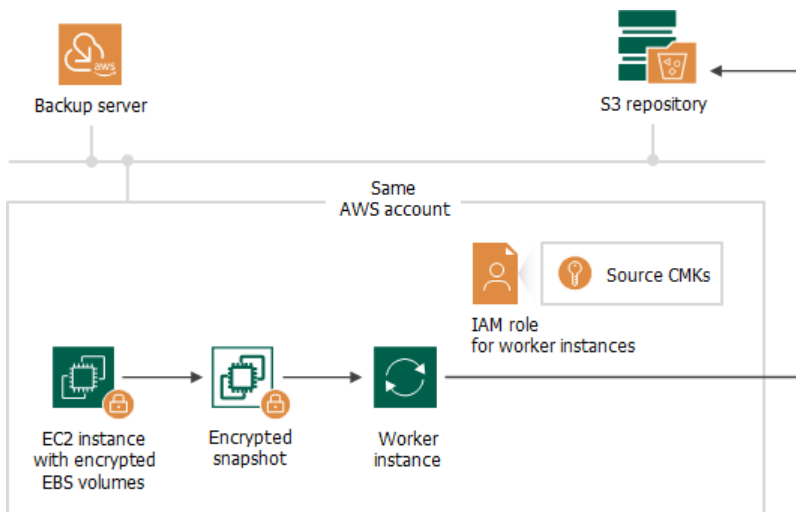
Creating Image-Level Backup in Same AWS Account

If a worker instance is deployed in the same AWS account to which the processed EC2 instance belongs, Veeam Backup for AWS performs the following steps:

- Creates an encrypted cloud-native snapshot of the EC2 instance.

2. Creates encrypted EBS volumes from the snapshot, and then attaches them to the worker instance for reading and further transferring EBS volume data to a backup repository.

To access the data, Veeam Backup for AWS uses an IAM role specified to deploy worker instances, as described in section [Managing Worker Configurations](#). The IAM role must have permissions to access the KMS keys with which EBS volumes of the EC2 instance are encrypted (source KMS keys).



Creating Image-Level Backup in Another AWS Account

If a worker instance is deployed in an AWS account other than the AWS account to which the processed EC2 instance belongs, Veeam Backup for AWS performs the following steps:

1. Creates an encrypted cloud-native snapshot of the EC2 instance.
2. Shares the created snapshot with the AWS account where the worker instance is deployed.

To share the encrypted snapshot, Veeam Backup for AWS uses the IAM role specified at the **Sources** step of the **Add Policy** wizard, as described in section [Creating EC2 Backup Policies](#). The IAM role must have permissions to access the KMS keys with which EBS volumes of the EC2 instance are encrypted (source KMS keys).

IMPORTANT

If EBS volumes of the EC2 instance are encrypted with the [default key for EBS encryption \(aws/ebc alias\)](#), Veeam Backup for AWS will not be able to share the snapshot with another AWS account and the backup process will fail to complete successfully. To work around the issue, enable the worker deployment in production accounts functionality, as described in [Creating EC2 Backup Policies](#).

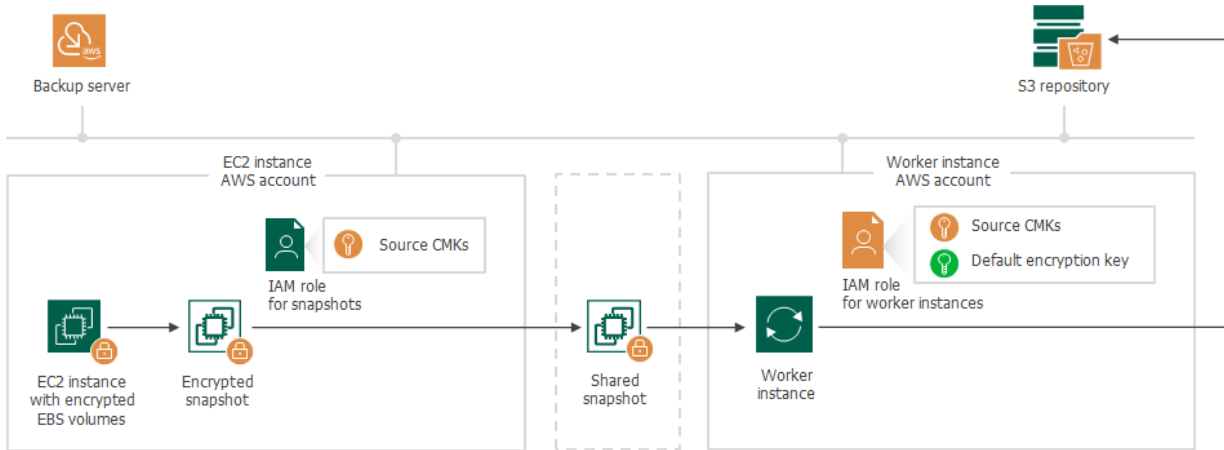
3. Creates encrypted EBS volumes from the shared encrypted snapshot, and then attaches them to the worker instance for reading and further transferring EBS volume data to a backup repository.

Due to AWS requirements, EBS volumes created from encrypted snapshots must also be encrypted. Thus, Veeam Backup for AWS encrypts re-created EBS volumes with the [default encryption key](#) specified for the AWS Region where the worker instance is deployed.

To access the data, Veeam Backup for AWS uses an IAM role specified to deploy worker instances, as described in section [Managing Worker Configurations](#). The IAM role must have permissions to access the following KMS keys:

- The KMS keys with which EBS volumes of the EC2 instance are encrypted (source KMS keys).

- The default encryption key specified for the AWS Region where the worker instance is deployed.



Restoring From Image-Level Backups

The process of restoring an EC2 instance with encrypted EBS volumes from an image-level backup differs depending on whether a worker instance is deployed in the same AWS account to which you perform restore or not:

- Performing restore from the image-level backup to the AWS account where the worker instance is deployed.
- Performing restore from the image-level backup to an AWS account other than the AWS account where the worker instance is deployed.

NOTE

Consider the following:

- An AWS account to which an IAM role specified for deploying worker instances belongs is also referred to as the source AWS account.
- An AWS account to which you restore an instance is also referred to as the target AWS account.
- To perform EC2 instance restore operations from image-level backups, Veeam Backup for AWS deploys worker instances in a target AWS Region specified in the restore settings.

Restore to Same AWS Account

If a worker instance is deployed in the same AWS account to which the restored EC2 instance will belong, to encrypt EBS volumes of the restored EC2 instance, Veeam Backup for AWS uses an IAM role specified to deploy worker instances, as described in section [Configuring Worker Instance Settings](#). The IAM role must have permissions to access to the KMS key with which you want to encrypt EBS volumes of the restored EC2 instance.

Restore to Another AWS Account

If a worker instance is deployed in an AWS account other than the AWS account to which the restored EC2 instance will belong, Veeam Backup for AWS performs the following steps:

1. Creates empty EBS volumes in the target AWS Region in the source AWS account and attaches them to the worker instance. To protect data that will be restored to these volumes, Veeam Backup for AWS encrypts the created EBS volumes with the [default encryption key](#) specified for the target AWS Region.

To encrypt the volumes, Veeam Backup for AWS uses an IAM role specified to deploy worker instances, as described in section [Managing Worker Configurations](#). The IAM role must have permissions to access to the default encryption key specified for the target AWS Region in the source AWS account.

2. Restores backed-up data to the empty EBS volumes on the worker instance.
3. Creates an encrypted cloud-native snapshot of the EBS volumes with the restored data.
4. Shares the created snapshot with the target AWS account.

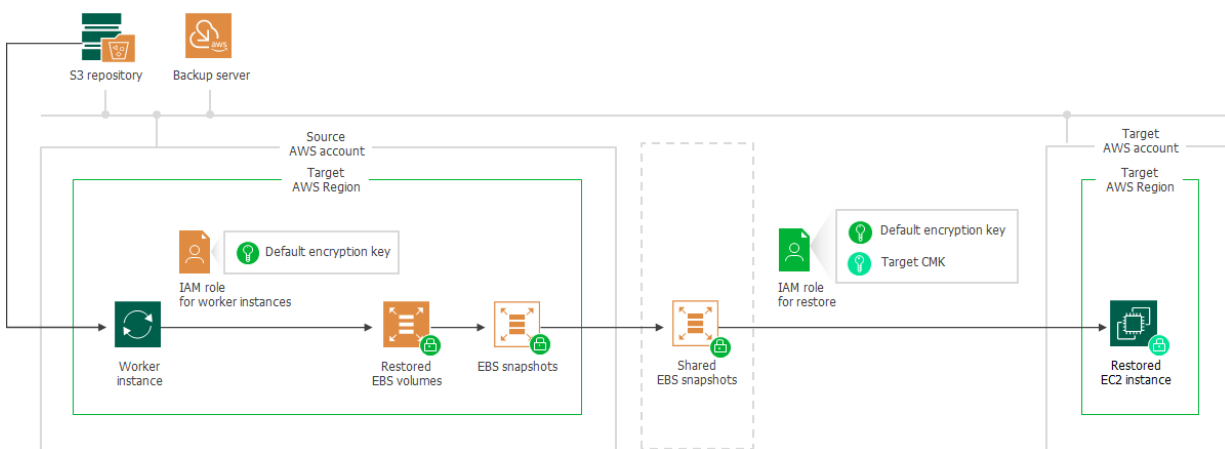
IMPORTANT

Due to AWS limitations, cloud-native snapshots encrypted with the [default key for EBS encryption \(aws/ebs alias\)](#) cannot be shared with other AWS accounts. Thus, if the default encryption key specified for the target AWS Region in the source AWS account is the default key for EBS encryption, Veeam Backup for AWS will not be able to share the snapshot and the restore process will fail to complete successfully. For more information, see [this Veeam KB article](#).

5. Creates an EC2 instance in the target AWS Region within the target AWS account.
6. Creates encrypted EBS volumes from the shared encrypted snapshot and attaches them to the created EC2 instance.

To create and encrypt EBS volumes, Veeam Backup for AWS uses an IAM role specified for the restore operation, as described in section [Performing Entire EC2 Instance Restore](#). The IAM role must have permissions to access the following KMS keys:

- The default encryption key specified for the target AWS Region in the source AWS account.
- A KMS key with which you want to encrypt EBS volumes of the restored EC2 instance (target KMS key).



Planning and Preparation

Before you start using Veeam Backup for AWS, consider the following requirements:

- [Hardware and software requirements.](#)
- [Network ports that must be open to ensure proper communication of Veeam Backup for AWS components.](#)
- [AWS services to which Veeam Backup for AWS must have outbound internet access.](#)
- [Permissions that must be assigned to accounts used to perform operations started using the Veeam Backup & Replication console.](#)
- [IAM permissions that must be assigned to IAM roles or IAM users used to perform operations started using the Web UI.](#)
- [Considerations and limitations that should be kept in mind before you deploy Veeam Backup for AWS.](#)
- [Sizing and Scalability Guidelines.](#)

System Requirements

When you plan to install AWS Plug-in for Veeam Backup & Replication, consider the following hardware and software requirements.

Backup Server

The machine where AWS Plug-in for Veeam Backup & Replication will run must meet system requirements described in the Veeam Backup & Replication User Guide, section [System Requirements](#). Additionally, the following software must be installed:

- Microsoft .NET Core Runtime 8.0.6
- Microsoft ASP.NET Core Shared Framework 8.0.6

IMPORTANT

If the version of Microsoft .NET Core Runtime differs from the version of Microsoft ASP.NET Core Shared Framework, AWS Plug-in for Veeam Backup & Replication services will not be able to start.

AWS Services

The backup appliance and worker instances must have outbound internet access to a number of AWS services. For the list of services, see [AWS Services](#).

Web Browsers

Internet Explorer is not supported. To access Veeam Backup for AWS, use Microsoft Edge (latest version), Mozilla Firefox (latest version) or Google Chrome (latest version).

Veeam Backup & Replication

AWS Plug-in for Veeam Backup & Replication version 12.8.0 supports integration with Veeam Backup & Replication version 12.1.2

Veeam Backup for AWS

AWS Plug-in for Veeam Backup & Replication version 12.8.0 supports integration with Veeam Backup for AWS version 8.x.

Version Compatibility

The following table lists compatible versions of Veeam Backup & Replication, AWS Plug-in for Veeam Backup & Replication and Veeam Backup for AWS.

Veeam Backup & Replication Build	AWS Plug-in for Veeam Backup & Replication Version	Veeam Backup for AWS Build	Backup Appliance OS Version
12.1.2.172	12.8.0.3264	8.0.0.845	Ubuntu 22.04 LTS
12.1.0.2131	12.7.0.1255	7.0.0, 7.0.1	
12.0.0.1420	12.2.6.5	6.1.0, 6.1.1, 6.1.2	
	12.1.6.93		
	12.0.6.956	6.0.0, 6.0.1, 6.0.2	Ubuntu 18.04 LTS
11.0.1.1261 including all cumulative patches starting from P20211211 (CP3)	11.0.5.553	5.0.0, 5.1.0, 5.1.1	
11.0.1.1261 including all cumulative patches prior to P20211211 (CP3)	11.0.4.305	4.0.0, 4.1.0, 4.1.1	
11.0.0.837	11.0.3.1132	3.0.0, 3.1.0, 3.1.1	
10.0.1.4854	10.0.3.825	3.0.0, 3.1.0, 3.1.1	
	10.0.1.661	2.0.0, 2.0.1	

Ports

As AWS Plug-in for Veeam Backup & Replication is installed on the same machine where Veeam Backup & Replication runs, it uses the same ports as those described in the Veeam Backup & Replication User Guide, section [Ports](#). In addition, AWS Plug-in for Veeam Backup & Replication also uses ports listed in the following table.

From	To	Protocol	Port	Notes
Web browser (local machine)	Backup appliance	TCP/HTTPS	443	Required to access the Web UI component from a user workstation.
		SSH	22	[Optional] Required to connect to the backup appliance using SSH.
		TCP/HTTPS	11005	[Optional] Default port required to communicate with the public REST API service running on the backup appliance. For more information on Veeam Backup for AWS REST API, see the Veeam Backup for AWS REST API Reference . To learn how to change the port number, see the Configuring Security Settings section in the Veeam Backup for AWS REST API Reference.
	Worker instances	TCP/HTTPS	443	Required to access the file-level recovery browser running on a worker instance during the file-level recovery process.
Backup appliance	SMTP server	TCP/SMTP	25	Default port used for sending email notifications.
	Veeam Update Repository (repository.veeam.com)	TCP/HTTPS	443	Required to download available product updates, worker deployment packages and restore utilities.
	Ubuntu Security Repository and OS Update repository (security.ubuntu.com, archive.ubuntu.com)	TCP/HTTP	80	Required to get OS security updates.
	Microsoft Package Repository (packages.microsoft.com, dotnetcli.blob.core.windows.net)	TCP/HTTPS	443	Required to get .NET package updates.

From	To	Protocol	Port	Notes
	PostgreSQL Apt Repository (apt.postgresql.org)	TCP/HTTP	80	Required to get PostgreSQL updates.
	PostgreSQL Website (postgresql.org)	TCP/HTTP	443	Required to download the repository key https://www.postgresql.org/media/keys/ACCC4CF8.asc .
	AWS services	TCP/HTTP	443	Required to perform data protection and disaster recovery operations.
	Route 53 Resolver	UDP	53	[Optional] Default port required to perform DNS resolution if you plan to use a custom DNS server for your VPC.
Worker instances	AWS services	TCP/HTTP	443	Required to perform data protection and disaster recovery operations.
	Route 53 Resolver	UDP	53	[Optional] Default port required to perform DNS resolution if you plan to use a custom DNS server for your VPC.
AWS Plug-in for Veeam Backup & Replication	Backup appliance, AWS services	TCP/HTTP	443	Port used for communication with AWS and Veeam Backup for AWS.
	Backup server	TCP	6172	Port used by AWS Plug-in for Veeam Backup & Replication to connect to a component that enables communication with the Veeam Backup & Replication database.
Veeam Backup & Replication console and Veeam ONE server	AWS Plug-in for Veeam Backup & Replication	TCP	9402	Port used to connect to AWS Plug-in for Veeam Backup & Replication.

To open network ports, you must add rules to security groups associated with Veeam Backup for AWS components:

- A security group associated with the backup appliance. For more information, see [Deploying Appliance from Console](#).

- Security groups associated with worker instances. For more information, see [Managing Worker Configurations](#).

To learn how to add security groups rules, see [AWS Documentation](#).

AWS Services

To perform backup and restore operations, the [AWS Plug-In for Veeam Backup & Replication](#), [backup appliance](#) and [worker instances](#) must have outbound internet access to the following AWS services.

AWS Services Required For AWS Plug-In for Veeam Backup & Replication

- [Amazon CloudWatch](#)
- [Amazon Data Lifecycle Manager](#)
- [Amazon Elastic Compute Cloud \(EC2\)](#)
- [Amazon Simple Storage Service \(S3\)](#)
- [AWS Identity and Access Management \(IAM\)](#)
- [AWS Key Management Service \(KMS\)](#)
- [AWS Systems Manager \(SSM\)](#)
- [AWS Security Token Service \(STS\)](#)
- [AWS Service Quotas](#)

AWS Services Required For Backup Appliance

- [Amazon CloudWatch](#)
- [Amazon CloudWatch Events](#)
- [Amazon Elastic Block Store \(EBS\)](#)
- [Amazon Elastic Compute Cloud \(EC2\)](#)
- [Amazon Kinesis Data Streams](#)
- [Amazon Relational Database Service \(RDS\)](#)
- [Amazon Redshift](#)
- [Amazon Elastic File System \(EFS\)](#)
- [Amazon FSx File Systems](#)
- [Amazon DynamoDB](#)
- [Amazon Simple Notification Service \(SNS\)](#)
- [Amazon Simple Queue Service \(SQS\)](#)
- [Amazon Simple Storage Service \(S3\)](#)
- [AWS Identity and Access Management \(IAM\)](#)
- [AWS Key Management Service \(KMS\)](#)
- [AWS Marketplace Metering Service](#)


- [AWS Resource Access Manager](#)
- [AWS Security Token Service \(STS\)](#)
- [AWS Service Quotas](#)
- [AWS Backup](#)
- [AWS Systems Manager \(SSM\)](#), including access to the *ec2messages* and *ssmmessages* endpoints
- [Elastic Load Balancing \(ELB\)](#)

AWS Services Required For Worker Instances

- [AWS Systems Manager \(SSM\)](#), including access to the *ec2messages* and *ssmmessages* endpoints
- [Amazon Simple Queue Service \(SQS\)](#)
- [Amazon Simple Storage Service \(S3\)](#)
- [Amazon Elastic Block Store \(EBS\)](#)
- [Amazon Kinesis Data Streams](#)

IMPORTANT

Proxy redirect and setting a proxy in the Veeam Backup for AWS configuration are not supported. Therefore, make sure that the security group associated with the backup appliance and worker instances allow direct network traffic required to communicate with the AWS services.



Plug-In Permissions

To perform backup and restore operations, accounts that AWS Plug-in for Veeam Backup & Replication uses to perform data protection and disaster recovery operations must be granted the following permissions.

Veeam Backup & Replication User Account Permissions

A user account that you plan to use when installing and working with Veeam Backup & Replication must have permissions described in the Veeam Backup & Replication User Guide, section [Installing and Using Veeam Backup & Replication](#).

Veeam Backup for AWS User Account Permissions

A user account that Veeam Backup & Replication will use to authenticate against the backup appliance and get access to the appliance functionality must be assigned the Portal Administrator role. For more information on user roles, see [Managing User Accounts](#).

NOTE

When you deploy a backup appliance from the Veeam Backup & Replication console, Veeam Backup & Replication will automatically create the necessary user account that will be assigned all the required permissions.

AWS IAM User Permissions

AWS Plug-in for Veeam Backup & Replication requires the following [IAM identities](#):

- An IAM user whose permissions are used to create, connect and manage backup appliances. To be able to perform these operations, the specified IAM user must have the following set of permissions:
 - List of permissions to deploy a new backup appliance

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "CreateResources",
      "Effect": "Allow",
      "Action": [
        "cloudwatch:PutMetricAlarm",
        "ec2:AllocateAddress",
        "ec2:AssociateAddress",
        "ec2:AttachInternetGateway",
        "ec2:AuthorizeSecurityGroupIngress",
        "ec2:CreateInternetGateway",
        "ec2:CreateKeyPair",
        "ec2:CreateRoute",
        "ec2:CreateSecurityGroup",
        "ec2:CreateSnapshot",
        "ec2:CreateSubnet",
        "ec2:CreateTags",
        "ec2:CreateVpc",
        "ec2>DeleteSnapshot",
        "ec2:DescribeAddresses",
        "ec2:DescribeAvailabilityZones",
        "ec2:DescribeIamInstanceProfileAssociations",
        "ec2:DescribeInstanceAttribute",
        "ec2:DescribeImages",
        "ec2:DescribeInstanceTypes",
        "ec2:DescribeInstances",
        "ec2:DescribeInternetGateways",
        "ec2:DescribeKeyPairs",
        "ec2:DescribeManagedPrefixLists",
        "ec2:DescribeRegions",
        "ec2:DescribeRouteTables",
        "ec2:DescribeSecurityGroups",
        "ec2:DescribeSnapshots",
        "ec2:DescribeSubnets",
        "ec2:DescribeVolumes",
        "ec2:DescribeVpcs",
        "ec2:ModifyInstanceAttribute",
        "ec2:GetManagedPrefixListEntries",
        "ec2:ModifyNetworkInterfaceAttribute",
        "ec2:ModifySubnetAttribute",
        "ec2:ModifyVpcAttribute",
        "ec2:RunInstances",
        "iam:AddRoleToInstanceProfile",
        "iam:AttachRolePolicy",
        "iam:CreateInstanceProfile",
        "iam:CreatePolicy",
        "iam:CreatePolicyVersion",
        "iam:CreateRole",
        "iam:CreateServiceLinkedRole",
        "iam:DetachRolePolicy",
        "iam:GetAccountSummary",
        "iam:GetInstanceProfile",

```

```

        "iam:GetPolicy",
        "iam:GetPolicyVersion",
        "iam:GetRole",
        "iam:ListAttachedRolePolicies",
        "iam:ListInstanceProfilesForRole",
        "iam:ListPolicyVersions",
        "iam:PassRole",
        "iam:PutRolePolicy",
        "iam:SimulatePrincipalPolicy",
        "s3:CreateBucket",
        "s3:DeleteBucket",
        "s3:DeleteObject",
        "s3:GetObject",
        "s3:ListAllMyBuckets",
        "s3:ListBucket",
        "s3:PutObject",
        "servicequotas:ListServiceQuotas",
        "ssm:DescribeInstanceInformation",
        "ssm:GetCommandInvocation",
        "ssm:GetParameter",
        "ssm:SendCommand"
    ],
    "Resource": "*"
},
{
    "Sid": "CleanupResources",
    "Effect": "Allow",
    "Action": [
        "cloudwatch:DeleteAlarms",
        "ec2:DeleteInternetGateway",
        "ec2:DeleteSecurityGroup",
        "ec2:DeleteSubnet",
        "ec2:DeleteVpc",
        "ec2:DetachInternetGateway",
        "ec2:DisassociateAddress",
        "ec2:ReleaseAddress",
        "ec2:TerminateInstances",
        "iam:DeleteInstanceProfile",
        "iam:DeletePolicy",
        "iam:DeletePolicyVersion",
        "iam:DeleteRole",
        "iam:DeleteRolePolicy",
        "iam:DetachRolePolicy",
        "iam:RemoveRoleFromInstanceProfile"
    ],
    "Resource": [
        "*"
    ]
}
]
}

```


NOTE

If EBS encryption is enabled for your AWS account, all the newly created volumes will be encrypted using the default KMS key specified in the EC2 console. Therefore, for the IAM user to be able to encrypt the EBS volumes of the appliance, the following conditions must be met:

- The symmetric encryption KMS key must be specified in the EC2 console.
- The IAM user must have permissions to use the default KMS key.
- The IAM user must be assigned the `kms:GenerateDataKeyWithoutPlaintext` permission.

For more information on EBS encryption, see [AWS Documentation](#).

› List of permissions to connect an existing backup appliance

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "ec2:AttachVolume",
        "ec2:CreateSnapshot",
        "ec2:CreateVolume",
        "ec2:DescribeAddresses",
        "ec2:DescribeAvailabilityZones",
        "ec2:DescribeInstances",
        "ec2:DescribeRegions",
        "ec2:DescribeVolumes",
        "ec2:DescribeSnapshots",
        "ec2:DescribeIamInstanceProfileAssociations",
        "ec2:DescribeInstanceAttribute",
        "ec2:DescribeImages",
        "ec2:DescribeVolumeAttribute",
        "ec2>DeleteSnapshot",
        "ec2:DetachVolume",
        "ec2>DeleteVolume",
        "ec2:ModifyInstanceAttribute",
        "ec2:RunInstances",
        "ec2:StopInstances",
        "ec2:StartInstances",
        "ec2:TerminateInstances",
        "iam:AddRoleToInstanceProfile",
        "iam:AttachRolePolicy",
        "iam:CreateInstanceProfile",
        "iam:CreatePolicy",
        "iam:CreatePolicyVersion",
        "iam:GetAccountSummary",
        "iam:GetPolicy",
        "iam:GetPolicyVersion",
        "iam:GetRole",
        "iam:GetInstanceProfile",
        "iam>ListAttachedRolePolicies",
        "iam>ListInstanceProfilesForRole",
        "iam>ListRolePolicies",
        "iam:PutRolePolicy",
        "iam:SimulatePrincipalPolicy",
        "iam>ListPolicyVersions",
        "iam:UpdateAssumeRolePolicy",
        "sts:GetCallerIdentity"
      ],
      "Resource": "*"
    }
  ]
}
```

› List of permissions to add a repository

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "ec2:DescribeRegions",
        "ec2:DescribeAddresses",
        "ec2:DescribeInstances",
        "iam:GetRole",
        "iam:SimulatePrincipalPolicy",
        "s3:CreateBucket",
        "s3:DeleteObject",
        "s3:DeleteObjectVersion",
        "s3:GetBucketLocation",
        "s3:GetBucketVersioning",
        "s3:GetBucketObjectLockConfiguration",
        "s3:GetObject",
        "s3:GetObjectVersion",
        "s3:GetObjectRetention",
        "s3:ListBucket",
        "s3:ListAllMyBuckets",
        "s3:ListBucketVersions",
        "s3:PutBucketVersioning",
        "s3:PutBucketObjectLockConfiguration",
        "s3:PutObjectRetention",
        "s3:PutObject"
      ],
      "Resource": "*"
    }
  ]
}
```

› List of permissions to encrypt repositories using AWS KMS keys

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "kms:Decrypt",
        "kms:DescribeKey",
        "kms:Encrypt",
        "kms:ListAliases",
        "kms:ListKeys"
      ],
      "Resource": "*"
    }
  ]
}
```

> List of permissions to upgrade backup appliance to version 8.0

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "iam:GetRole",
        "iam:SimulatePrincipalPolicy",
        "ec2:AttachVolume",
        "ec2:CreateVolume",
        "ec2:CreateSnapshot",
        "ec2:DescribeAddresses",
        "ec2:DescribeInstances",
        "ec2:DescribeVolumes",
        "ec2:DescribeSnapshots",
        "ec2:DescribeAvailabilityZones",
        "ec2:DescribeRegions",
        "ec2:DetachVolume",
        "ec2>DeleteVolume",
        "ec2:DescribeVolumeAttribute",
        "ec2>DeleteSnapshot",
        "ec2:DescribeInstanceAttribute",
        "ec2:DescribeImages",
        "ec2:ModifyInstanceAttribute",
        "ec2:RunInstances",
        "ec2:StartInstances",
        "ec2:StopInstances",
        "ec2:TerminateInstances",
        "ec2:RunInstances",
        "sts:GetCallerIdentity"
      ],
      "Resource": "*"
    }
  ]
}
```

NOTE

For Veeam Backup & Replication to be able to upgrade permissions of the [Default Backup Restore IAM role](#) when upgrading to version 8.0, add the necessary permissions listed below to the IAM policy.

➤ List of permissions to upgrade the *Default Backup Restore* IAM role

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "iam:AttachUserPolicy",
        "iam:AddRoleToInstanceProfile",
        "iam:CreatePolicyVersion",
        "iam:CreateInstanceProfile",
        "iam:GetAccountSummary",
        "iam:GetInstanceProfile",
        "iam:GetPolicyVersion",
        "iam:ListAttachedRolePolicies",
        "iam:ListPolicyVersions",
        "iam:ListInstanceProfilesForRole",
        "iam:ListRolePolicies",
        "iam:PutRolePolicy",
        "iam:UpdateAssumeRolePolicy"
      ],
      "Resource": "*"
    }
  ]
}
```

➤ Full list of permissions

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "cloudwatch:DeleteAlarms",
        "cloudwatch:PutMetricAlarm",
        "ec2:AllocateAddress",
        "ec2:AssociateAddress",
        "ec2:AttachInternetGateway",
        "ec2:AttachVolume",
        "ec2:AuthorizeSecurityGroupIngress",
        "ec2:CreateInternetGateway",
        "ec2:CreateKeyPair",
        "ec2:CreateRoute",
        "ec2:CreateSecurityGroup",
        "ec2:CreateSnapshot",
        "ec2:CreateSubnet",
        "ec2:CreateTags",
        "ec2:CreateVolume",
        "ec2:CreateVpc",
        "ec2>DeleteInternetGateway",
        "ec2>DeleteSecurityGroup",
        "ec2>DeleteSnapshot",
        "ec2>DeleteSubnet",
        "ec2>DeleteVolume",
        "ec2>DeleteVpc",
        "ec2:DescribeAddresses",
        "ec2:DescribeAvailabilityZones",
        "ec2:DescribeIamInstanceProfileAssociations",
        "ec2:DescribeImages",
        "ec2:DescribeInstanceAttribute",
        "ec2:DescribeInstances",
        "ec2:DescribeInstanceTypes",
        "ec2:DescribeInternetGateways",
        "ec2:DescribeKeyPairs",
        "ec2:DescribeManagedPrefixLists",
        "ec2:DescribeRegions",
        "ec2:DescribeRouteTables",
        "ec2:DescribeSecurityGroups",
        "ec2:DescribeSnapshots",
        "ec2:DescribeSubnets",
        "ec2:DescribeVolumeAttribute",
        "ec2:DescribeVolumes",
        "ec2:DescribeVpcs",
        "ec2:DetachInternetGateway",
        "ec2:DetachVolume",
        "ec2:DisassociateAddress",
        "ec2:GetManagedPrefixListEntries",
        "ec2:ModifyInstanceAttribute",
        "ec2:ModifyNetworkInterfaceAttribute",
        "ec2:ModifySubnetAttribute",
        "ec2:ModifyVpcAttribute",

```

```
"ec2:ReleaseAddress",
"ec2:RunInstances",
"ec2:StartInstances",
"ec2:StopInstances",
"ec2:TerminateInstances",
"iam:AddRoleToInstanceProfile",
"iam:AttachRolePolicy",
"iam:AttachUserPolicy",
"iam:CreateInstanceProfile",
"iam:CreatePolicy",
"iam:CreatePolicyVersion",
"iam:CreateRole",
"iam:CreateServiceLinkedRole",
"iam>DeleteInstanceProfile",
"iam>DeletePolicy",
"iam>DeletePolicyVersion",
"iam>DeleteRole",
"iam>DeleteRolePolicy",
"iam:DetachRolePolicy",
"iam:GetAccountSummary",
"iam:GetInstanceProfile",
"iam:GetPolicy",
"iam:GetPolicyVersion",
"iam:GetRole",
"iam>ListAttachedRolePolicies",
"iam>ListInstanceProfilesForRole",
"iam>ListPolicyVersions",
"iam>ListRolePolicies",
"iam:PassRole",
"iam:PutRolePolicy",
"iam:RemoveRoleFromInstanceProfile",
"iam:SimulatePrincipalPolicy",
"iam:UpdateAssumeRolePolicy",
"kms:Decrypt",
"kms:DescribeKey",
"kms:Encrypt",
"kms>ListAliases",
"kms>ListKeys",
"s3:CreateBucket",
"s3>DeleteBucket",
"s3>DeleteObject",
"s3>DeleteObjectVersion",
"s3:GetBucketLocation",
"s3:GetBucketObjectLockConfiguration",
"s3:GetBucketVersioning",
"s3:GetObject",
"s3:GetObjectRetention",
"s3:GetObjectVersion",
"s3>ListAllMyBuckets",
"s3>ListBucket",
"s3>ListBucketVersions",
"s3:PutBucketObjectLockConfiguration",
"s3:PutBucketVersioning",
"s3:PutObject",
```

```

        "s3:PutObjectRetention",
        "servicequotas:ListServiceQuotas",
        "ssm:DescribeInstanceInformation",
        "ssm:GetCommandInvocation",
        "ssm:GetParameter",
        "ssm:SendCommand",
        "sts:GetCallerIdentity"
    ],
    "Resource": "*"
}
]
}

```

IMPORTANT

Note that the following permissions are only required to remove created resources during appliance deployment in case of deployment failure or removal of the backup appliance from the backup infrastructure: `ec2:DeleteSubnet`, `ec2:DeleteSecurityGroup`, `ec2:DetachInternetGateway`, `ec2:DeleteInternetGateway`, `ec2:DeleteVpc`. If you have not added these permissions for security reasons, remove the resources manually using the AWS Management Console as described in [AWS Documentation](#).

- IAM roles whose permissions are used to perform data protection and disaster recovery operations with AWS resources.

When you deploy a new backup appliance, the *Default Backup Restore IAM role* is automatically created and added to the appliance. The *Default Backup Restore IAM role* is assigned all permissions required to perform data protection and disaster recovery operations in the same AWS account where the backup appliance resides. For more information on the *Default Backup Restore IAM role* permissions, see [Full List of IAM Permissions](#). However, you can create additional IAM roles with granular permissions and add them to the appliance as described in section [Managing IAM Roles](#).

- IAM users whose access keys are specified to access standard backup repositories where the image-level backups are stored must have permissions described in the [Using Amazon S3 Object Storage](#) section in the Veeam Backup & Replication User Guide if plan to [copy image-level backups](#) or to [restore guest OS files from image-level backups](#). To learn how to specify access keys of IAM users, see sections [Connecting to Existing Appliance](#) and [Creating New Repositories](#).
- IAM users whose one-time access keys are used to automatically grant missing permissions to IAM users and roles must have the following permissions:

```

"iam:AttachRolePolicy",
"iam:CreatePolicy"
"iam:GetAccountSummary",
"iam:GetPolicy",
"iam:GetPolicyVersion"
"iam:ListPolicyVersions",
"iam:ListAttachedUserPolicies"

```

Veeam Backup & Replication neither saves nor stores these one-time access keys in the configuration database.

Virtualization Servers and Hosts Service Account Permissions

If you plan to copy backups to on-premises backup repositories, to perform restore to VMware vSphere and Microsoft Hyper-V environments, or to perform other tasks related to virtualization servers and hosts, you must check whether the service account specified for these servers and hosts has the required permissions described in the [Veeam Backup & Replication User Guide for VMware vSphere](#) and [Veeam Backup & Replication User Guide for Microsoft Hyper-V](#), section *Using Virtualization Servers and Hosts*.

Microsoft Azure Account Permissions

An Azure AD application that you plan to use to restore EC2 instances to Microsoft Azure must have permissions described in the Veeam Backup & Replication User Guide, section [Permissions](#).

Google Cloud Service Account Permissions

A service account that you plan to use to restore EC2 instances to Google Cloud must have permissions described in the Veeam Backup & Replication User Guide, section [Google Compute Engine IAM User Permissions](#).

IAM Permissions

To perform data protection and disaster recovery operations, you must specify IAM roles whose permissions Veeam Backup for AWS will use to access AWS services and resources.

When you deploy Veeam Backup for AWS, the *Default Backup Restore* IAM role is automatically created and added to the backup appliance. This IAM role is assigned all permissions required to perform operations in the same AWS account where the backup appliance resides. However, you can manually create additional IAM roles with granular permissions to perform specific operations in this or in other AWS accounts [using the AWS Management Console](#), and then add them to Veeam Backup for AWS.

For more information on IAM roles in Veeam Backup for AWS, see [Managing IAM Roles](#).

Worker IAM Permissions

Depending on whether you plan to deploy worker instances in the backup account or in production accounts, IAM roles used for worker instance deployment and communication with the instances must have a specific set of permissions:

- [IAM role permissions required in the backup account.](#)
- [IAM role permissions required in production accounts.](#)

For more information on AWS accounts in which Veeam Backup for AWS deploys worker instances, see [Worker Deployment Options](#).

Worker Deployment Role Permissions in Backup Account

The worker deployment role (service IAM role) is used to deploy worker instances in the [backup account](#) to perform backup and restore operations, and to create IAM roles that are attached to the deployed instances and used by Veeam Backup for AWS to communicate with them. The IAM role is specified in the [worker instance settings](#) and must be granted the following permissions:

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "ebs:ListChangedBlocks",
        "ebs:ListSnapshotBlocks",
        "ec2:AttachVolume",
        "ec2:CopySnapshot",
        "ec2:CreateKeyPair",
        "ec2:CreateSnapshot",
        "ec2:CreateTags",
        "ec2:CreateVolume",
        "ec2>DeleteKeyPair",
        "ec2>DeleteSnapshot",
        "ec2>DeleteVolume",
        "ec2:DescribeAccountAttributes",
        "ec2:DescribeAvailabilityZones",
        "ec2:DescribeImages",
        "ec2:DescribeInstanceAttribute",
        "ec2:DescribeInstances",
        "ec2:DescribeKeyPairs",
        "ec2:DescribeRegions",
        "ec2:DescribeRouteTables",
        "ec2:DescribeSecurityGroups",
        "ec2:DescribeSnapshots",
        "ec2:DescribeSnapshotAttribute",
        "ec2:DescribeSubnets",
        "ec2:DescribeVolumes",
        "ec2:DescribeVpcEndpoints",
        "ec2:DescribeVpcs",
        "ec2:DetachVolume",
        "ec2:GetEbsDefaultKmsKeyId",
        "ec2:ModifyInstanceAttribute",
        "ec2:ModifySnapshotAttribute",
        "ec2:ModifyVolume",
        "ec2:RunInstances",
        "ec2:StartInstances",
        "ec2:StopInstances",
        "ec2:TerminateInstances",
        "iam:AddRoleToInstanceProfile",
        "iam:AttachRolePolicy",
        "iam:CreateInstanceProfile",
        "iam:CreateRole",
        "iam>DeleteInstanceProfile",
        "iam>DeleteRole",
        "iam>DeleteRolePolicy",
        "iam:DetachRolePolicy",
        "iam:GetContextKeysForPrincipalPolicy",
        "iam:GetInstanceProfile",
        "iam:GetRole",
        "iam:ListAccountAliases",
        "iam:ListAttachedRolePolicies",

```

```

        "iam:ListInstanceProfilesForRole",
        "iam:ListRolePolicies",
        "iam:PassRole",
        "iam:PutRolePolicy",
        "iam:RemoveRoleFromInstanceProfile",
        "iam:SimulatePrincipalPolicy",
        "kinesis:CreateStream",
        "kinesis>DeleteStream",
        "kinesis:DescribeStream",
        "kinesis:PutRecord",
        "kms:CreateGrant",
        "kms:DescribeKey",
        "kms:GenerateDataKeyWithoutPlaintext",
        "kms:GetKeyPolicy",
        "kms:ListAliases",
        "kms:ListKeys",
        "kms:ReEncryptFrom",
        "kms:ReEncryptTo",
        "servicequotas:ListServiceQuotas",
        "sqs:CreateQueue",
        "sqs>DeleteMessage",
        "sqs>DeleteQueue",
        "sqs:ListQueues",
        "sqs:ReceiveMessage",
        "sqs:SendMessage",
        "sqs:SetQueueAttributes",
        "ssm:GetCommandInvocation",
        "ssm:GetParameter",
        "ssm:SendCommand"
    ],
    "Resource": "*"
}
]
}

```

To learn how to create IAM roles and assign them the required permissions, see [Appendix A. Creating IAM Roles in AWS](#).

Worker Deployment Role Permissions in Production Accounts

IMPORTANT

When you instruct Veeam Backup for AWS to deploy worker instances in [production accounts](#) to perform RDS and EC2 backup and restore operations, as well as EFS indexing operations, Veeam Backup for AWS uses the permissions of IAM roles specified for backup and restore operations. That is why you must assign to these IAM roles additional permissions listed in sections [RDS Backup IAM Role Permissions](#), [EC2 Backup IAM Role Permissions](#), [EC2 Restore IAM Permissions](#), [RDS Database Restore IAM Permissions](#) and [EFS Backup IAM Role Permissions](#).

Veeam Backup for AWS uses IAM roles that are attached to worker instances deployed in production accounts, which are further used by Veeam Backup for AWS to communicate with these instances to perform the following operations:

- To create indexes of the backed up EFS file systems.
- To perform image-level backup and restore operations with PostgreSQL DB instances.
- To perform image-level backup, entire instance and volume-level restore operations with EC2 instances.

To perform these operations, IAM roles specified in the [EFS backup policy settings](#), [RDS backup policy settings](#), [EC2 backup policy settings](#), [EC2 entire instance restore](#), [EC2 volume-level restore](#) and [RDS database restore settings](#) must meet the following requirements:

- The IAM roles must be included at least in one instance profile. For more information on instance profiles, see [AWS Documentation](#).
- The backup appliance must be granted permissions to assume the IAM roles. For more information on the requirements for adding IAM roles, see [Before You Begin](#).
- The Amazon EC2 service must be granted permissions to assume the IAM roles.

To allow the Amazon EC2 service to assume an IAM role, configure trust relationships for the role and add the following statement to the trust policy:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "sts:AssumeRole",
      "Principal": {
        "Service": "ec2.amazonaws.com"
      }
    }
  ]
}
```

To learn how to modify role trust policies, see [AWS Documentation](#).

- The IAM roles must be granted the following permissions:

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": [
        "ec2messages:AcknowledgeMessage",
        "ec2messages>DeleteMessage",
        "ec2messages:FailMessage",
        "ec2messages:GetEndpoint",
        "ec2messages:GetMessages",
        "ec2messages:SendReply",
        "iam:GetContextKeysForPrincipalPolicy",
        "iam:GetRole",
        "iam:ListInstanceProfilesForRole",
        "iam:SimulatePrincipalPolicy",
        "sqs>DeleteMessage",
        "sqs:ListQueues",
        "sqs:ReceiveMessage",
        "sqs:SendMessage",
        "ssm:DescribeAssociation",
        "ssm:DescribeDocument",
        "ssm:GetDeployablePatchSnapshotForInstance",
        "ssm:GetDocument",
        "ssm:GetManifest",
        "ssm:GetParameter",
        "ssm:GetParameters",
        "ssm:ListAssociations",
        "ssm:ListInstanceAssociations",
        "ssm:PutComplianceItems",
        "ssm:PutConfigurePackageResult",
        "ssm:PutInventory",
        "ssm:UpdateAssociationStatus",
        "ssm:UpdateInstanceAssociationStatus",
        "ssm:UpdateInstanceInformation",
        "ssmmessages:CreateControlChannel",
        "ssmmessages:CreateDataChannel",
        "ssmmessages:OpenControlChannel",
        "ssmmessages:OpenDataChannel"
      ],
      "Resource": "*",
      "Effect": "Allow"
    }
  ]
}

```

To learn how to create IAM roles and assign them the required permissions, see [Appendix A. Creating IAM Roles in AWS](#).

NOTE

Since you do not choose an IAM role for file-level recovery operations, the role that you specify [when enabling worker deployment in production accounts](#) in the restore settings is also used by Veeam Backup for AWS to deploy worker instances. That is why this role must be assigned permissions listed in section [FLR Worker IAM Role Permissions](#).

Worker Configuration IAM Role Permissions

When creating a new [worker configuration](#), you specify an IAM role whose permissions will be used to list network settings available in AWS Regions of production AWS accounts. The specified IAM role must be granted the following permissions:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": [
        "ec2:DescribeAvailabilityZones",
        "ec2:DescribeVpcs",
        "ec2:DescribeRegions",
        "ec2:DescribeAccountAttributes",
        "ec2:DescribeSubnets",
        "ec2:DescribeSecurityGroups"
      ],
      "Resource": "*",
      "Effect": "Allow"
    }
  ]
}
```

FLR Worker IAM Role Permissions

To allow Veeam Backup for AWS to deploy worker instances in production accounts to perform EC2 file-level recovery operations, attach IAM roles to the instances and further to communicate with these instances, [IAM roles specified in the file-level recovery settings](#) must meet the following requirements:

1. The IAM roles must be included at least in one instance profile. For more information on instance profiles, see [AWS Documentation](#).
2. The backup appliance must be granted permissions to assume the IAM roles. For more information on the requirements for adding IAM roles, see [Before You Begin](#).

3. The Amazon EC2 service must be granted permissions to assume the IAM roles.

To allow the Amazon EC2 service to assume an IAM role, configure trust relationships for the role and add the following statement to the trust policy:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "sts:AssumeRole",
      "Principal": {
        "Service": "ec2.amazonaws.com"
      }
    }
  ]
}
```

To learn how to modify role trust policies, see [AWS Documentation](#).

4. The IAM roles must be granted the following permissions:

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "ec2:AttachNetworkInterface",
        "ec2:AttachVolume",
        "ec2:CopySnapshot",
        "ec2:CreateKeyPair",
        "ec2:CreateNetworkInterface",
        "ec2:CreateSnapshot",
        "ec2:CreateTags",
        "ec2:CreateVolume",
        "ec2>DeleteKeyPair",
        "ec2>DeleteNetworkInterface",
        "ec2>DeleteSnapshot",
        "ec2>DeleteVolume",
        "ec2:DescribeAccountAttributes",
        "ec2:DescribeAddresses",
        "ec2:DescribeAvailabilityZones",
        "ec2:DescribeInstanceAttribute",
        "ec2:DescribeInstances",
        "ec2:DescribeKeyPairs",
        "ec2:DescribeNetworkInterfaces",
        "ec2:DescribeRegions",
        "ec2:DescribeRouteTables",
        "ec2:DescribeSecurityGroups",
        "ec2:DescribeSnapshotAttribute",
        "ec2:DescribeSnapshots",
        "ec2:DescribeSubnets",
        "ec2:DescribeVolumeAttribute",
        "ec2:DescribeVolumes",
        "ec2:DescribeVpcs",
        "ec2:DetachVolume",
        "ec2:ModifyInstanceAttribute",
        "ec2:RunInstances",
        "ec2:StartInstances",
        "ec2:StopInstances",
        "ec2:TerminateInstances",
        "ec2messages:AcknowledgeMessage",
        "ec2messages>DeleteMessage",
        "ec2messages:FailMessage",
        "ec2messages:GetEndpoint",
        "ec2messages:GetMessages",
        "ec2messages:SendReply",
        "iam:GetContextKeysForPrincipalPolicy",
        "iam:GetRole",
        "iam:ListAccountAliases",
        "iam:ListInstanceProfilesForRole",
        "iam:PassRole",
        "iam:SimulatePrincipalPolicy",
        "kms:CreateGrant",
        "kms:DescribeKey",

```

```

        "kms:GetKeyPolicy",
        "kms:ListAliases",
        "kms:ListKeys",
        "kms:ReEncryptFrom",
        "kms:ReEncryptTo",
        "servicequotas:ListServiceQuotas",
        "sqs:CreateQueue",
        "sqs:DeleteMessage",
        "sqs:DeleteQueue",
        "sqs:ListQueues",
        "sqs:ReceiveMessage",
        "sqs:SetQueueAttributes",
        "sqs:SendMessage",
        "ssm:DescribeAssociation",
        "ssm:DescribeDocument",
        "ssm:GetCommandInvocation",
        "ssm:GetDeployablePatchSnapshotForInstance",
        "ssm:GetDocument",
        "ssm:GetManifest",
        "ssm:GetParameter",
        "ssm:GetParameters",
        "ssm:ListAssociations",
        "ssm:ListInstanceAssociations",
        "ssm:PutComplianceItems",
        "ssm:PutConfigurePackageResult",
        "ssm:PutInventory",
        "ssm:SendCommand",
        "ssm:UpdateAssociationStatus",
        "ssm:UpdateInstanceAssociationStatus",
        "ssm:UpdateInstanceInformation",
        "ssmmessages:CreateControlChannel",
        "ssmmessages:CreateDataChannel",
        "ssmmessages:OpenControlChannel",
        "ssmmessages:OpenDataChannel"
    ],
    "Resource": "*"
}
]
}

```

To learn how to create IAM roles and assign them the required permissions, see [Appendix A. Creating IAM Roles in AWS](#).

Repository IAM Permissions

To allow Veeam Backup for AWS to create backup repositories in Amazon S3 buckets and to access the repository when performing backup and restore operations, IAM roles specified in the [repository settings](#) must meet the following requirements:

1. The backup appliance must be granted permissions to assume the IAM roles. For more information on the requirements for adding IAM roles, see [Before You Begin](#).
2. The Amazon S3 Batch Operations service must be granted permissions to assume the IAM roles.

To allow the AWS service to assume an IAM role, configure trust relationships for the role and add the following statement to the trust policy:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "sts:AssumeRole",
      "Principal": {
        "Service": "batchoperations.s3.amazonaws.com"
      }
    }
  ]
}
```

To learn how to modify role trust policies, see [AWS Documentation](#).

2. The IAM roles must be granted the following permissions:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "ec2:DescribeRegions",
        "ec2:DescribeVpcEndpoints",
        "iam:GetContextKeysForPrincipalPolicy",
        "iam:SimulatePrincipalPolicy",
        "iam:GetRole",
        "iam:PassRole",
        "iam:ListAccountAliases",
        "kms:ListKeys",
        "kms:ListAliases",
        "kms:DescribeKey",
        "kms:Encrypt",
        "kms:Decrypt",
        "s3:ListAllMyBuckets",
        "s3:CreateJob",
        "s3:DescribeJob",
        "s3:PutObject",
        "s3:GetObject",
        "s3>DeleteObject",
        "s3:RestoreObject",
        "s3:GetObjectRetention",
        "s3:PutObjectRetention",
        "s3:GetObjectVersion",
        "s3>DeleteObjectVersion",
        "s3:ListBucket",
        "s3:ListBucketVersions",
        "s3:GetBucketLocation",
        "s3:GetBucketVersioning",
        "s3:GetBucketObjectLockConfiguration"
      ],
      "Resource": "*"
    }
  ]
}
```

IMPORTANT

If you plan to use KMS key encryption for backup repositories, consider the following:

- The key policy of an AWS KMS key that will be used to encrypt a repository must allow the IAM role specified in the repository settings access to the key.
- AWS managed keys cannot be used to encrypt repositories due to [AWS limitations](#).

To learn how to create IAM roles and assign them the required permissions, see [Appendix A. Creating IAM Roles in AWS](#).

Backup IAM Permissions

To allow Veeam Backup for AWS to perform backup of AWS resources, IAM roles specified for backup operations must be granted specific permissions that depend on the type of AWS resources being backed up:

- [EC2 Backup IAM Role Permissions](#)
- [RDS Backup IAM Role Permissions](#)
- [DynamoDB Backup IAM Role Permissions](#)
- [Redshift Backup IAM Role Permissions](#)
- [EFS Backup IAM Role Permissions](#)
- [FSx Backup IAM Role Permissions](#)
- [VPC Configuration Backup IAM Role Permissions](#)

EC2 Backup IAM Role Permissions

Veeam Backup for AWS uses *EC2 Backup* IAM roles to perform the following operations:

- To enumerate resources added to a backup policy.
- To create cloud-native snapshots of EC2 instances.
- To create snapshot replicas, and so on.

NOTE

The same scope of permissions is required for IAM roles used to perform backup operations automatically as described in section [Creating EC2 Backup Policies](#), and IAM roles used to perform backup operations manually as described in section [Creating EC2 Snapshots Manually](#).

To perform these operations, [IAM roles specified in the EC2 backup policy settings](#) must be granted the following permissions:

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "ebs:ListSnapshotBlocks",
        "ebs:ListChangedBlocks",
        "ec2:CreateSnapshots",
        "ec2:CreateSnapshot",
        "ec2:CopySnapshot",
        "ec2:CreateTags",
        "ec2:DescribeInstanceAttribute",
        "ec2:DeleteSnapshot",
        "ec2:DescribeSnapshotAttribute",
        "ec2:DescribeInstances",
        "ec2:DeleteTags",
        "ec2:DescribeRegions",
        "ec2:DescribeSnapshots",
        "ec2:DescribeVpcs",
        "ec2:DescribeImages",
        "ec2:ModifySnapshotAttribute",
        "ec2:DescribeAvailabilityZones",
        "ec2:DescribeVolumes",
        "ec2:DescribeInstanceTypes",
        "ec2:DescribeAddresses",
        "ec2:DescribeNetworkInterfaces",
        "ec2:DescribeSubnets",
        "ec2:DescribeConversionTasks",
        "ec2:DescribeVolumeAttribute",
        "ec2:DescribeTags",
        "ec2:GetEbsDefaultKmsKeyId",
        "events:DescribeRule",
        "events:RemoveTargets",
        "events:PutTargets",
        "events>DeleteRule",
        "events:ListTargetsByRule",
        "events:PutRule",
        "iam:GetContextKeysForPrincipalPolicy",
        "iam:GetInstanceProfile",
        "iam:SimulatePrincipalPolicy",
        "iam:ListAccountAliases",
        "iam:ListInstanceProfiles",
        "kms:ListKeys",
        "kms:ListAliases",
        "kms:GetKeyPolicy",
        "kms:ReEncryptTo",
        "kms:DescribeKey",
        "kms:ReEncryptFrom",
        "kms:CreateGrant",
        "servicequotas:ListServiceQuotas",
        "sqs>DeleteMessage",
        "sqs:ListQueues",
        "sqs:ReceiveMessage",

```

```
        "sqs:DeleteQueue",
        "sqs:CreateQueue",
        "sqs:SetQueueAttributes",
        "sns:ListSubscriptionsByTopic",
        "sns:DeleteTopic",
        "sns:CreateTopic",
        "sns:ListTopics",
        "sns:Unsubscribe",
        "sns:SetTopicAttributes",
        "sns:Subscribe",
        "ssm:SendCommand",
        "ssm:GetCommandInvocation",
        "ssm:DescribeInstanceInformation"
    ],
    "Resource": "*"
}
]
```

To learn how to create IAM roles and assign them the required permissions, see [Appendix A. Creating IAM Roles in AWS](#).

Permissions Required to Deploy Worker Instances in Production Account

[Applies only to IAM roles specified in the backup policy settings] If you plan to instruct Veeam Backup for AWS to deploy worker instances in [production accounts](#), [IAM roles specified in the backup policy settings](#) must be granted the following additional permissions:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "ec2:AttachVolume",
        "ec2:CreateKeyPair",
        "ec2:CreateVolume",
        "ec2>DeleteKeyPair",
        "ec2>DeleteVolume",
        "ec2:DescribeAccountAttributes",
        "ec2:DescribeKeyPairs",
        "ec2:DescribeSecurityGroups",
        "ec2:DetachVolume",
        "ec2:ModifyInstanceAttribute",
        "ec2:RunInstances",
        "ec2:TerminateInstances",
        "iam:GetRole",
        "iam:ListInstanceProfilesForRole",
        "iam:PassRole",
        "ssm:GetParameter",
        "sqs:SendMessage"
      ],
      "Resource": "*"
    }
  ]
}
```

RDS Backup IAM Role Permissions

Veeam Backup for AWS uses *RDS Backup* IAM roles to perform the following operations:

- To enumerate resources added to a backup policy.
- To create cloud-native snapshots of RDS resources.
- To create snapshot replicas, and so on.

NOTE

The same scope of permissions is required for IAM roles used to perform backup operations automatically as described in section [Creating RDS Backup Policies](#), and IAM roles used to perform backup operations manually as described in section [Creating RDS Snapshots Manually](#).

To perform backup operations, [IAM roles specified in the RDS backup policy settings](#) must be granted the following permissions:

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "ec2:DescribeRegions",
        "ec2:DescribeAvailabilityZones",
        "ec2:DescribeVpcs",
        "events:DescribeRule",
        "events:PutRule",
        "events:PutTargets",
        "events>DeleteRule",
        "events:RemoveTargets",
        "events:ListTargetsByRule",
        "iam:GetContextKeysForPrincipalPolicy",
        "iam:SimulatePrincipalPolicy",
        "iam:ListAccountAliases",
        "kms:DescribeKey",
        "kms:CreateGrant",
        "kms:GetKeyPolicy",
        "kms:ListKeys",
        "kms:ListAliases",
        "rds:AddTagsToResource",
        "rds:ListTagsForResource",
        "rds:DescribeDBSnapshots",
        "rds:CreateDBSnapshot",
        "rds:DescribeDBInstances",
        "rds>DeleteDBSnapshot",
        "rds:ModifyDBSnapshotAttribute",
        "rds:RemoveTagsFromResource",
        "rds:CopyDBSnapshot",
        "rds:DescribeDBClusters",
        "rds:CreateDBClusterSnapshot",
        "rds:DescribeDBClusterSnapshots",
        "rds>DeleteDBClusterSnapshot",
        "rds:CopyDBClusterSnapshot",
        "rds:ModifyDBClusterSnapshotAttribute",
        "rds:DescribeDBSubnetGroups",
        "sns:ListSubscriptionsByTopic",
        "sns>DeleteTopic",
        "sns:CreateTopic",
        "sns:ListTopics",
        "sns:Unsubscribe",
        "sns:SetTopicAttributes",
        "sns:Subscribe",
        "sqs>DeleteQueue",
        "sqs:CreateQueue",
        "sqs:SetQueueAttributes",
        "sqs>DeleteMessage",
        "sqs:ListQueues",
        "sqs:ReceiveMessage"
      ],
      "Resource": "*"
    }
  ]
}

```

```
}
  ]
}
```

To learn how to create IAM roles and assign them the required permissions, see [Appendix A. Creating IAM Roles in AWS](#).

Permissions Required to Create RDS Image-Level Backups

[Applies only to IAM roles specified in the backup policy settings] For Veeam Backup for AWS to be able to create [RDS image-level backups](#), [IAM roles specified in the backup policy settings](#) must be granted the following additional permissions to deploy worker instances in [production accounts](#):

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "ec2:AuthorizeSecurityGroupEgress",
        "ec2:AuthorizeSecurityGroupIngress",
        "ec2:TerminateInstances",
        "ec2:StartInstances",
        "ec2:RunInstances",
        "ec2:CreateKeyPair",
        "ec2:CreateSecurityGroup",
        "ec2:CreateTags",
        "ec2>DeleteKeyPair",
        "ec2>DeleteSecurityGroup",
        "ec2:DescribeAccountAttributes",
        "ec2:DescribeImages",
        "ec2:DescribeInstances",
        "ec2:DescribeInstanceAttribute",
        "ec2:DescribeInternetGateways",
        "ec2:DescribeKeyPairs",
        "ec2:DescribeRouteTables",
        "ec2:DescribeSecurityGroups",
        "ec2:DescribeSubnets",
        "ec2:DescribeVpcEndpoints",
        "ec2:ModifyInstanceAttribute",
        "ec2:RevokeSecurityGroupEgress",
        "ec2:RevokeSecurityGroupIngress",
        "iam:GetInstanceProfile",
        "iam:GetRole",
        "iam:ListInstanceProfilesForRole",
        "iam:PassRole",
        "rds:ModifyDBInstance",
        "servicequotas:listServiceQuotas",
        "ssm:GetCommandInvocation",
        "ssm:GetParameter",
        "ssm:SendCommand",
        "sqs:SendMessage"
      ],
      "Resource": "*"
    }
  ]
}
```


DynamoDB Backup IAM Role Permissions

Veeam Backup for AWS uses *DynamoDB Backup* IAM roles to perform the following operations:

- To enumerate resources added to a backup policy.
- To create backups of DynamoDB tables.
- To create backup copies, and so on.

To perform these operations, IAM roles specified in the DynamoDB backup policy settings must meet the following requirements:

1. The backup appliance must be granted permissions to assume the IAM roles. For more information on the requirements for adding IAM roles, see [Before You Begin](#).
2. The AWS Backup service must be granted permissions to assume the IAM roles.

To allow the AWS Backup service to assume an IAM role, configure trust relationships for the role and add the following statement to the trust policy:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "sts:AssumeRole",
      "Principal": {
        "Service": "backup.amazonaws.com"
      }
    }
  ]
}
```

To learn how to modify role trust policies, see [AWS Documentation](#).

3. The IAM roles must be granted the following permissions:
 - IAM roles specified in the [backup policy settings](#):

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "backup:CopyFromBackupVault",
        "backup:CopyIntoBackupVault",
        "backup>DeleteRecoveryPoint",
        "backup:DescribeBackupJob",
        "backup:DescribeCopyJob",
        "backup:DescribeRecoveryPoint",
        "backup:DescribeRegionSettings",
        "backup>ListBackupVaults",
        "backup>ListRecoveryPointsByBackupVault",
        "backup>ListTags",
        "backup:StartBackupJob",
        "backup:StartCopyJob",
        "backup:StopBackupJob",
        "backup:TagResource",
        "backup:UntagResource",
        "backup:UpdateRegionSettings",
        "dynamodb:DescribeContinuousBackups",
        "dynamodb:DescribeTable",
        "dynamodb:DescribeTimeToLive",
        "dynamodb>ListTables",
        "dynamodb>ListTagsOfResource",
        "dynamodb:StartAwsBackupJob",
        "ec2:DescribeRegions",
        "ec2:DescribeAvailabilityZones",
        "events>DeleteRule",
        "events:DescribeRule",
        "events>ListTargetsByRule",
        "events:PutRule",
        "events:PutTargets",
        "events:RemoveTargets",
        "iam:GetContextKeysForPrincipalPolicy",
        "iam:GetRole",
        "iam>ListAccountAliases",
        "iam:PassRole",
        "iam:SimulatePrincipalPolicy",
        "kms:Decrypt",
        "sns:CreateTopic",
        "sns>DeleteTopic",
        "sns>ListSubscriptionsByTopic",
        "sns>ListTopics",
        "sns:SetTopicAttributes",
        "sns:Subscribe",
        "sns:Unsubscribe",
        "sqs:CreateQueue",
        "sqs>DeleteMessage",
        "sqs>DeleteQueue",
        "sqs>ListQueues",
        "sqs:ReceiveMessage",

```

```

        "sqs:SetQueueAttributes"
    ],
    "Resource": "*"
}
]
}

```

- o IAM roles used to perform backup operations manually as described in section [Creating DynamoDB Backups Manually](#):

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "backup-storage:MountCapsule",
        "backup:CopyFromBackupVault",
        "backup:CopyIntoBackupVault",
        "backup:CreateBackupVault",
        "backup>DeleteBackupVault",
        "backup>DeleteRecoveryPoint",
        "backup:DescribeBackupJob",
        "backup:DescribeCopyJob",
        "backup:DescribeRecoveryPoint",
        "backup:DescribeRegionSettings",
        "backup:ListBackupVaults",
        "backup:ListTags",
        "backup:StartBackupJob",
        "backup:StartCopyJob",
        "backup:StopBackupJob",
        "backup:TagResource",
        "backup:UntagResource",
        "backup:UpdateRegionSettings",
        "dynamodb:DescribeContinuousBackups",
        "dynamodb:DescribeTable",
        "dynamodb:DescribeTimeToLive",
        "dynamodb:ListTagsOfResource",
        "dynamodb:StartAwsBackupJob",
        "ec2:DescribeRegions",
        "iam:GetContextKeysForPrincipalPolicy",
        "iam:GetRole",
        "iam:PassRole",
        "iam:SimulatePrincipalPolicy",
        "kms:Decrypt",
        "kms:DescribeKey"
      ],
      "Resource": "*"
    }
  ]
}

```

To learn how to create IAM roles and assign them the required permissions, see [Appendix A. Creating IAM Roles in AWS](#).

Redshift Backup IAM Role Permissions

Veeam Backup for AWS uses *Redshift Backup* IAM roles to perform the following operations:

- To enumerate resources added to a backup policy.
- To create backups of Redshift clusters.

To perform these operations, IAM roles specified in the Redshift backup policy settings must meet the following requirements:

1. The backup appliance must be granted permissions to assume the IAM roles. For more information on the requirements for adding IAM roles, see [Before You Begin](#).
2. The AWS Backup service must be granted permissions to assume the IAM roles.

To allow the AWS Backup service to assume an IAM role, configure trust relationships for the role and add the following statement to the trust policy:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "sts:AssumeRole",
      "Principal": {
        "Service": "backup.amazonaws.com"
      }
    }
  ]
}
```

To learn how to modify role trust policies, see [AWS Documentation](#).

3. The IAM roles must be granted the following permissions:
 - IAM roles specified in the [backup policy settings](#):

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "backup:DeleteRecoveryPoint",
        "backup:DescribeBackupJob",
        "backup:DescribeRegionSettings",
        "backup:DescribeRecoveryPoint",
        "backup:ListBackupVaults",
        "backup:ListRecoveryPointsByBackupVault",
        "backup:ListTags",
        "backup:StartBackupJob",
        "backup:StopBackupJob",
        "backup:UpdateRegionSettings",
        "ec2:DescribeRegions",
        "ec2:DescribeAvailabilityZones",
        "events:DeleteRule",
        "events:DescribeRule",
        "events:ListTargetsByRule",
        "events:PutRule",
        "events:PutTargets",
        "events:RemoveTargets",
        "iam:GetRole",
        "iam:GetContextKeysForPrincipalPolicy",
        "iam:PassRole",
        "iam:SimulatePrincipalPolicy",
        "iam:ListAccountAliases",
        "redshift:DescribeClusters",
        "redshift:CreateTags",
        "redshift:DescribeTags",
        "redshift>DeleteTags",
        "redshift>CreateClusterSnapshot",
        "redshift>DeleteClusterSnapshot",
        "redshift:DescribeClusterSnapshots",
        "sns:CreateTopic",
        "sns>DeleteTopic",
        "sns:ListSubscriptionsByTopic",
        "sns:ListTopics",
        "sns:SetTopicAttributes",
        "sns:Subscribe",
        "sns:Unsubscribe",
        "sqs:CreateQueue",
        "sqs>DeleteMessage",
        "sqs>DeleteQueue",
        "sqs:ListQueues",
        "sqs:ReceiveMessage",
        "sqs:SetQueueAttributes"
      ],
      "Resource": "*"
    }
  ]
}

```

- IAM roles used to perform backup operations manually as described in section [Creating Redshift Backups Manually](#):

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "backup:DeleteRecoveryPoint",
        "backup:DescribeBackupJob",
        "backup:DescribeRecoveryPoint",
        "backup:DescribeRegionSettings",
        "backup:ListBackupVaults",
        "backup:ListTags",
        "backup:StartBackupJob",
        "backup:StopBackupJob",
        "backup:UpdateRegionSettings",
        "ec2:DescribeAvailabilityZones",
        "ec2:DescribeRegions",
        "iam:GetContextKeysForPrincipalPolicy",
        "iam:GetRole",
        "iam:ListAccountAliases",
        "iam:SimulatePrincipalPolicy",
        "iam:PassRole",
        "redshift:CreateClusterSnapshot",
        "redshift:CreateTags",
        "redshift>DeleteTags",
        "redshift:DescribeClusters",
        "redshift:DescribeClusterSnapshots",
        "redshift:DescribeTags",
        "redshift>DeleteClusterSnapshot"
      ],
      "Resource": "*"
    }
  ]
}
```

To learn how to create IAM roles and assign them the required permissions, see [Appendix A. Creating IAM Roles in AWS](#).

EFS Backup IAM Role Permissions

Veeam Backup for AWS uses *EFS Backup* IAM roles to perform the following operations:

- To enumerate resources added to a backup policy.
- To create backups of EFS file systems.
- To create backup copies, and so on.

To perform these operations, IAM roles specified in the EFS backup policy settings must meet the following requirements:

1. The backup appliance must be granted permissions to assume the IAM roles. For more information on the requirements for adding IAM roles, see [Before You Begin](#).
2. The AWS Backup service must be granted permissions to assume the IAM roles.

To allow the AWS Backup service to assume an IAM role, configure trust relationships for the role and add the following statement to the trust policy.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "sts:AssumeRole",
      "Principal": {
        "Service": "backup.amazonaws.com"
      }
    }
  ]
}
```

To learn how to modify role trust policies, see [AWS Documentation](#).

3. The IAM roles must be granted the following permissions:
 - IAM roles specified in the [backup policy settings](#):

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "backup:CopyFromBackupVault",
        "backup:CopyIntoBackupVault",
        "backup>DeleteRecoveryPoint",
        "backup:DescribeBackupJob",
        "backup:DescribeCopyJob",
        "backup:DescribeRecoveryPoint",
        "backup>ListBackupVaults",
        "backup>ListRecoveryPointsByBackupVault",
        "backup>ListTags",
        "backup:StartBackupJob",
        "backup:StartCopyJob",
        "backup:StopBackupJob",
        "backup:TagResource",
        "backup:UntagResource",
        "ec2:CreateKeyPair",
        "ec2>DeleteKeyPair",
        "ec2:DescribeAvailabilityZones",
        "ec2:DescribeImages",
        "ec2:DescribeInstances",
        "ec2:DescribeInternetGateways",
        "ec2:DescribeKeyPairs",
        "ec2:DescribeNetworkInterfaceAttribute",
        "ec2:DescribeRegions",
        "ec2:DescribeRouteTables",
        "ec2:DescribeSecurityGroups",
        "ec2:DescribeSubnets",
        "ec2:DescribeVpcEndpoints",
        "ec2:DescribeVpcs",
        "elasticfilesystem:Backup",
        "elasticfilesystem:DescribeAccessPoints",
        "elasticfilesystem:DescribeBackupPolicy",
        "elasticfilesystem:DescribeFileSystemPolicy",
        "elasticfilesystem:DescribeFileSystems",
        "elasticfilesystem:DescribeLifecycleConfiguration",
        "elasticfilesystem:DescribeMountTargetSecurityGroups",
        "elasticfilesystem:DescribeMountTargets",
        "elasticfilesystem:DescribeTags",
        "elasticfilesystem:ListTagsForResource",
        "events>DeleteRule",
        "events:DescribeRule",
        "events>ListTargetsByRule",
        "events:PutRule",
        "events:PutTargets",
        "events:RemoveTargets",
        "iam:GetContextKeysForPrincipalPolicy",
        "iam:GetInstanceProfile",
        "iam:GetRole",
        "iam>ListAccountAliases",

```

```
        "iam:ListInstanceProfilesForRole",
        "iam:PassRole",
        "iam:SimulatePrincipalPolicy",
        "sns:CreateTopic",
        "sns:DeleteTopic",
        "sns:ListSubscriptionsByTopic",
        "sns:ListTopics",
        "sns:SetTopicAttributes",
        "sns:Subscribe",
        "sns:Unsubscribe",
        "sqs:CreateQueue",
        "sqs:DeleteMessage",
        "sqs:DeleteQueue",
        "sqs:ListQueues",
        "sqs:ReceiveMessage",
        "sqs:SendMessage",
        "sqs:SetQueueAttributes",
        "ssm:GetCommandInvocation",
        "ssm:GetParameter",
        "ssm:SendCommand"
    ],
    "Resource": "*"
}
}
```

- o IAM roles used to perform backup operations manually as described in section [Creating EFS Backups Manually](#):

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "backup-storage:MountCapsule",
        "backup:CopyFromBackupVault",
        "backup:CopyIntoBackupVault",
        "backup:CreateBackupVault",
        "backup>DeleteBackupVault",
        "backup>DeleteRecoveryPoint",
        "backup:DescribeBackupJob",
        "backup:DescribeCopyJob",
        "backup:DescribeRecoveryPoint",
        "backup>ListBackupVaults",
        "backup>ListTags",
        "backup:StartBackupJob",
        "backup:StartCopyJob",
        "backup:StopBackupJob",
        "backup:TagResource",
        "backup:UntagResource",
        "ec2:DescribeAvailabilityZones",
        "ec2:DescribeNetworkInterfaceAttribute",
        "ec2:DescribeRegions",
        "elasticfilesystem:Backup",
        "elasticfilesystem:DescribeAccessPoints",
        "elasticfilesystem:DescribeBackupPolicy",
        "elasticfilesystem:DescribeFileSystemPolicy",
        "elasticfilesystem:DescribeFileSystems",
        "elasticfilesystem:DescribeLifecycleConfiguration",
        "elasticfilesystem:DescribeMountTargetSecurityGroups",
        "elasticfilesystem:DescribeMountTargets",
        "elasticfilesystem:DescribeTags",
        "elasticfilesystem:ListTagsForResource",
        "iam:GetContextKeysForPrincipalPolicy",
        "iam:GetRole",
        "iam:PassRole",
        "iam:SimulatePrincipalPolicy",
        "kms:DescribeKey"
      ],
      "Resource": "*"
    }
  ]
}
```

To learn how to create IAM roles and assign them the required permissions, see [Appendix A. Creating IAM Roles in AWS](#).

Permissions Required to Create Indexes of EFS File Systems

[Applies only to IAM roles specified in the backup policy settings] If you plan to instruct Veeam Backup for AWS to perform [indexing of the processed file systems](#), [IAM roles specified in the backup policy settings](#) must be granted the following additional permissions to deploy worker instances in [production accounts](#):

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "ec2:TerminateInstances",
        "ec2:StartInstances"
      ],
      "Resource": "*",
      "Condition": {
        "StringEquals": {
          "ec2:ResourceTag/EfsIndexWorker": "EfsIndexWorker"
        }
      }
    },
    {
      "Effect": "Allow",
      "Action": "ec2:CreateTags",
      "Resource": "*",
      "Condition": {
        "StringEquals": {
          "ec2:CreateAction": "RunInstances",
          "aws:RequestTag/EfsIndexWorker": "EfsIndexWorker"
        }
      }
    },
    {
      "Effect": "Allow",
      "Action": [
        "ec2:RunInstances",
        "servicequotas:listservicequotas"
      ],
      "Resource": "*"
    }
  ]
}
```

FSx Backup IAM Role Permissions

Veeam Backup for AWS uses *FSx Backup* IAM roles to perform the following operations:

- To enumerate resources added to a backup policy.
- To create backups of FSx file systems.
- To create backup copies, and so on.

To perform these operations, IAM roles specified in the FSx backup policy settings must meet the following requirements:

1. The backup appliance must be granted permissions to assume the IAM roles. For more information on the requirements for adding IAM roles, see [Before You Begin](#).
2. The AWS Backup service must be granted permissions to assume the IAM roles.

To allow the AWS Backup service to assume an IAM role, configure trust relationships for the role and add the following statement to the trust policy.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "sts:AssumeRole",
      "Principal": {
        "Service": "backup.amazonaws.com"
      }
    }
  ]
}
```

To learn how to modify role trust policies, see [AWS Documentation](#).

3. The IAM roles must be granted the following permissions:
 - IAM roles specified in the [backup policy settings](#):


```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "backup:CopyFromBackupVault",
        "backup:CopyIntoBackupVault",
        "backup:DescribeCopyJob",
        "backup:DescribeRegionSettings",
        "backup>DeleteRecoveryPoint",
        "backup:DescribeBackupJob",
        "backup:DescribeRecoveryPoint",
        "backup:ListTags",
        "backup:ListBackupVaults",
        "backup:ListRecoveryPointsByBackupVault",
        "backup:StartBackupJob",
        "backup:StartCopyJob",
        "backup:UpdateRegionSettings",
        "events>DeleteRule",
        "events:DescribeRule",
        "events:ListTargetsByRule",
        "events:PutTargets",
        "events:PutRule",
        "events:RemoveTargets",
        "ec2:DescribeRegions",
        "ec2:DescribeNetworkInterfaces",
        "ec2:DescribeAvailabilityZones",
        "fsx:CopyBackup",
        "fsx:DescribeDataRepositoryAssociations",
        "fsx:DescribeFileSystems",
        "fsx:DescribeBackups",
        "fsx:ListTagsForResource",
        "fsx:CreateBackup",
        "fsx:TagResource",
        "fsx:UntagResource",
        "iam:GetContextKeysForPrincipalPolicy",
        "iam:GetRole",
        "iam:ListAccountAliases",
        "iam:PassRole",
        "iam:SimulatePrincipalPolicy",
        "kms:DescribeKey",
        "kms:CreateGrant",
        "sns:CreateTopic",
        "sns>DeleteTopic",
        "sns:ListSubscriptionsByTopic",
        "sns:Subscribe",
        "sns:Unsubscribe",
        "sqs:CreateQueue",
        "sqs>DeleteQueue",
        "sqs>DeleteMessage",
        "sqs:ListQueues",
        "sns:ListTopics",
        "sns:SetTopicAttributes",

```

```

        "sqs:SetQueueAttributes",
        "sqs:ReceiveMessage"
    ],
    "Resource": "*"
}
]
}

```

- o IAM roles used to perform backup operations manually as described in section [Creating FSx Backups Manually](#):

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "backup:DescribeRegionSettings",
        "backup:DeleteRecoveryPoint",
        "backup:DeleteBackupVault",
        "backup:DescribeBackupJob",
        "backup:DescribeCopyJob",
        "backup:DescribeRecoveryPoint",
        "backup:CopyIntoBackupVault",
        "backup:CopyFromBackupVault",
        "backup:CreateBackupVault",
        "backup:UpdateRegionSettings",
        "backup:ListBackupVaults",
        "backup:ListTags",
        "backup:StartBackupJob",
        "backup:StartCopyJob",
        "backup-storage:MountCapsule",
        "backup:TagResource",
        "ec2:DescribeNetworkInterfaces",
        "ec2:DescribeRegions",
        "fsx:CopyBackup",
        "fsx:CreateBackup",
        "fsx:DescribeBackups",
        "fsx:DescribeFileSystems",
        "fsx:DescribeDataRepositoryAssociations",
        "fsx:ListTagsForResource",
        "fsx:TagResource",
        "fsx:UntagResource",
        "iam:PassRole",
        "kms:CreateGrant",
        "kms:DescribeKey"
      ],
      "Resource": "*"
    }
  ]
}

```

To learn how to create IAM roles and assign them the required permissions, see [Appendix A. Creating IAM Roles in AWS](#).

VPC Configuration Backup IAM Role Permissions

Veeam Backup for AWS uses *VPC Configuration Backup* IAM roles to perform the following operations:

- To enumerate resources added to a backup policy.
- To create VPC configuration backups of AWS Regions.
- To create backup copies, and so on.

To perform these operations, [IAM roles specified in the VPC configuration backup policy settings](#) must be granted the following permissions:

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "ec2:DescribeAddresses",
        "ec2:DescribeClientVpnAuthorizationRules",
        "ec2:DescribeClientVpnEndpoints",
        "ec2:DescribeClientVpnRoutes",
        "ec2:DescribeClientVpnTargetNetworks",
        "ec2:DescribeCustomerGateways",
        "ec2:DescribeDhcpOptions",
        "ec2:DescribeEgressOnlyInternetGateways",
        "ec2:DescribeInstances",
        "ec2:DescribeInternetGateways",
        "ec2:DescribeManagedPrefixLists",
        "ec2:DescribeNatGateways",
        "ec2:DescribeNetworkAcls",
        "ec2:DescribeNetworkInterfaces",
        "ec2:DescribeRegions",
        "ec2:DescribeRouteTables",
        "ec2:DescribeSecurityGroups",
        "ec2:DescribeSubnets",
        "ec2:DescribeTransitGatewayAttachments",
        "ec2:DescribeTransitGatewayMulticastDomains",
        "ec2:DescribeTransitGatewayPeeringAttachments",
        "ec2:DescribeTransitGatewayRouteTables",
        "ec2:DescribeTransitGatewayVpcAttachments",
        "ec2:DescribeTransitGateways",
        "ec2:DescribeVpcAttribute",
        "ec2:DescribeVpcEndpointServiceConfigurations",
        "ec2:DescribeVpcEndpoints",
        "ec2:DescribeVpcPeeringConnections",
        "ec2:DescribeVpcs",
        "ec2:DescribeVpnConnections",
        "ec2:DescribeVpnGateways",
        "ec2:GetManagedPrefixListEntries",
        "ec2:GetTransitGatewayPrefixListReferences",
        "ec2:GetTransitGatewayRouteTableAssociations",
        "ec2:GetTransitGatewayRouteTablePropagations",
        "ec2:SearchTransitGatewayRoutes",
        "elasticloadbalancing:DescribeListeners",
        "elasticloadbalancing:DescribeLoadBalancers",
        "elasticloadbalancing:DescribeTags",
        "elasticloadbalancing:DescribeTargetGroups",
        "elasticloadbalancing:DescribeTargetHealth",
        "iam:GetContextKeysForPrincipalPolicy",
        "iam:ListAccountAliases",
        "iam:SimulatePrincipalPolicy",
        "ram:GetResourceShares",
        "ram:ListPrincipals",
        "ram:ListResourceSharePermissions",
        "ram:ListResources"
      ]
    }
  ]
}

```

```
    ],  
    "Resource": "*"  }  
  ]  
}
```

To learn how to create IAM roles and assign them the required permissions, see [Appendix A. Creating IAM Roles in AWS](#).

Restore IAM Permissions

To allow Veeam Backup for AWS to perform restore of AWS resources, IAM roles and IAM users whose one-time access keys are specified for restore operations must have specific permissions that depend on the type of AWS resources being restored:

- [EC2 Restore IAM Permissions](#)
- [RDS Instance Restore IAM Permissions](#)
- [RDS Database Restore IAM Permissions](#)
- [DynamoDB Restore IAM Permissions](#)
- [Redshift Restore IAM Permissions](#)
- [EFS Restore IAM Permissions](#)
- [FSx Restore IAM Permissions](#)
- [VPC Configuration Restore IAM Permissions](#)

EC2 Restore IAM Permissions

To perform EC2 restore operations, IAM roles and IAM users specified in the [entire EC2 instance](#) and [volume-level restore](#) settings must be granted the following permissions:


```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "ec2:AllocateAddress",
        "ec2:AssignPrivateIpAddresses",
        "ec2:AssociateAddress",
        "ec2:AttachNetworkInterface",
        "ec2:AttachVolume",
        "ec2:CopySnapshot",
        "ec2>CreateNetworkInterface",
        "ec2:CreateTags",
        "ec2>CreateVolume",
        "ec2>DeleteNetworkInterface",
        "ec2>DeleteSnapshot",
        "ec2>DeleteTags",
        "ec2>DeleteVolume",
        "ec2:DescribeAddresses",
        "ec2:DescribeAvailabilityZones",
        "ec2:DescribeConversionTasks",
        "ec2:DescribeImages",
        "ec2:DescribeInstanceAttribute",
        "ec2:DescribeInstanceState",
        "ec2:DescribeInstanceTypes",
        "ec2:DescribeInstances",
        "ec2:DescribeKeyPairs",
        "ec2:DescribeNetworkInterfaces",
        "ec2:DescribeRegions",
        "ec2:DescribeSecurityGroups",
        "ec2:DescribeSnapshotAttribute",
        "ec2:DescribeSnapshots",
        "ec2:DescribeSubnets",
        "ec2:DescribeVolumeAttribute",
        "ec2:DescribeVolumes",
        "ec2:DescribeVpcs",
        "ec2:DetachVolume",
        "ec2:DisassociateAddress",
        "ec2:GetEbsDefaultKmsKeyId",
        "ec2:ModifyInstanceAttribute",
        "ec2:ModifyNetworkInterfaceAttribute",
        "ec2:ModifyVolume",
        "ec2:RunInstances",
        "ec2:StartInstances",
        "ec2:StopInstances",
        "ec2:TerminateInstances",
        "ec2:AssociateIamInstanceProfile",
        "iam:GetInstanceProfile",
        "iam:GetContextKeysForPrincipalPolicy",
        "iam:ListAccountAliases",
        "iam:listInstanceProfiles",
        "iam:PassRole",
        "iam:SimulatePrincipalPolicy",

```

```

        "kms:CreateGrant",
        "kms:DescribeKey",
        "kms:GenerateDataKeyWithoutPlaintext",
        "kms:GetKeyPolicy",
        "kms:ListAliases",
        "kms:ListKeys",
        "kms:ReEncryptFrom",
        "kms:ReEncryptTo"
    ],
    "Resource": "*"
}
]
}

```

To learn how to create IAM roles and assign them the required permissions, see [Appendix A. Creating IAM Roles in AWS](#).

Permissions Required to Deploy Worker Instances in Production Account

If you plan to instruct Veeam Backup for AWS to deploy worker instances in production accounts to perform [entire EC2 instance](#) or [volume-level restore](#), the IAM roles specified in the [entire EC2 instance](#) and [volume-level restore](#) settings must be granted the following additional permissions:

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "ec2:CreateKeyPair",
        "ec2>DeleteKeyPair",
        "ec2:DescribeAccountAttributes",
        "iam:GetRole",
        "iam:ListInstanceProfilesForRole",
        "servicequotas:ListServiceQuotas",
        "sqs:CreateQueue",
        "sqs>DeleteMessage",
        "sqs>DeleteQueue",
        "sqs:ListQueues",
        "sqs:ReceiveMessage",
        "sqs:SendMessage",
        "sqs:SetQueueAttributes",
        "ssm:GetCommandInvocation",
        "ssm:GetParameter",
        "ssm:SendCommand"
      ],
      "Resource": "*"
    }
  ]
}

```

RDS Instance Restore IAM Permissions

To perform RDS instance restore operations, [IAM roles and IAM users specified in the restore settings](#) must be granted the following permissions:

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "ec2:DescribeAvailabilityZones",
        "ec2:DescribeDhcpOptions",
        "ec2:DescribeInternetGateways",
        "ec2:DescribeRegions",
        "ec2:DescribeSecurityGroups",
        "ec2:DescribeSubnets",
        "ec2:DescribeVpcAttribute",
        "ec2:DescribeVpcs",
        "iam:CreateServiceLinkedRole",
        "iam:GetContextKeysForPrincipalPolicy",
        "iam:ListAccountAliases",
        "iam:PassRole",
        "iam:SimulatePrincipalPolicy",
        "kms:CreateGrant",
        "kms:DescribeKey",
        "kms:GetKeyPolicy",
        "kms:ListAliases",
        "kms:ListKeys",
        "rds:AddTagsToResource",
        "rds:CopyDBClusterSnapshot",
        "rds:CopyDBSnapshot",
        "rds:CreateDbInstance",
        "rds:CreateTenantDatabase",
        "rds>DeleteDBClusterSnapshot",
        "rds>DeleteDBInstance",
        "rds>DeleteDBSnapshot",
        "rds>DeleteDbCluster",
        "rds:DescribeAccountAttributes",
        "rds:DescribeDBClusterSnapshots",
        "rds:DescribeDBClusters",
        "rds:DescribeDBEngineVersions",
        "rds:DescribeDBInstances",
        "rds:DescribeDBParameterGroups",
        "rds:DescribeDBSnapshots",
        "rds:DescribeDBSubnetGroups",
        "rds:DescribeDbClusterParameterGroups",
        "rds:DescribeDbClusterParameters",
        "rds:DescribeOptionGroups",
        "rds:DescribeOrderableDbInstanceOptions",
        "rds:ListTagsForResource",
        "rds:ModifyDBClusterSnapshotAttribute",
        "rds:ModifyDBInstance",
        "rds:ModifyDbCluster",
        "rds:RemoveTagsFromResource",
        "rds:RestoreDBInstanceFromDBSnapshot",
        "rds:RestoreDbClusterFromSnapshot",
        "servicequotas:ListServiceQuotas"
      ],
    }
  ],
}

```

```
        "Resource": "*"
    }
]
}
```

To learn how to create IAM roles and assign them the required permissions, see [Appendix A. Creating IAM Roles in AWS](#).

RDS Database Restore IAM Permissions

To perform RDS database restore operations, [IAM roles specified in the restore settings](#) must be granted the following permissions:

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "ec2:AuthorizeSecurityGroupEgress",
        "ec2:AuthorizeSecurityGroupIngress",
        "ec2:CreateKeyPair",
        "ec2:CreateSecurityGroup",
        "ec2:CreateTags",
        "ec2>DeleteKeyPair",
        "ec2>DeleteSecurityGroup",
        "ec2:DescribeAccountAttributes",
        "ec2:DescribeAvailabilityZones",
        "ec2:DescribeImages",
        "ec2:DescribeInstanceAttribute",
        "ec2:DescribeInstances",
        "ec2:DescribeInternetGateways",
        "ec2:DescribeKeyPairs",
        "ec2:DescribeRegions",
        "ec2:DescribeRouteTables",
        "ec2:DescribeSecurityGroups",
        "ec2:DescribeSubnets",
        "ec2:DescribeVpcs",
        "ec2:DescribeVpcEndpoints",
        "ec2:ModifyInstanceAttribute",
        "ec2:RevokeSecurityGroupEgress",
        "ec2:RevokeSecurityGroupIngress",
        "ec2:RunInstances",
        "ec2:StartInstances",
        "ec2:TerminateInstances",
        "iam:GetContextKeysForPrincipalPolicy",
        "iam:GetInstanceProfile",
        "iam:GetRole",
        "iam:ListInstanceProfilesForRole",
        "iam:PassRole",
        "iam:SimulatePrincipalPolicy",
        "rds:DescribeDBInstances",
        "rds:DescribeDBSubnetGroups",
        "rds:ModifyDBInstance",
        "servicequotas:listServiceQuotas",
        "sqs:CreateQueue",
        "sqs>DeleteMessage",
        "sqs>DeleteQueue",
        "sqs:ListQueues",
        "sqs:ReceiveMessage",
        "sqs:SendMessage",
        "sqs:SetQueueAttributes",
        "ssm:GetCommandInvocation",
        "ssm:GetParameter",
        "ssm:SendCommand"
      ],
      "Resource": "*"
    }
  ]
}

```

```
}
]
}
```

To learn how to create IAM roles and assign them the required permissions, see [Appendix A. Creating IAM Roles in AWS](#).

DynamoDB Restore IAM Permissions

To perform DynamoDB restore operations, IAM roles and IAM users specified in the [restore settings](#) must meet the following requirements:

1. The backup appliance must be granted permissions to assume the IAM roles. For more information on the requirements for adding IAM roles, see [Before You Begin](#).
2. The AWS Backup service must be granted permissions to assume the IAM roles.

To allow the AWS Backup service to assume an IAM role, configure trust relationships for the role and add the following statement to the trust policy:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "sts:AssumeRole",
      "Principal": {
        "Service": "backup.amazonaws.com"
      }
    }
  ]
}
```

To learn how to modify role trust policies, see [AWS Documentation](#).

2. The IAM roles must be granted the following permissions:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "backup-storage:MountCapsule",
        "backup:CopyFromBackupVault",
        "backup:CopyIntoBackupVault",
        "backup:CreateBackupVault",
        "backup>DeleteBackupVault",
        "backup>DeleteRecoveryPoint",
        "backup:DescribeCopyJob",
        "backup:DescribeRecoveryPoint",
        "backup:DescribeRestoreJob",
        "backup:ListBackupVaults",
        "backup:ListTags",
        "backup:StartCopyJob",
        "backup:StartRestoreJob",
        "backup:TagResource",
        "dynamodb>DeleteTable",
        "dynamodb:DescribeContinuousBackups",
        "dynamodb:DescribeTable",
        "dynamodb:DescribeTimeToLive",
        "dynamodb:ListTables",
        "dynamodb:RestoreTableFromAwsBackup",
        "dynamodb:TagResource",
        "dynamodb:UpdateContinuousBackups",
        "dynamodb:UpdateTable",
        "dynamodb:UpdateTimeToLive",
        "ec2:DescribeRegions",
        "iam:GetContextKeysForPrincipalPolicy",
        "iam:GetRole",
        "iam:ListAccountAliases",
        "iam:PassRole",
        "iam:SimulatePrincipalPolicy",
        "kms:CreateGrant",
        "kms:Decrypt",
        "kms:DescribeKey",
        "kms:ListAliases",
        "kms:ListKeys"
      ],
      "Resource": "*"
    }
  ]
}
```

To learn how to create IAM roles and assign them the required permissions, see [Appendix A. Creating IAM Roles in AWS](#).

Redshift Restore IAM Permissions

To perform Redshift restore operations, IAM roles and IAM users specified in the [restore settings](#) must meet the following requirements:

1. The backup appliance must be granted permissions to assume the IAM roles. For more information on the requirements for adding IAM roles, see [Before You Begin](#).
2. The AWS Backup service must be granted permissions to assume the IAM roles.

To allow the AWS Backup service to assume an IAM role, configure trust relationships for the role and add the following statement to the trust policy:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "sts:AssumeRole",
      "Principal": {
        "Service": "backup.amazonaws.com"
      }
    }
  ]
}
```

To learn how to modify role trust policies, see [AWS Documentation](#).

2. The IAM roles must be granted the following permissions:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "backup:DescribeRestoreJob",
        "backup:ListTags",
        "backup:StartRestoreJob",
        "ec2:DescribeInternetGateways",
        "ec2:DescribeRegions",
        "ec2:DescribeSecurityGroups",
        "ec2:DescribeSubnets",
        "ec2:DescribeVpcs",
        "iam:GetContextKeysForPrincipalPolicy",
        "iam:GetRole",
        "iam:ListAccountAliases",
        "iam:ListRoles",
        "iam:PassRole",
        "iam:SimulatePrincipalPolicy",
        "kms:CreateGrant",
        "kms:Decrypt",
        "kms:DescribeKey",
        "kms:GenerateDataKey",
        "kms:ListAliases",
        "kms:ListKeys",
        "redshift:CreateTags",
        "redshift>DeleteCluster",
        "redshift:DescribeClusterParameterGroups",
        "redshift:DescribeClusters",
        "redshift:DescribeClusterSnapshots",
        "redshift:DescribeClusterSubnetGroups",
        "redshift:DescribeNodeConfigurationOptions",
        "redshift:DescribeTags",
        "redshift:ModifyCluster",
        "redshift:RestoreFromClusterSnapshot",
        "secretsmanager:CreateSecret",
        "secretsmanager:DescribeSecret",
        "secretsmanager:TagResource"
      ],
      "Resource": "*"
    }
  ]
}
```

To learn how to create IAM roles and assign them the required permissions, see [Appendix A. Creating IAM Roles in AWS](#).

EFS Restore IAM Permissions

To perform EFS restore operations, IAM roles and IAM users specified in the [restore settings](#) must meet the following requirements:

1. The backup appliance must be granted permissions to assume the IAM roles. For more information on the requirements for adding IAM roles, see [Before You Begin](#).
2. The AWS Backup service must be granted permissions to assume the IAM roles.

To allow the AWS Backup service to assume an IAM role, configure trust relationships for the role and add the following statement to the trust policy.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "sts:AssumeRole",
      "Principal": {
        "Service": "backup.amazonaws.com"
      }
    }
  ]
}
```

To learn how to modify role trust policies, see [AWS Documentation](#).

2. The IAM roles must be granted the following permissions:

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "backup-storage:MountCapsule",
        "backup:CopyFromBackupVault",
        "backup:CopyIntoBackupVault",
        "backup:CreateBackupVault",
        "backup>DeleteBackupVault",
        "backup>DeleteRecoveryPoint",
        "backup:DescribeCopyJob",
        "backup:DescribeRecoveryPoint",
        "backup:DescribeRestoreJob",
        "backup:ListBackupVaults",
        "backup:ListTags",
        "backup:StartCopyJob",
        "backup:StartRestoreJob",
        "backup:TagResource",
        "ec2:DescribeAccountAttributes",
        "ec2:DescribeAvailabilityZones",
        "ec2:DescribeRegions",
        "ec2:DescribeSecurityGroups",
        "ec2:DescribeSubnets",
        "ec2:DescribeVpcs",
        "elasticfilesystem:CreateAccessPoint",
        "elasticfilesystem:CreateFileSystem",
        "elasticfilesystem:CreateMountTarget",
        "elasticfilesystem>DeleteAccessPoint",
        "elasticfilesystem>DeleteFileSystem",
        "elasticfilesystem>DeleteMountTarget",
        "elasticfilesystem:DescribeAccessPoints",
        "elasticfilesystem:DescribeFileSystemPolicy",
        "elasticfilesystem:DescribeFileSystems",
        "elasticfilesystem:DescribeLifecycleConfiguration",
        "elasticfilesystem:DescribeMountTargetSecurityGroups",
        "elasticfilesystem:DescribeMountTargets",
        "elasticfilesystem:DescribeReplicationConfigurations",
        "elasticfilesystem:PutBackupPolicy",
        "elasticfilesystem:PutFileSystemPolicy",
        "elasticfilesystem:PutLifecycleConfiguration",
        "elasticfilesystem:Restore",
        "elasticfilesystem:TagResource",
        "elasticfilesystem:UntagResource",
        "elasticfilesystem:UpdateFileSystem",
        "iam:GetContextKeysForPrincipalPolicy",
        "iam:GetRole",
        "iam:ListAccountAliases",
        "iam:PassRole",
        "iam:SimulatePrincipalPolicy",
        "kms:CreateGrant",
        "kms:DescribeKey",
        "kms:GenerateDataKeyWithoutPlaintext",

```

```
        "kms:ListAliases",
        "kms:ListKeys"
    ],
    "Resource": "*"
}
]
```

To learn how to create IAM roles and assign them the required permissions, see [Appendix A. Creating IAM Roles in AWS](#).

FSx Restore IAM Permissions

To perform FSx restore operations, IAM roles and IAM users specified in the [restore settings](#) must meet the following requirements:

1. The backup appliance must be granted permissions to assume the IAM roles. For more information on the requirements for adding IAM roles, see [Before You Begin](#).
2. The AWS Backup service must be granted permissions to assume the IAM roles.

To allow the AWS Backup service to assume an IAM role, configure trust relationships for the role and add the following statement to the trust policy.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "sts:AssumeRole",
      "Principal": {
        "Service": "backup.amazonaws.com"
      }
    }
  ]
}
```

To learn how to modify role trust policies, see [AWS Documentation](#).

2. The IAM roles must be granted the following permissions:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "ds:DescribeDirectories",
        "ec2:CreateTags",
        "ec2:DescribeRegions",
        "ec2:DescribeVpcs",
        "ec2:DescribeSecurityGroups",
        "ec2:DescribeSubnets",
        "ec2:DescribeRouteTables",
        "backup:CopyFromBackupVault",
        "backup:CopyIntoBackupVault",
        "backup:CreateBackupVault",
        "backup:DescribeCopyJob",
        "backup>DeleteRecoveryPoint",
        "backup>DeleteBackupVault",
        "backup:DescribeRestoreJob",
        "backup:DescribeRecoveryPoint",
        "backup:DescribeBackupVault",
        "backup:ListTags",
        "backup:StartCopyJob",
        "backup:StartRestoreJob",
        "backup:TagResource",
        "backup:ListBackupVaults",
        "backup-storage:MountCapsule",
        "fsx:CreateFileSystemFromBackup",
        "fsx:CopyBackup",
        "fsx:DescribeFileSystems",
        "fsx>DeleteFileSystem",
        "fsx:TagResource",
        "fsx:UntagResource",
        "fsx:DescribeBackups",
        "fsx:ListTagsForResource",
        "iam:GetRole",
        "iam:GetContextKeysForPrincipalPolicy",
        "iam:SimulatePrincipalPolicy",
        "iam:PassRole",
        "iam:ListAccountAliases",
        "kms:ListKeys",
        "kms:DescribeKey",
        "kms:CreateGrant",
        "kms:ListAliases",
        "servicequotas:ListServiceQuotas"
      ],
      "Resource": "*"
    }
  ]
}
```


To learn how to create IAM roles and assign them the required permissions, see [Appendix A. Creating IAM Roles in AWS](#).

VPC Configuration Restore IAM Permissions

To perform VPC configuration restore operations, [IAM roles and IAM users specified in the restore settings](#) must be granted the following permissions:

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "ec2:AcceptVpcEndpointConnections",
        "ec2:AllocateAddress",
        "ec2:AssociateAddress",
        "ec2:AssociateClientVpnTargetNetwork",
        "ec2:AssociateDhcpOptions",
        "ec2:AssociateRouteTable",
        "ec2:AssociateSubnetCidrBlock",
        "ec2:AssociateTransitGatewayMulticastDomain",
        "ec2:AssociateTransitGatewayRouteTable",
        "ec2:AssociateVpcCidrBlock",
        "ec2:AttachInternetGateway",
        "ec2:AttachVpnGateway",
        "ec2:AuthorizeClientVpnIngress",
        "ec2:AuthorizeSecurityGroupEgress",
        "ec2:AuthorizeSecurityGroupIngress",
        "ec2:CreateClientVpnEndpoint",
        "ec2:CreateClientVpnRoute",
        "ec2:CreateCustomerGateway",
        "ec2:CreateDefaultSubnet",
        "ec2:CreateDefaultVpc",
        "ec2:CreateDhcpOptions",
        "ec2:CreateEgressOnlyInternetGateway",
        "ec2:CreateInternetGateway",
        "ec2:CreateManagedPrefixList",
        "ec2:CreateNatGateway",
        "ec2:CreateNetworkAcl",
        "ec2:CreateNetworkAclEntry",
        "ec2:CreateNetworkInterface",
        "ec2:CreateRoute",
        "ec2:CreateRouteTable",
        "ec2:CreateSecurityGroup",
        "ec2:CreateSubnet",
        "ec2:CreateTags",
        "ec2:CreateTransitGateway",
        "ec2:CreateTransitGatewayMulticastDomain",
        "ec2:CreateTransitGatewayPeeringAttachment",
        "ec2:CreateTransitGatewayPrefixListReference",
        "ec2:CreateTransitGatewayRoute",
        "ec2:CreateTransitGatewayRouteTable",
        "ec2:CreateTransitGatewayVpcAttachment",
        "ec2:CreateVpc",
        "ec2:CreateVpcEndpoint",
        "ec2:CreateVpcEndpointServiceConfiguration",
        "ec2:CreateVpcPeeringConnection",
        "ec2:CreateVpnConnection",
        "ec2:CreateVpnGateway",
        "ec2>DeleteClientVpnEndpoint",
        "ec2>DeleteClientVpnRoute",

```

```
"ec2:DeleteCustomerGateway",
"ec2:DeleteDhcpOptions",
"ec2:DeleteEgressOnlyInternetGateway",
"ec2:DeleteInternetGateway",
"ec2:DeleteManagedPrefixList",
"ec2:DeleteNatGateway",
"ec2:DeleteNetworkAcl",
"ec2:DeleteNetworkAclEntry",
"ec2:DeleteNetworkInterface",
"ec2:DeleteRoute",
"ec2:DeleteRouteTable",
"ec2:DeleteSecurityGroup",
"ec2:DeleteSubnet",
"ec2:DeleteTransitGateway",
"ec2:DeleteTransitGatewayMulticastDomain",
"ec2:DeleteTransitGatewayPeeringAttachment",
"ec2:DeleteTransitGatewayPrefixListReference",
"ec2:DeleteTransitGatewayRoute",
"ec2:DeleteTransitGatewayRouteTable",
"ec2:DeleteTransitGatewayVpcAttachment",
"ec2:DeleteVpc",
"ec2:DeleteVpcEndpointServiceConfigurations",
"ec2:DeleteVpcEndpoints",
"ec2:DeleteVpcPeeringConnection",
"ec2:DeleteVpnConnection",
"ec2:DeleteVpnGateway",
"ec2:DescribeAddresses",
"ec2:DescribeAvailabilityZones",
"ec2:DescribeClientVpnAuthorizationRules",
"ec2:DescribeClientVpnEndpoints",
"ec2:DescribeClientVpnRoutes",
"ec2:DescribeClientVpnTargetNetworks",
"ec2:DescribeCustomerGateways",
"ec2:DescribeDhcpOptions",
"ec2:DescribeEgressOnlyInternetGateways",
"ec2:DescribeInstances",
"ec2:DescribeInternetGateways",
"ec2:DescribeManagedPrefixLists",
"ec2:DescribeNatGateways",
"ec2:DescribeNetworkAcls",
"ec2:DescribeNetworkInterfaces",
"ec2:DescribeRegions",
"ec2:DescribeRouteTables",
"ec2:DescribeSecurityGroups",
"ec2:DescribeSubnets",
"ec2:DescribeTransitGatewayAttachments",
"ec2:DescribeTransitGatewayMulticastDomains",
"ec2:DescribeTransitGatewayPeeringAttachments",
"ec2:DescribeTransitGatewayRouteTables",
"ec2:DescribeTransitGatewayVpcAttachments",
"ec2:DescribeTransitGateways",
"ec2:DescribeVpcAttribute",
"ec2:DescribeVpcEndpointServiceConfigurations",
"ec2:DescribeVpcEndpoints",
```

```
"ec2:DescribeVpcPeeringConnections",
"ec2:DescribeVpcs",
"ec2:DescribeVpnConnections",
"ec2:DescribeVpnGateways",
"ec2:DetachInternetGateway",
"ec2:DetachVpnGateway",
"ec2:DisableTransitGatewayRouteTablePropagation",
"ec2:DisableVgwRoutePropagation",
"ec2:DisassociateAddress",
"ec2:DisassociateClientVpnTargetNetwork",
"ec2:DisassociateRouteTable",
"ec2:DisassociateTransitGatewayMulticastDomain",
"ec2:DisassociateTransitGatewayRouteTable",
"ec2:EnableTransitGatewayRouteTablePropagation",
"ec2:EnableVgwRoutePropagation",
"ec2:GetManagedPrefixListEntries",
"ec2:GetTransitGatewayMulticastDomainAssociations",
"ec2:GetTransitGatewayPrefixListReferences",
"ec2:GetTransitGatewayRouteTableAssociations",
"ec2:GetTransitGatewayRouteTablePropagations",
"ec2:ModifyClientVpnEndpoint",
"ec2:ModifyManagedPrefixList",
"ec2:ModifyNetworkInterfaceAttribute",
"ec2:ModifySubnetAttribute",
"ec2:ModifyTransitGateway",
"ec2:ModifyTransitGatewayVpcAttachment",
"ec2:ModifyVpcAttribute",
"ec2:ModifyVpcEndpoint",
"ec2:ModifyVpcEndpointServiceConfiguration",
"ec2:ModifyVpcPeeringConnectionOptions",
"ec2:ModifyVpnConnection",
"ec2:RejectVpcEndpointConnections",
"ec2:ReleaseAddress",
"ec2:ReplaceNetworkAclAssociation",
"ec2:ReplaceRouteTableAssociation",
"ec2:RevokeClientVpnIngress",
"ec2:RevokeSecurityGroupEgress",
"ec2:RevokeSecurityGroupIngress",
"ec2:SearchTransitGatewayRoutes",
"elasticloadbalancing:AddTags",
"elasticloadbalancing:CreateListener",
"elasticloadbalancing:CreateLoadBalancer",
"elasticloadbalancing:CreateTargetGroup",
"elasticloadbalancing>DeleteListener",
"elasticloadbalancing>DeleteLoadBalancer",
"elasticloadbalancing>DeleteTargetGroup",
"elasticloadbalancing:DeregisterTargets",
"elasticloadbalancing:DescribeListeners",
"elasticloadbalancing:DescribeLoadBalancers",
"elasticloadbalancing:DescribeTags",
"elasticloadbalancing:DescribeTargetGroups",
"elasticloadbalancing:DescribeTargetHealth",
"elasticloadbalancing:ModifyTargetGroup",
"elasticloadbalancing:RegisterTargets",
```

```

        "elasticloadbalancing:RemoveTags",
        "elasticloadbalancing:SetSecurityGroups",
        "elasticloadbalancing:SetSubnets",
        "iam:GetContextKeysForPrincipalPolicy",
        "iam:ListAccountAliases",
        "iam:SimulatePrincipalPolicy",
        "lambda:ListFunctions",
        "ram:AssociateResourceShare",
        "ram:CreateResourceShare",
        "ram>DeleteResourceShare",
        "ram:DisassociateResourceShare",
        "ram:GetResourceShareAssociations",
        "ram:GetResourceShares",
        "ram:ListPrincipals",
        "ram:ListResourceSharePermissions",
        "ram:ListResources",
        "ram:TagResource",
        "ram:UntagResource",
        "s3:GetBucketLocation",
        "s3:ListAllMyBuckets",
        "s3:ListBucket",
        "s3:PutObject",
        "servicequotas:ListServiceQuotas"
    ],
    "Resource": "*"
}
]
}

```

To learn how to create IAM roles and assign them the required permissions, see [Appendix A. Creating IAM Roles in AWS](#).

Full List of IAM Permissions

If you want Veeam Backup for AWS to use a single IAM role to perform all restore and backup operations, you can use the *Default Backup Restore* IAM role created during Veeam Backup for AWS installation or a custom IAM role that must meet the following requirements:

1. The IAM role must be included at least in one instance profile. For more information on instance profiles, see [AWS Documentation](#).
2. The backup appliance must be granted permissions to assume the IAM roles. For more information on the requirements for adding IAM roles, see [Before You Begin](#).
3. The Amazon EC2, Amazon S3 Batch Operations and Amazon Backup services must be granted permissions to assume the IAM roles.

To allow an Amazon service to assume an IAM role, configure trust relationships for the role and add the following statement to the trust policy:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "Service": [
          "backup.amazonaws.com",
          "batchoperations.s3.amazonaws.com",
          "ec2.amazonaws.com"
        ]
      },
      "Action": "sts:AssumeRole"
    }
  ]
}
```

To learn how to modify role trust policies, see [AWS Documentation](#).

3. The IAM roles must be granted the following permissions:

IMPORTANT

Since the size of an IAM policy added to an IAM role cannot exceed 6.144 characters, it is recommended to create 2 IAM policies that will cover all the required permissions. For more information on IAM character limits, see [AWS Documentation](#).

- Permissions, part 1


```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "backup-storage:MountCapsule",
        "backup:CopyFromBackupVault",
        "backup:CopyIntoBackupVault",
        "backup:CreateBackupVault",
        "backup>DeleteBackupVault",
        "backup>DeleteRecoveryPoint",
        "backup:DescribeBackupVault",
        "backup:DescribeBackupJob",
        "backup:DescribeCopyJob",
        "backup:DescribeRecoveryPoint",
        "backup:DescribeRegionSettings",
        "backup:DescribeRestoreJob",
        "backup>ListBackupVaults",
        "backup>ListRecoveryPointsByBackupVault",
        "backup>ListTags",
        "backup:StartBackupJob",
        "backup:StartCopyJob",
        "backup:StartRestoreJob",
        "backup:StopBackupJob",
        "backup:TagResource",
        "backup:UntagResource",
        "backup:UpdateRegionSettings",
        "ds:DescribeDirectories",
        "dynamodb>DeleteTable",
        "dynamodb:DescribeContinuousBackups",
        "dynamodb:DescribeTable",
        "dynamodb:DescribeTimeToLive",
        "dynamodb>ListTables",
        "dynamodb>ListTagsOfResource",
        "dynamodb:RestoreTableFromAwsBackup",
        "dynamodb:StartAwsBackupJob",
        "dynamodb:TagResource",
        "dynamodb:UpdateContinuousBackups",
        "dynamodb:UpdateTable",
        "dynamodb:UpdateTimeToLive",
        "ebs>ListChangedBlocks",
        "ebs>ListSnapshotBlocks",
        "ec2:AcceptVpcEndpointConnections",
        "ec2:AllocateAddress",
        "ec2:AssignPrivateIpAddresses",
        "ec2:AssociateAddress",
        "ec2:AssociateClientVpnTargetNetwork",
        "ec2:AssociateDhcpOptions",
        "ec2:AssociateIamInstanceProfile",
        "ec2:AssociateRouteTable",
        "ec2:AssociateSubnetCidrBlock",
        "ec2:AssociateTransitGatewayMulticastDomain",
        "ec2:AssociateTransitGatewayRouteTable",

```

```
"ec2:AssociateVpcCidrBlock",
"ec2:AttachInternetGateway",
"ec2:AttachNetworkInterface",
"ec2:AttachVolume",
"ec2:AttachVpnGateway",
"ec2:AuthorizeClientVpnIngress",
"ec2:AuthorizeSecurityGroupEgress",
"ec2:AuthorizeSecurityGroupIngress",
"ec2:CopySnapshot",
"ec2:CreateClientVpnEndpoint",
"ec2:CreateClientVpnRoute",
"ec2:CreateCustomerGateway",
"ec2:CreateDefaultSubnet",
"ec2:CreateDefaultVpc",
"ec2:CreateDhcpOptions",
"ec2:CreateEgressOnlyInternetGateway",
"ec2:CreateInternetGateway",
"ec2:CreateKeyPair",
"ec2:CreateManagedPrefixList",
"ec2:CreateNatGateway",
"ec2:CreateNetworkAcl",
"ec2:CreateNetworkAclEntry",
"ec2:CreateNetworkInterface",
"ec2:CreateRoute",
"ec2:CreateRouteTable",
"ec2:CreateSecurityGroup",
"ec2:CreateSnapshot",
"ec2:CreateSnapshots",
"ec2:CreateSubnet",
"ec2:CreateTags",
"ec2:CreateTransitGateway",
"ec2:CreateTransitGatewayMulticastDomain",
"ec2:CreateTransitGatewayPeeringAttachment",
"ec2:CreateTransitGatewayPrefixListReference",
"ec2:CreateTransitGatewayRoute",
"ec2:CreateTransitGatewayRouteTable",
"ec2:CreateTransitGatewayVpcAttachment",
"ec2:CreateVolume",
"ec2:CreateVpc",
"ec2:CreateVpcEndpoint",
"ec2:CreateVpcEndpointServiceConfiguration",
"ec2:CreateVpcPeeringConnection",
"ec2:CreateVpnConnection",
"ec2:CreateVpnGateway",
"ec2>DeleteClientVpnEndpoint",
"ec2>DeleteClientVpnRoute",
"ec2>DeleteCustomerGateway",
"ec2>DeleteDhcpOptions",
"ec2>DeleteEgressOnlyInternetGateway",
"ec2>DeleteInternetGateway",
"ec2>DeleteKeyPair",
"ec2>DeleteManagedPrefixList",
"ec2>DeleteNatGateway",
"ec2>DeleteNetworkAcl",
```

```
"ec2:DeleteNetworkAclEntry",
"ec2:DeleteNetworkInterface",
"ec2:DeleteRoute",
"ec2:DeleteRouteTable",
"ec2:DeleteSecurityGroup",
"ec2:DeleteSnapshot",
"ec2:DeleteSubnet",
"ec2:DeleteTags",
"ec2:DeleteTransitGateway",
"ec2:DeleteTransitGatewayMulticastDomain",
"ec2:DeleteTransitGatewayPeeringAttachment",
"ec2:DeleteTransitGatewayPrefixListReference",
"ec2:DeleteTransitGatewayRoute",
"ec2:DeleteTransitGatewayRouteTable",
"ec2:DeleteTransitGatewayVpcAttachment",
"ec2:DeleteVolume",
"ec2:DeleteVpc",
"ec2:DeleteVpcEndpointServiceConfigurations",
"ec2:DeleteVpcEndpoints",
"ec2:DeleteVpcPeeringConnection",
"ec2:DeleteVpnConnection",
"ec2:DeleteVpnGateway",
"ec2:DescribeAccountAttributes",
"ec2:DescribeAddresses",
"ec2:DescribeAvailabilityZones",
"ec2:DescribeClientVpnAuthorizationRules",
"ec2:DescribeClientVpnEndpoints",
"ec2:DescribeClientVpnRoutes",
"ec2:DescribeClientVpnTargetNetworks",
"ec2:DescribeConversionTasks",
"ec2:DescribeCustomerGateways",
"ec2:DescribeDhcpOptions",
"ec2:DescribeEgressOnlyInternetGateways",
"ec2:DescribeImages",
"ec2:DescribeInstanceAttribute",
"ec2:DescribeInstanceStatus",
"ec2:DescribeInstanceTypes",
"ec2:DescribeInstances",
"ec2:DescribeInternetGateways",
"ec2:DescribeKeyPairs",
"ec2:DescribeManagedPrefixLists",
"ec2:DescribeNatGateways",
"ec2:DescribeNetworkAcls",
"ec2:DescribeNetworkInterfaceAttribute",
"ec2:DescribeNetworkInterfaces",
"ec2:DescribeRegions",
"ec2:DescribeRouteTables",
"ec2:DescribeSecurityGroups",
"ec2:DescribeSnapshotAttribute",
"ec2:DescribeSnapshots",
"ec2:DescribeSubnets",
"ec2:DescribeTags",
"ec2:DescribeTransitGatewayAttachments",
"ec2:DescribeTransitGatewayMulticastDomains",
```

```

        "ec2:DescribeTransitGatewayPeeringAttachments",
        "ec2:DescribeTransitGatewayRouteTables",
        "ec2:DescribeTransitGatewayVpcAttachments",
        "ec2:DescribeTransitGateways",
        "ec2:DescribeVolumeAttribute",
        "ec2:DescribeVolumes",
        "ec2:DescribeVpcAttribute",
        "ec2:DescribeVpcEndpointServiceConfigurations",
        "ec2:DescribeVpcEndpoints",
        "ec2:DescribeVpcPeeringConnections",
        "ec2:DescribeVpcs",
        "ec2:DescribeVpnConnections",
        "ec2:DescribeVpnGateways",
        "ec2:DetachInternetGateway",
        "ec2:DetachVolume",
        "ec2:DetachVpnGateway",
        "ec2:DisableTransitGatewayRouteTablePropagation",
        "ec2:DisableVgwRoutePropagation",
        "ec2:DisassociateAddress",
        "ec2:DisassociateClientVpnTargetNetwork",
        "ec2:DisassociateRouteTable",
        "ec2:DisassociateTransitGatewayMulticastDomain",
        "ec2:DisassociateTransitGatewayRouteTable",
        "ec2:EnableTransitGatewayRouteTablePropagation",
        "ec2:EnableVgwRoutePropagation",
        "ec2:GetEbsDefaultKmsKeyId",
        "ec2:GetManagedPrefixListEntries",
        "ec2:GetTransitGatewayMulticastDomainAssociations",
        "ec2:GetTransitGatewayPrefixListReferences",
        "ec2:GetTransitGatewayRouteTableAssociations",
        "ec2:GetTransitGatewayRouteTablePropagations"
    ],
    "Resource": "*"
}
]
}

```

› Permissions, part 2

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "ec2:ModifyClientVpnEndpoint",
        "ec2:ModifyInstanceAttribute",
        "ec2:ModifyManagedPrefixList",
        "ec2:ModifyNetworkInterfaceAttribute",
        "ec2:ModifySnapshotAttribute",
        "ec2:ModifySubnetAttribute",
        "ec2:ModifyTransitGateway",
        "ec2:ModifyTransitGatewayVpcAttachment",
        "ec2:ModifyVolume",
        "ec2:ModifyVpcAttribute",
        "ec2:ModifyVpcEndpoint",
        "ec2:ModifyVpcEndpointServiceConfiguration",
        "ec2:ModifyVpcPeeringConnectionOptions",
        "ec2:ModifyVpnConnection",
        "ec2:RejectVpcEndpointConnections",
        "ec2:ReleaseAddress",
        "ec2:ReplaceNetworkAclAssociation",
        "ec2:ReplaceRouteTableAssociation",
        "ec2:RevokeClientVpnIngress",
        "ec2:RevokeSecurityGroupEgress",
        "ec2:RevokeSecurityGroupIngress",
        "ec2:RunInstances",
        "ec2:SearchTransitGatewayRoutes",
        "ec2:StartInstances",
        "ec2:StopInstances",
        "ec2:TerminateInstances",
        "ec2messages:AcknowledgeMessage",
        "ec2messages>DeleteMessage",
        "ec2messages:FailMessage",
        "ec2messages:GetEndpoint",
        "ec2messages:GetMessages",
        "ec2messages:SendReply",
        "elasticfilesystem:Backup",
        "elasticfilesystem:CreateAccessPoint",
        "elasticfilesystem:CreateFileSystem",
        "elasticfilesystem:CreateMountTarget",
        "elasticfilesystem>DeleteAccessPoint",
        "elasticfilesystem>DeleteFileSystem",
        "elasticfilesystem>DeleteMountTarget",
        "elasticfilesystem:DescribeAccessPoints",
        "elasticfilesystem:DescribeBackupPolicy",
        "elasticfilesystem:DescribeFileSystemPolicy",
        "elasticfilesystem:DescribeFileSystems",
        "elasticfilesystem:DescribeLifecycleConfiguration",
        "elasticfilesystem:DescribeMountTargetSecurityGroups",
        "elasticfilesystem:DescribeMountTargets",
        "elasticfilesystem:DescribeReplicationConfigurations",
        "elasticfilesystem:DescribeTags",

```

```
"elasticfilesystem:ListTagsForResource",
"elasticfilesystem:PutBackupPolicy",
"elasticfilesystem:PutFileSystemPolicy",
"elasticfilesystem:PutLifecycleConfiguration",
"elasticfilesystem:Restore",
"elasticfilesystem:TagResource",
"elasticfilesystem:UntagResource",
"elasticfilesystem:UpdateFileSystem",
"elasticloadbalancing:AddTags",
"elasticloadbalancing:CreateListener",
"elasticloadbalancing:CreateLoadBalancer",
"elasticloadbalancing:CreateTargetGroup",
"elasticloadbalancing>DeleteListener",
"elasticloadbalancing>DeleteLoadBalancer",
"elasticloadbalancing>DeleteTargetGroup",
"elasticloadbalancing:DeregisterTargets",
"elasticloadbalancing:DescribeListeners",
"elasticloadbalancing:DescribeLoadBalancers",
"elasticloadbalancing:DescribeTags",
"elasticloadbalancing:DescribeTargetGroups",
"elasticloadbalancing:DescribeTargetHealth",
"elasticloadbalancing:ModifyTargetGroup",
"elasticloadbalancing:RegisterTargets",
"elasticloadbalancing:RemoveTags",
"elasticloadbalancing:SetSecurityGroups",
"elasticloadbalancing:SetSubnets",
"events>DeleteRule",
"events:DescribeRule",
"events:ListTargetsByRule",
"events:PutRule",
"events:PutTargets",
"events:RemoveTargets",
"fsx:CopyBackup",
"fsx:CreateBackup",
"fsx:CreateFileSystemFromBackup",
"fsx:DescribeBackups",
"fsx:DescribeDataRepositoryAssociations",
"fsx:DescribeFileSystems",
"fsx>DeleteFileSystem",
"fsx:ListTagsForResource",
"fsx:TagResource",
"fsx:UntagResource",
"iam:AddRoleToInstanceProfile",
"iam:AttachRolePolicy",
"iam:CreateInstanceProfile",
"iam:CreateRole",
"iam:CreateServiceLinkedRole",
"iam>DeleteInstanceProfile",
"iam>DeleteRole",
"iam>DeleteRolePolicy",
"iam:DetachRolePolicy",
"iam:GetContextKeysForPrincipalPolicy",
"iam:GetInstanceProfile",
"iam:GetRole",
```

```
"iam:ListAccountAliases",
"iam:ListAttachedRolePolicies",
"iam:ListInstanceProfiles",
"iam:ListInstanceProfilesForRole",
"iam:ListRoles",
"iam:ListRolePolicies",
"iam:PassRole",
"iam:PutRolePolicy",
"iam:RemoveRoleFromInstanceProfile",
"iam:SimulatePrincipalPolicy",
"kinesis:CreateStream",
"kinesis>DeleteStream",
"kinesis:DescribeStream",
"kinesis:PutRecord",
"kms:CreateGrant",
"kms:Decrypt",
"kms:DescribeKey",
"kms:Encrypt",
"kms:GenerateDataKey",
"kms:GenerateDataKeyWithoutPlaintext",
"kms:GetKeyPolicy",
"kms:ListAliases",
"kms:ListKeys",
"kms:ReEncryptFrom",
"kms:ReEncryptTo",
"lambda:ListFunctions",
"ram:AssociateResourceShare",
"ram:CreateResourceShare",
"ram>DeleteResourceShare",
"ram:DisassociateResourceShare",
"ram:GetResourceShareAssociations",
"ram:GetResourceShares",
"ram:ListPrincipals",
"ram:ListResourceSharePermissions",
"ram:ListResources",
"ram:TagResource",
"ram:UntagResource",
"redshift:CreateClusterSnapshot",
"redshift:CreateTags",
"redshift>DeleteCluster",
"redshift>DeleteClusterSnapshot",
"redshift>DeleteTags",
"redshift:DescribeClusterParameterGroups",
"redshift:DescribeClusters",
"redshift:DescribeClusterSnapshots",
"redshift:DescribeClusterSubnetGroups",
"redshift:DescribeNodeConfigurationOptions",
"redshift:DescribeTags",
"redshift:ModifyCluster",
"redshift:RestoreFromClusterSnapshot",
"rds:AddTagsToResource",
"rds:CopyDBClusterSnapshot",
"rds:CopyDBSnapshot",
"rds>CreateDBClusterSnapshot",
```



```

        "rds:CreateDBSnapshot",
        "rds:CreateDbInstance",
        "rds:CreateTenantDatabase",
        "rds>DeleteDBClusterSnapshot",
        "rds>DeleteDBInstance",
        "rds>DeleteDBSnapshot",
        "rds>DeleteDbCluster",
        "rds:DescribeAccountAttributes",
        "rds:DescribeDBClusterSnapshots",
        "rds:DescribeDBClusters",
        "rds:DescribeDBEngineVersions",
        "rds:DescribeDBInstances",
        "rds:DescribeDBParameterGroups",
        "rds:DescribeDBSnapshots",
        "rds:DescribeDBSubnetGroups",
        "rds:DescribeDbClusterParameterGroups",
        "rds:DescribeDbClusterParameters",
        "rds:DescribeOptionGroups",
        "rds:DescribeOrderableDbInstanceOptions",
        "rds:ListTagsForResource",
        "rds:ModifyDBClusterSnapshotAttribute",
        "rds:ModifyDBInstance",
        "rds:ModifyDBSnapshotAttribute",
        "rds:ModifyDbCluster",
        "rds:RemoveTagsFromResource",
        "rds:RestoreDBInstanceFromDBSnapshot",
        "rds:RestoreDbClusterFromSnapshot",
        "s3:CreateJob",
        "s3>DeleteObject",
        "s3>DeleteObjectVersion",
        "s3:DescribeJob",
        "s3:GetBucketLocation",
        "s3:GetBucketObjectLockConfiguration",
        "s3:GetBucketVersioning",
        "s3:GetObject",
        "s3:GetObjectRetention",
        "s3:GetObjectVersion",
        "s3:ListAllMyBuckets",
        "s3:ListBucket",
        "s3:ListBucketVersions",
        "s3:PutObject",
        "s3:PutObjectRetention",
        "s3:RestoreObject"
    ],
    "Resource": "*"
}
]
}

```

> Permissions, part 3

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "secretsmanager:CreateSecret",
        "secretsmanager:DescribeSecret",
        "secretsmanager:TagResource",
        "servicequotas:ListServiceQuotas",
        "sns:CreateTopic",
        "sns:DeleteTopic",
        "sns:ListSubscriptionsByTopic",
        "sns:ListTopics",
        "sns:SetTopicAttributes",
        "sns:Subscribe",
        "sns:Unsubscribe",
        "sqs:CreateQueue",
        "sqs:DeleteMessage",
        "sqs:DeleteQueue",
        "sqs:ListQueues",
        "sqs:ReceiveMessage",
        "sqs:SendMessage",
        "sqs:SetQueueAttributes",
        "ssm:DescribeAssociation",
        "ssm:DescribeDocument",
        "ssm:DescribeInstanceInformation",
        "ssm:GetCommandInvocation",
        "ssm:GetDeployablePatchSnapshotForInstance",
        "ssm:GetDocument",
        "ssm:GetManifest",
        "ssm:GetParameter",
        "ssm:GetParameters",
        "ssm:ListAssociations",
        "ssm:ListInstanceAssociations",
        "ssm:PutComplianceItems",
        "ssm:PutConfigurePackageResult",
        "ssm:PutInventory",
        "ssm:SendCommand",
        "ssm:UpdateAssociationStatus",
        "ssm:UpdateInstanceAssociationStatus",
        "ssm:UpdateInstanceInformation",
        "ssmmessages:CreateControlChannel",
        "ssmmessages:CreateDataChannel",
        "ssmmessages:OpenControlChannel",
        "ssmmessages:OpenDataChannel"
      ],
      "Resource": "*"
    }
  ]
}
```

To learn how to create IAM roles and assign them the required permissions, see [Appendix A. Creating IAM Roles in AWS](#).

IAM Permissions Changelog

This section describes the latest changes in IAM permissions required for Veeam Backup for AWS to perform operations.

When you update Veeam Backup for AWS version 7.0 to version 8.0, consider that additional permissions must be granted to the following IAM roles:

- For Veeam Backup for AWS to be able to deploy worker instances in the backup account when performing backup and restore operations, the [worker deployment role](#) (service IAM role) must be granted the following additional permission:

```
"sqs:SetQueueAttributes"
```

- For Veeam Backup for AWS to be able to back up EC2 instances with instance profiles attached, the IAM role specified in the [backup policy settings](#) must be granted the following additional permission:

```
"iam:GetInstanceProfile"
```

- For Veeam Backup for AWS to be able to restore EC2 instances with instance profiles attached and deploy worker instances in production accounts, the IAM role specified in the [restore settings](#) must be granted the following additional permissions:

```
"ec2:AssociateIamInstanceProfile",  
"iam:GetInstanceProfile",  
"iam:ListInstanceProfiles",  
"servicequotas:ListServiceQuotas",  
"sqs:SetQueueAttributes"
```

- For Veeam Backup for AWS to be able to perform EC2 file-level recovery operations and deploy worker instances in production accounts, the IAM role specified in the [restore settings](#) must be granted the following additional permissions:

```
"servicequotas:ListServiceQuotas",  
"sqs:SetQueueAttributes"
```

- For Veeam Backup for AWS to be able to perform RDS backup, the IAM role specified in the [backup policy settings](#) must be granted the following additional permissions:

```
"ec2:DescribeVpcs",  
"servicequotas:ListServiceQuotas"
```

- For Veeam Backup for AWS to be able to perform RDS database restore operations, the IAM role specified in the [restore settings](#) must be granted the following additional permissions:

```
"ec2:DescribeVpcs",  
"servicequotas:ListServiceQuotas",  
"sqs:SetQueueAttributes"
```

- For Veeam Backup for AWS to be able to back up EFS file systems, the IAM role specified in the [backup policy settings](#) must be granted the following additional permission:

```
"sqs:SendMessage",  
"servicequotas:ListServiceQuotas"
```

Also, the IAM role that is attached to worker instances and specified in the [backup policy settings](#) must be granted the following additional permissions:

```
"sqs:DeleteMessage",  
"sqs:ListQueues",  
"sqs:ReceiveMessage",  
"sqs:SendMessage"
```

The following permission has been removed from the [Worker Deployment Role Permissions in Production Accounts](#) list:

```
"sts:AssumeRole"
```

- For Veeam Backup for AWS to be able to perform DynamoDB backup, the IAM role specified in the [backup policy settings](#) must be granted the following additional permission:

```
"ec2:DescribeAvailabilityZones"
```

You can [update the roles manually using the AWS Management Console](#) or instruct Veeam Backup for AWS to do it, as described in section [Updating IAM Roles](#).

IMPORTANT

Veeam Backup for AWS version 8.0 comes with 2 major features – the ability to protect Redshift clusters and FSx file systems. For the list of permissions required to perform backup and restore operations with these workloads, see sections [Redshift Backup IAM Role Permissions](#), [FSx Backup IAM Role Permissions](#), [Redshift Restore IAM Permissions](#) and [FSx Restore IAM Permissions](#).

Considerations and Limitations

When you plan to deploy and configure Veeam Backup for AWS, keep in mind the following limitations and considerations.

Deployment

When deploying backup appliances, consider the following:

- Veeam Backup for AWS is available only in AWS Global and AWS GovCloud (US) regions.
- You can deploy Veeam Backup for AWS within a single Availability Zone only.
- To ensure successful deployment and installation of Veeam Backup for AWS, customers are encouraged to make sure they are operating within AWS service quotas. For more information, see [AWS Documentation](#).

Licensing

If the license file is not installed, Veeam Backup for AWS will operate in the *Free* edition allowing you to protect up to 10 instances free of charge.

Hardware

The minimum recommended EC2 instance type for the backup appliance is *t3.medium*. For the list of all existing instance types, see [AWS Documentation](#).

Software

To access Veeam Backup for AWS, use Microsoft Edge (latest version), Mozilla Firefox (latest version) or Google Chrome (latest version). Internet Explorer is not supported.

Security Certificates

Veeam Backup for AWS supports certificates only in the .PFX and .P12 formats.

Backup Repositories

When managing backup repositories, consider the following:

- Amazon S3 buckets with S3 Object Lock and S3 Versioning enabled can be used only for creating backup repositories with enabled immutability settings.
- Amazon S3 buckets using server-side encryption with AWS KMS keys (SSE-KMS) are not supported.
- Veeam Backup for AWS allows you to store backups only in the S3 Standard, S3 Glacier Flexible Retrieval and S3 Glacier Deep Archive storage classes. The S3 Standard-IA and S3 One Zone-IA storage classes are not supported.
- You cannot change Amazon S3 buckets, folders and storage classes for backup repositories already added to Veeam Backup for AWS.

- Veeam Backup for AWS does not support changing immutability settings for existing repositories specified in backup policies. However, you can change immutability settings for the newly created bucket in the AWS Management Console after creating the bucket, and then specify this bucket as the target location for image-level backups when creating a new repository with immutability enabled.
- If you enable S3 Object Lock for a bucket that is already used as a target location for image-level backups, Veeam Backup for AWS will not be able to continue creating backups and storing them in an existing backup repository.
- When you add a backup repository of the S3 Glacier Flexible Retrieval or S3 Glacier Deep Archive storage class, Veeam Backup for AWS does not create any S3 Glacier vaults in your AWS environment – it assigns the selected storage class to backups stored in the repository. That is why these backups remain in Amazon S3 and cannot be accessed directly through the Amazon S3 Glacier service.
- If you plan to use [AWS Key Management Service \(KMS\) keys](#) to encrypt backup repositories, note that only symmetric KMS keys are supported.

If you use a KMS key to encrypt a repository, do not disable or delete this key. Otherwise, Veeam Backup for AWS will not be able to encrypt and decrypt data stored in the repository.

- After you create a repository with encryption enabled, you will not be able to disable encryption for this repository. However, you will still be able to change the encryption settings as described in section [Editing Backup Repository Settings](#).
- A backup repository must not be managed by multiple backup appliances simultaneously. Retention sessions running on different backup appliances may corrupt backups stored in the repository, which may result in unpredictable data loss.
- Even though an Amazon S3 bucket is no longer used as a backup repository, Veeam Backup for AWS preserves all backup files previously stored in the repository and keeps these files in Amazon S3.

If you no longer need the backed-up data, either delete it as described in sections [Removing EC2 Backups and Snapshots](#), [Removing RDS Backups and Snapshots](#) and [Removing VPC Configuration Backups](#) before you remove the repository from Veeam Backup for AWS, or [use the AWS Management Console](#) to delete the data if the repository has already been removed.

EC2 Backup

When protecting EC2 instances, consider the following:

- Veeam Backup for AWS protects only EC2 instances that run in VPCs. EC2-Classic instances are not supported. For more information, see [this Veeam KB article](#).
- When Veeam Backup for AWS backs up EC2 instances with IPv6 addresses assigned, it does not save the addresses. That is why when you restore these instances, IP addresses are assigned according to the settings specified in AWS for the subnet to which the restored instances will be connected.
- Veeam Backup for AWS may fail to create image-level backups of EC2 instances with [product codes](#) if the AMIs that were used to deploy the instances do not support the type of worker instances deployed for the backup operation. To work around the issue, modify the worker profile to choose another instance type, as described in section [Managing Worker Profiles](#).
- [Applies only to image-level backups and file-level recovery from cloud-native snapshots] Veeam Backup for AWS does not support backup and restore of EC2 instances with [product codes](#) that have vendor restrictions preventing root EBS volumes from being attached to worker instances as secondary volumes. To learn how Veeam Backup for AWS performs EC2 backup, see [Protecting EC2 Instances](#).
- When configuring policy scheduling, consider that Veeam Backup for AWS runs retention sessions at 4:00 AM by default, according to the time zone set on the backup appliance. If you schedule backup policies to execute at 4:00 AM, the backup policies and retention tasks will be queued.

RDS Backup

When protecting RDS resources, consider the following:

- Veeam Backup for AWS does not support backup of Oracle DB instances with multi-tenant architecture, as well as backup of PostgreSQL DB clusters with Multi-AZ DB cluster deployment and IBM Db2 DB instances.
- Veeam Backup for AWS allows you to create image-level backups of PostgreSQL DB instances only. For the list of supported PostgreSQL versions, see [Protecting RDS Resources](#).
- For Veeam Backup for AWS to be able to create image-level backups of PostgreSQL DB instances, make sure that [security groups associated with worker instances](#) allow outbound HTTPS traffic from the worker instances through port 443 to download a certificate bundle for establishing SSL/TLS connections. For more information on certificate bundles for AWS Regions, see [AWS Documentation](#).
- When configuring policy scheduling, consider that Veeam Backup for AWS runs retention sessions at 4:00 AM by default, according to the time zone set on the backup appliance. If you schedule backup policies to execute at 4:00 AM, the backup policies and retention tasks will be queued.

Redshift Backup

When protecting Redshift clusters, consider the following:

- Veeam Backup for AWS does not support backup of Amazon Redshift Serverless.
- Veeam Backup for AWS does not support storing backups in [logically air gapped vaults](#).
- Veeam Backup for AWS supports backup of Redshift clusters only to the same AWS accounts to which the source clusters belong and the same AWS Region where the source cluster resides.
- For Veeam Backup for AWS to be able to back up Redshift clusters, you must enable the Opt-in service for the Redshift resource type in the AWS Backup settings. Otherwise, Veeam Backup for AWS will automatically enable the service for each AWS Region specified in the backup policy settings in your AWS account while performing backup operations. For more information on considerations for using AWS Backup with Amazon Redshift, see [AWS Documentation](#).
- When configuring policy scheduling, consider that Veeam Backup for AWS runs retention sessions at 4:00 AM by default, according to the time zone set on the backup appliance. If you schedule backup policies to execute at 4:00 AM, the backup policies and retention tasks will be queued.

DynamoDB Backup

When protecting DynamoDB tables, consider the following:

- Veeam Backup for AWS does not support storing backups in [logically air gapped vaults](#).
- Veeam Backup for AWS supports backup of DynamoDB tables only to the same AWS accounts where the source tables belong.
- Veeam Backup for AWS uses the [AWS Backup](#) service to create DynamoDB backups and backup copies. The [DynamoDB backup](#) service is not supported.
- For Veeam Backup for AWS to be able to back up DynamoDB tables, you must configure the AWS Backup settings to enable both the Opt-in service and the advanced features for Amazon DynamoDB backups. Otherwise, Veeam Backup for AWS will automatically enable these settings for each AWS Region specified in the backup policy settings in your AWS account while performing backup operations. For more information on advanced DynamoDB backup, see [AWS Documentation](#).

- When configuring policy scheduling, consider that Veeam Backup for AWS runs retention sessions at 4:00 AM by default, according to the time zone set on the backup appliance. If you schedule backup policies to execute at 4:00 AM, the backup policies and retention tasks will be queued.

EFS Backup

When protecting EFS file systems, consider the following:

- Veeam Backup for AWS does not support storing backups in [logically air gapped vaults](#).
- Veeam Backup for AWS supports backup of EFS file systems only to the same AWS accounts where the source file systems belong.
- Indexing of the backed up EFS file systems is not supported in the *Free* edition of Veeam Backup for AWS. For more information on license editions, see [Licensing](#).
- When configuring policy scheduling, consider that Veeam Backup for AWS runs retention sessions at 4:00 AM by default, according to the time zone set on the backup appliance. If you schedule backup policies to execute at 4:00 AM, the backup policies and retention tasks will be queued.

FSx Backup

When protecting FSx file systems, consider the following:

- Veeam Backup for AWS does not support storing backups in [logically air gapped vaults](#).
- Veeam Backup for AWS supports backup of FSx file systems only to the same AWS accounts where the source file systems belong.
- Veeam Backup for AWS does not support backup of Amazon FSx for NetApp ONTAP file systems.
- Veeam Backup for AWS does not support backup of Amazon FSx for Lustre file systems with the [Scratch deployment type](#).
- Veeam Backup for AWS does not support backup of Amazon FSx for Lustre with the [data repository association](#).
- Veeam Backup for AWS uses the [AWS Backup](#) service to create FSx backups and backup copies. The AWS Backup service does not support creating backup copies of FSx backups stored in [Opt-in Regions](#).
- For Veeam Backup for AWS to be able to back up FSx file systems, you must enable the Opt-in service for the FSx resource type in the AWS Backup settings. Otherwise, Veeam Backup for AWS will automatically enable the service for each AWS Region specified in the backup policy settings in your AWS account while performing backup operations.
- When configuring policy scheduling, consider that Veeam Backup for AWS runs retention sessions at 4:00 AM by default, according to the time zone set on the backup appliance. If you schedule backup policies to execute at 4:00 AM, the backup policies and retention tasks will be queued.

VPC Backup

When protecting VPC configurations, consider the following:

- Veeam Backup for AWS does not support backup of the following VPC configuration components: VPC Traffic Mirroring, AWS Network Firewall, Route 53 Resolver DNS Firewall, AWS Verified Access, VPC Flow Logs, carrier gateways, customer IP pools, transit gateway policy tables, and core networks in route tables.

- When configuring policy scheduling, consider that Veeam Backup for AWS runs retention sessions at 4:00 AM by default, according to the time zone set on the backup appliance. If you schedule backup policies to execute at 4:00 AM, the backup policies and retention tasks will be queued.

EC2 Restore

When restoring EC2 instances, consider the following:

- When restoring multiple EC2 instances that have the same EBS volume attached, Veeam Backup for AWS restores one volume per each instance and enables the Multi-Attach option for every restored volume. For more information on Amazon EBS Multi-Attach, see [AWS Documentation](#).
- Restore of files and folders is supported only for the following file systems: FAT, FAT32, NTFS, ext2, ext3, ext4, XFS, Btrfs.

For EC2 instances running Microsoft Windows OSes, Veeam Backup for AWS supports file-level recovery only for basic volumes.

- Restore of EC2 instances to the original location cannot be performed, if the source instances with termination protection and stop protection enabled still exist in AWS.

RDS Restore

When restoring Aurora DB clusters to a new location, Veeam Backup for AWS creates only primary DB instances in the restored clusters. Additional writer DB instances (for Aurora multi-master clusters) or Aurora Replicas (for Aurora DB clusters with single-master replication) must be added manually in the AWS Management Console after the restore operation completes. To learn how to add DB instances to Amazon Aurora DB clusters, see [AWS Documentation](#).

Redshift Restore

When restoring Redshift clusters, consider the following:

- Veeam Backup for AWS supports restore of Amazon Redshift clusters only to the same AWS accounts to which the source clusters belong and the same AWS Region where the source cluster resides.
- Veeam Backup for AWS does not support restore of Amazon Redshift clusters with the Multi-AZ deployment.
- Veeam Backup for AWS does not support restore of Amazon Redshift Serverless.

DynamoDB Restore

When restoring DynamoDB tables, consider the following:

- The [AWS Backup](#) service does not support copying DynamoDB backups stored in a cold storage tier to another AWS Region. These means that you will only be able to use these backups to restore tables to the same AWS Region in which the backups reside after being transitioned from a warm storage tier.
- Veeam Backup for AWS supports restore of DynamoDB tables only to the same AWS account to which the source tables belong.
- You can change the Time to Live (TTL) setting for DynamoDB tables only an hour after the restore operation completes.

EFS Restore

Veeam Backup for AWS supports restore of EFS file systems only to the same AWS account to which the source file systems belong.

FSx Restore

When restoring FSx file systems, consider the following:

- Veeam Backup for AWS supports restore of FSx file systems only to the same AWS accounts to which the source file systems belong.
- Veeam Backup for AWS does not support restore to a new location or with different settings for Amazon FSx for OpenZFS file systems with the Single-AZ 1 (HA) and Single-AZ 2 (HA) deployment types.
- Veeam Backup for AWS does not support restore of file system properties described in section [Protecting FSx File Systems](#).

VPC Restore

When restoring VPC configurations, consider the following:

- Restore of entire VPC configurations to a new location is not supported for the following VPC configuration items: Client VPN endpoints, customer gateways and load balancer listeners that use authentication certificates and specific components of route tables (core networks, routes to AWS Outpost local gateways, network interfaces, instances and carrier gateways).
- Restore of specific VPC configuration items to a new location is not supported.

Sizing and Scalability Guidelines

This section is intended for professionals who search for a best practice answer to sizing-related issues, and assumes you have already read the whole Veeam Backup for AWS User Guide.

Be aware that a best practice is not the only answer available. It will fit in the majority of cases but can also be totally wrong under different circumstances. Make sure you understand the implications of the recommended practices, or request assistance. If in doubt, reach out to Veeam professionals on [Veeam R&D Forums](#).

IMPORTANT

You must also consider the [AWS service quotas](#) associated with your AWS accounts, as well as the performance of [AWS instances of specific types](#). Some of the options may look good; however, make sure to take into account disk size, speed and burst credits.

Backup Appliance

You can choose the type of the EC2 instance running Veeam Backup for AWS during the deployment, or change it later as the environment grows.

General Recommendations

The following recommendations and examples apply to the latest Veeam Backup for AWS builds.

Instance Type	Recommended Maximum Number of Protected EC2 Instances
T3.medium (default - 2 vCPU, 4 GB RAM)	1,000
T3.2xlarge (medium - 8 vCPU, 32 GB RAM)	5,000
C5.9xlarge (large - 36 vCPU, 72 GB RAM)	10,000

When defining the instance type and amount of RAM required for proper functioning of the backup appliance, take into account the following:

- The average amount of RAM consumed in the idle state (approximately 1.5 GB).
- 5% of the total backup appliance RAM required for the Veeam Backup for AWS Web UI and REST API service.
- The maximum amount of RAM consumed by running backup policies. For more information, see [Backup Policies](#).

The RAM consumed by a backup policy depends on the data protection scenario.

Backup Policy Configuration	RAM Utilization (Default)	Additional RAM (per Workload)
EC2 Backup Policy		
Snapshots only	105 MB	1 MB
Snapshots and snapshot replicas	130 MB	1 MB
Snapshots and backups	170 MB	3 MB
Snapshots, snapshot replicas and backups	170 MB	3 MB
RDS Backup Policy		
Snapshots only	105 MB	1 MB

Backup Policy Configuration	RAM Utilization (Default)	Additional RAM (per Workload)
Snapshots and snapshot replicas	110 MB	1 MB
Snapshots and backups	180 MB	3 MB
Snapshots, snapshot replicas and backups	185 MB	3 MB
EFS Backup Policy		
Snapshots only	105 MB	3 MB
Snapshots and backup copies	120 MB	3 MB
Snapshots and indexing	165 MB	3 MB
Snapshots, backup copies and indexing	165 MB	3 MB
FSx Backup Policy		
Backups only	110 MB	3 MB
Backups and backup copies	125 MB	3 MB
DynamoDB Backup Policy		
Snapshots only	110 MB	3 MB
Snapshots and backup copies	125 MB	3 MB
Redshift Backup Policy		
Backups only	110 MB	3 MB

Note that these values are provided for demonstration purposes only. For production environments, it is recommended that you allocate an additional margin of 20% RAM.

RAM Sizing Examples

Consider the following example. You configure a number of backup policies to protect your workloads by regularly creating snapshots, snapshot replicas and backups. In this case, we advise to allocate minimum 150 MB per 1 policy.

The amount of RAM utilized by policies running on a backup appliance (Utilized RAM) depends on the total amount of RAM allocated to the backup appliance, the number of configured backup policies and the number of workloads protected by one policy. However, consider that the actual amount of RAM available for policy execution (Free RAM) will also be affected by the OS and Veeam services operation.

Total RAM	Number of Backup Policies	Workloads per Backup Policy	Utilized RAM ¹	Free RAM ²
4 GB	5	50	$(150 + (50 * 3)) * 5$ = ~ 1.5 GB	4 GB - 1.5 GB - 4 GB * 0.05 = 2.3 GB
8 GB	20	50	$(150 + (50 * 3)) * 20$ = ~ 6 GB	8 GB - 1.5 GB - 8 GB * 0.05 = 6.1 GB
16 GB	50	30	$(150 + (30 * 3)) * 50$ = ~ 12 GB	16 GB - 1.5 GB - 16 GB * 0.05 = 13.7 GB
32 GB	75	75	$(150 + (75 * 3)) * 75$ = ~ 28.2 GB	32 GB - 1.5 GB - 32 GB * 0.05 = 28.9 GB
72 GB	250	25	$(150 + (25 * 3)) * 250$ = ~ 56.25 GB	72 GB - 1.5 GB - 72 GB * 0.05 = 66.9 GB

¹The table shows the maximum amount of RAM utilization when all backup policies run at the same time.

²Additional RAM required for any other software must be calculated separately.

CPU Sizing Examples

Amount of vCPUs	Number of Snapshots Taken Simultaneously
EC2 CPU	
2 vCPU	< 300
4 vCPU	< 600
8 vCPU	< 1,600

Amount of vCPUs	Number of Snapshots Taken Simultaneously
16 vCPU	> 1,600
RDS CPU	
2 vCPU	< 300
4 vCPU	< 800
8 vCPU	< 1,600
16 vCPU	> 1,600
EFS CPU	
> 4 vCPU	> 25
DynamoDB CPU	
> 4 vCPU	> 100

*The examples apply only to workloads protected by snapshots and snapshot replicas, as the backup process is performed by worker instances.

Configuration Restore Recommendations

The following is recommended for large-scale deployments.

- The root EBS volume attached to the backup appliance must have at least twice as much free space as the size of the configuration backup file. If the backup file grows too large, you can increase the volume size as described in [AWS Documentation](#). Alternatively, open a [support case](#) to remove the unnecessary data from the configuration database.
- The EBS volume where Veeam Backup for AWS stores its configuration database must have at least twice as much free space as the size of the database. During configuration restore, Veeam Backup for AWS first creates the restored database and then deletes the original one.

Logging Recommendations

You can modify the following logging options in the configuration file `/etc/veeam/awsbackup/config.ini`:

```
[LogOptions]
LogLevel = "Normal"
LogsArchivesMaxCount = 100
LogsArchivesMaxSizeMb = 1000
WorkerLogsLifeTime = "36500:00:00:00"
WorkerLogsMaxArchivesCount = 2147483647
WorkerLogsMaxSizeMb = 2147483647
```

If the log files grow too large, you can remove them from the `/mnt/vcb-storage/logs` or `/var/log/veeam` folder, or open a [support case](#) to remove the unnecessary data.

Veeam Backup & Replication Integration

When you connect a backup appliance to the backup infrastructure, its backup policies, cloud-native snapshots, image-level backups, backup repositories and sessions are imported into the Veeam Backup & Replication database.

Time Consumption

When you connect an existing backup appliance to the backup infrastructure, the integration process includes the following steps:

- Retrieving data from the backup appliance.
- Saving the retrieved data to the Veeam Backup & Replication database.

Protected Workloads	Snapshots	Backups	Backup Policy Sessions	Workload Processing Sessions	Time Consumption
1,000	100,000	100,000	8,000	400,000	about 2 hours*
2,000	200,000	200,000	16,000	800,000	about 3 hours*
4,000	400,000	400,000	32,000	1,600,000	about 5 hours*

*The results were obtained when testing the backup appliance (c5.4xlarge, 16-core CPU, 32 GB RAM), the Veeam Backup & Replication server (PGSQL, 16-core CPU, 16 GB RAM) and Veeam Backup & Replication server (MSSQL, 16-core CPU, 16 GB RAM) and are approximate.

NOTE

The process of synchronizing data between the backup appliance and Veeam Backup & Replication database runs every 2 minutes after you add the backup appliance to the backup infrastructure. Creating new backup policies, updating policy settings, running backup and restore sessions may also trigger the synchronization process.

Object Storage

Veeam Backup for AWS compresses all backed-up data when saving it to object storage. The compression rate depends on the type and structure of source data and usually varies from 50% to 60%. This means that the compressed data typically consumes 50% less storage space than the source data.

Parameter	Value
Average size of backed-up data	40%-50% of source data
Compression rate	50%-60%

Object Sizes

Depending on whether you choose to keep backed-up data in short-term or long-term storage, Veeam Backup for AWS saves different objects to S3 buckets.

Object Type	S3 Storage Type	Block Size
Backup	S3 Standard	1 MB (compressed to ~512 KB)
Archive	S3 Glacier and S3 Glacier Deep Archive	512 MB
Metadata	S3 Standard	4 KB (per 1 GB of source data)

Amazon S3 Bucket Limits

You can send 3,500 PUT/COPY/POST/DELETE and 5,500 GET/HEAD requests per second per prefix in an Amazon S3 bucket. Veeam Backup for AWS has built-in mechanisms to assure you do not exceed the recommended maximums. While you could use 1 bucket to store all your data, it is recommended to use multiple buckets and S3 Glacier for cost-effective long-term archiving. For more information on Amazon S3 pricing, see [AWS Documentation](#).

It is also recommended to use dedicated IAM roles for backup repositories, as described in section [Repository IAM Permissions](#).

Cost Estimation

Veeam Backup for AWS comes with a built-in cost calculator that allows you to estimate your AWS expenses. It uses publicly available AWS price lists, so it may not reflect your exact cost in case of custom pricing or an enterprise agreement. Full details can be found at the cost estimation step of the **Add Policy** wizard.

Backup Policies

Since one backup policy can be used to protect multiple workloads at the same time, it is recommended that you limit the number of processed workloads to simplify the backup schedule and to optimize the backup performance.

General Recommendations

This section provides best practices for the maximum number of workloads per policy. This number depends on the EC2 instance type of the backup appliance.

NOTE

This section does not apply to the [VPC Configuration Backup policy](#) that protects the Amazon VPC configuration and settings.

Instance Type: T3.medium*

Resource	Maximum Workloads	Maximum Workloads per Backup Policy
EC2 instance	1,000	250
RDS instance	500	100
EFS file system	250	25
DynamoDB table	250	100

*Provided that a maximum of 100 AWS accounts is added to the backup appliance.

Instance Type: C5.9xlarge

Resource	Maximum Workloads	Maximum Workloads per Backup Policy
EC2 instance	10,000	1,000
RDS instance	2,500	1,000
EFS file system	1,000	100
DynamoDB table	1,000	150

*Provided that a maximum of 300 AWS accounts is added to the backup appliance.

Maximizing Throughput

The number of worker instances simultaneously deployed to process workloads added to a backup policy is defined by the speed of data upload to the backup repository specified for the policy. To maximize policy processing throughput, consider that every backup and archive session started during policy execution requires a separate worker instance to be deployed. For more details, see [Worker Instances](#).

Worker Instances

If you want initial full backups to be processed quickly, it is recommended to use a larger worker instance profile, and then change it to a smaller profile for incremental backup. You can change worker instance profile settings on a regional basis, so make sure that the worker instance size is appropriate to process the largest workload within the required time.

Each worker instance is deployed as an `amzn-linux-v2` image, and the binaries are downloaded from the connected S3 bucket. Instance types of worker instances sizes depend on the total EBS volume size.

Worker Profile	Default Instance Type	Usage
Small	c5.large	Processing EBS volumes under 1024 GB (default)
Medium	c5.2xlarge	Processing EBS volumes between 1024 GB and 16 TB (default)
Large	c5.4xlarge	Processing EBS volumes over 16 TB (default)
Archiving	c5.2xlarge	Processing EBS volumes under 6 TB
	c5.4xlarge	Processing EBS volumes over 6 TB

For more information on AWS pricing, see [AWS Documentation](#).

Deployment

To deploy Veeam Backup for AWS, do the following:

1. Deploy the backup server as described in the Veeam Backup & Replication User Guide, section [Installing Veeam Backup & Replication](#).

Alternatively, you can use a backup server that already exists in your backup infrastructure if it meets the AWS Plug-in for Veeam Backup & Replication [system requirements](#).

2. [Install AWS Plug-in for Veeam Backup & Replication on the backup server](#).
3. [Deploy a backup appliance in AWS](#).

Deploying Plug-In

If your installation package of Veeam Backup & Replication does not provide features that allow you to protect AWS resources, you must install AWS Plug-in for Veeam Backup & Replication on the backup server to be able to add your backup appliances to the backup infrastructure.

Installing Plug-In

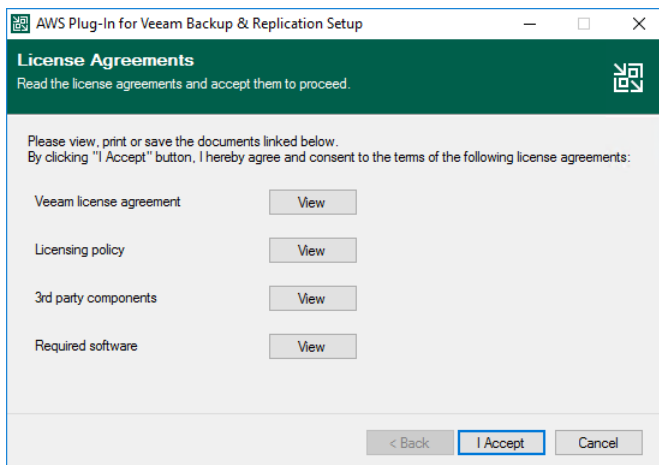
NOTE

Before you install AWS Plug-in for Veeam Backup & Replication, stop all running backup policies, disable all jobs, and close the Veeam Backup & Replication console.

To install AWS Plug-in for Veeam Backup & Replication, do the following:

1. Log in to the backup server using an account with the local Administrator permissions.
2. In a web browser, navigate to the [Veeam Backup & Replication: Download](#) page, switch to the **Cloud Plug-ins** in the **Additional Downloads** section, and click the **Download** icon to download AWS Plug-in for Veeam Backup & Replication.
3. Open the downloaded `AWSPugin_12.8.0.3264.zip` file and launch the `AWSPugin_12.8.0.3264.exe` installation file.
4. Complete the **AWS Plug-In for Veeam Backup & Replication Setup** wizard:
 - a. At the **License Agreements** step, read and accept the Veeam license agreement and licensing policy, as well as the license agreements of 3rd party components that Veeam incorporates, and the license agreements of required software. If you reject the agreements, you will not be able to continue installation.

To read the terms of the agreements, click **View**.
 - b. At the **Installation Path** step of the wizard, you can specify the installation directory. To do that, click **Browse**. In the **Browse for folder** window, select the installation directory for the product or create a new one, and click **OK**.
 - c. At the **Ready to Install** step, click **Install** to begin installation.



Installing and Uninstalling Plug-In in Unattended Mode

You can install or uninstall AWS Plug-in for Veeam Backup & Replication in the unattended mode using the command line interface. The unattended mode does not require user interaction – the installation runs automatically in the background, and you do not have to respond to the installation wizard prompts. You can use it to automate processes in large-scale environments.

To install AWS Plug-in for Veeam Backup & Replication in unattended mode, use either of the following options:

- If AWS Plug-in for Veeam Backup & Replication is a part of Veeam Backup & Replication installation package, follow the instructions provided in the Veeam Backup & Replication User Guide, section [Installing Veeam Backup & Replication in Unattended Mode](#).
- If AWS Plug-in for Veeam Backup & Replication is delivered as a separate .EXE file, use the instructions from this subsection.

Before You Begin

Before you start unattended installation, do the following:

1. Download the AWS Plug-in for Veeam Backup & Replication .EXE file as described in section [Installing Plug-In](#) (steps 1-2).
2. Check compatibility of the AWS Plug-in for Veeam Backup & Replication and Veeam Backup & Replication versions. For more information, see [System Requirements](#).

Installation Command-Line Syntax

Open the command prompt and run the .EXE file using the following parameters:

```
%path% /silent /accepteula /acceptlicensingpolicy /acceptthirdpartylicenses /acceptrequiredsoftware [/uninstall]
```

The following command-line parameters are used to run the setup file:

Parameter	Required	Description
%path%	Yes	Specifies a path to the installation .EXE file on the backup server or in a network shared folder.
/silent	Yes	Sets the user interface level to <i>None</i> , which means no user interaction is needed during installation.
/accepteula	Yes	Confirms that you accept the terms of the Veeam license agreement.
/acceptlicensingpolicy	Yes	Confirms that you accept the Veeam licensing policy.

Parameter	Required	Description
/acceptthirdpartylicenses	Yes	Confirms that you accept the license agreement for 3rd party components that Veeam incorporates.
/acceptrequiredsoftware	Yes	Confirms that you accept the license agreements for each required software that Veeam will install.
/uninstall	No	Uninstalls the plug-in. Example: "AWSPlugin_12.8.0.3264.exe /silent /accepteula /acceptlicensingpolicy /acceptthirdpartylicenses /acceptrequiredsoftware /uninstall"
/repair	No	Replaces missing files, firewall rules and registry keys. Example: "AWSPlugin_12.8.0/3264.exe /silent /accepteula /acceptlicensingpolicy /acceptthirdpartylicenses /acceptrequiredsoftware /repair"

Upgrading Plug-In

To upgrade AWS Plug-in for Veeam Backup & Replication, do the following:

1. Install the new version of AWS Plug-in for Veeam Backup & Replication as described in section [Installing Plug-In](#).
2. Upgrade backup appliances from the Veeam Backup & Replication console as described in section [Updating Appliances Using Console](#).


Uninstalling Plug-In

Before you uninstall AWS Plug-in for Veeam Backup & Replication, it is recommended to [remove all connected backup appliances](#) from the backup infrastructure. If you keep the backup appliances in the backup infrastructure, the following will happen:

- You will be able to see information on snapshots of EC2 instances and RDS resources, as well as backups of DynamoDB tables, Redshift clusters, EFS file systems, FSx file systems, and VPC configurations in the Veeam Backup & Replication console. However, you will not be able to perform any operations with these snapshots and backups.
- You will be able to see information on image-level backups of EC2 and DB instances and perform data recovery operations using these backups. However, restore of entire EC2 instances to AWS will start working as described in the Veeam Backup & Replication User Guide, section [How Restore to Amazon EC2 Works](#).
- You will be able to see information on backup policies. However, you will only be able to remove these policies from the Veeam Backup & Replication console.

To uninstall AWS Plug-in for Veeam Backup & Replication, do the following:

1. Log in to the backup server using an account with local Administrator permissions.
2. Open the **Start** menu, navigate to **Control Panel > Programs > Programs and Features**.
3. In the program list, click **AWS Plug-in for Veeam Backup & Replication** and click **Uninstall**.
4. In the opened window, click **Remove**.



NOTE

After you uninstall AWS Plug-in for Veeam Backup & Replication, you will be no longer able to add backup appliances and new external repositories to the backup infrastructure.

Deploying Backup Appliance

Veeam Backup for AWS is installed on a Linux-based EC2 instance that is created in a selected AWS account during the product installation. You can deploy Veeam Backup for AWS only from the Veeam Backup & Replication console.

NOTE

This section is intended for IT managers, virtual infrastructure administrators, backup administrators and other IT professionals who plan to deploy and use Veeam Backup for AWS.

This section assumes that users have basic knowledge of AWS EC2, Managing VPCs and understanding AWS IAM.

When deploying Veeam Backup for AWS, Veeam Backup & Replication performs the following steps:

1. Deploys an EC2 instance from the Ubuntu 22.04 LTS image.
2. Creates a temporary Amazon S3 bucket in AWS and uploads Veeam Backup for AWS installation packages and their dependencies to the bucket.
3. Installs the [required software components](#) on the EC2 instance.
4. Creates the following IAM roles in AWS and adds them to the EC2 instance running Veeam Backup for AWS:
 - **Impersonation IAM role** – will be attached to the backup appliance and then used to assume other IAM roles added to Veeam Backup for AWS.
 - **Default Backup Restore IAM role** – will be used to perform data protection and recovery operations within the AWS account to which the backup appliance belongs. Out of the box, the role is already assigned all the required permissions listed in section [Full List of IAM Permissions](#).

You will be able to add other IAM roles later, after Veeam Backup for AWS installation. For more information, see [Managing IAM Roles](#).

5. Removes the temporary Amazon S3 bucket from AWS.

How to Perform Appliance Deployment


To deploy a new backup appliance from the Veeam Backup & Replication console, do the following:

1. [Launch the New Veeam Backup for AWS Appliance wizard](#).
2. [Choose a deployment mode](#).
3. [Specify an AWS account in which the appliance will be deployed](#).
4. [Specify a name and description for the appliance](#).
5. [Specify the connection type](#).
6. [Specify network settings for the appliance](#).
7. [Specify credentials for the default user account](#).
8. [Wait for the appliance to be added to the backup infrastructure](#).
9. [Finish working with the wizard](#).

Step 1. Launch New Veeam Backup for AWS Appliance Wizard

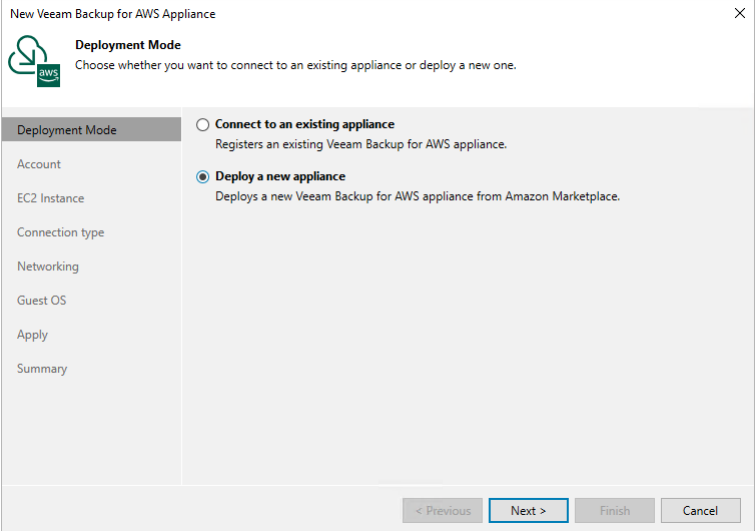
To launch the **New Veeam Backup for AWS Appliance** wizard, do one of the following:

1. In the Veeam Backup & Replication console, open the **Backup Infrastructure** view.
2. Navigate to **Managed Servers** and click **Add Server** on the ribbon.
Alternatively, you can right-click the **Managed Servers** node and select **Add Server**.
3. In the **Add Server** window:
 - a. [Applies only if you have several cloud plug-ins installed] Click **Veeam cloud-native backup appliance**.
 - b. Choose **Veeam Backup for AWS**.



Step 2. Choose Deployment Mode

At the **Deployment Mode** step of the wizard, select the **Deploy a new appliance** option.



Step 3. Specify AWS Account

At the **Account** step of the wizard, do the following:

1. From the **AWS account** drop-down list, select [access keys of an IAM user](#) that belongs to an AWS account in which the backup appliance will reside. Veeam Backup & Replication will use permissions of the specified IAM user to deploy the backup appliance, and further to connect to this appliance. For more information on the required permissions, see [Plug-in Permissions](#).

For access keys of an IAM user to be displayed in the **AWS account** drop-down list, the keys must be created in AWS and added to the Cloud Credentials Manager as described in the Veeam Backup & Replication User Guide, section [Access Keys for AWS Users](#). If you have not added the necessary keys to the Cloud Credentials Manager beforehand, you can do it without closing the **New Veeam Backup for AWS Appliance** wizard. To do that, click either the **Manage accounts** link or the **Add** button, and specify the access key and secret key in the **Credentials** window.

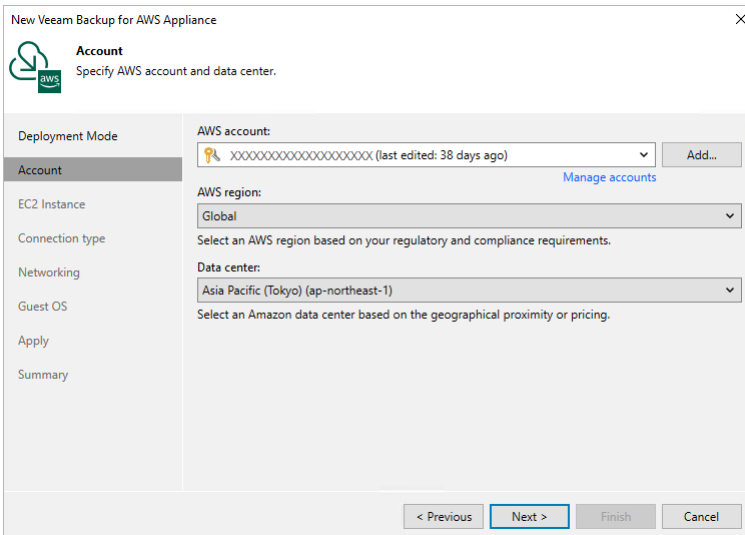
2. From the **AWS region** drop-down list, specify whether the backup appliance will reside in an AWS Global or AWS GovCloud (US) region.

IMPORTANT

To check the availability of the region, Veeam Backup & Replication by default establishes a temporary test connection with the US East (N. Virginia) region using endpoints of the [AWS Security Token Service \(STS\)](#) and [Amazon Elastic Compute Cloud \(EC2\)](#) AWS services. That is why the backup server must have access to this AWS Region. If you want to change the default region for a test connection, open a [support case](#).

3. From the **Data center** drop-down list, select an AWS Region where you want to deploy the backup appliance.

For more information on regions and availability zones, see [AWS Documentation](#).



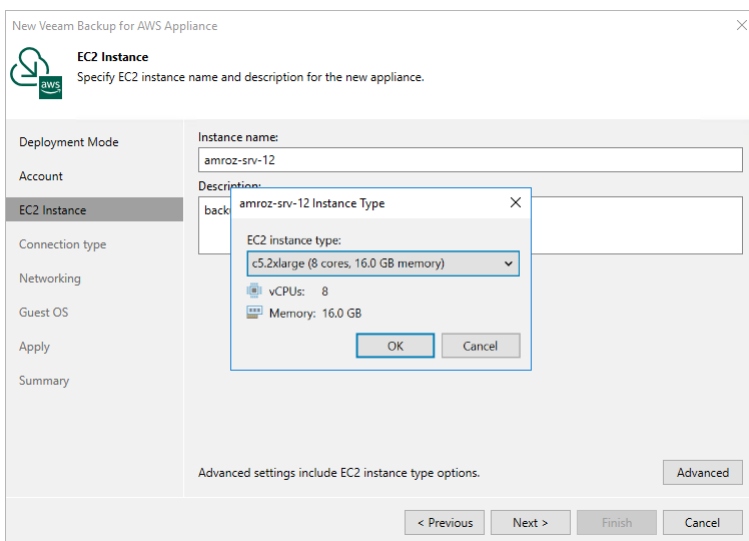
Step 4. Specify EC2 Instance Name and Description

At the **EC2 Instance** step of the wizard, specify a name and description for the EC2 instance where Veeam Backup for AWS will be deployed.

TIP

By default, Veeam Backup & Replication uses the recommended *t3.medium* EC2 instance type for the backup appliance. If you want to choose a specific machine type for the EC2 instance, click **Advanced** and select the necessary type in the **Instance Type** window.

For the list of all existing EC2 instance types, see [Sizing and Scalability Guidelines](#).



Step 5. Specify Connection Type

At the **Connection Type** step of the wizard, choose whether you want to assign a dynamic or a static (Elastic) public IP address, or a private IP address to the backup appliance. After the backup appliance is deployed, Veeam Backup & Replication will use the specified connection type to connect to the appliance.

To assign an Elastic IP address, you can either reserve a new address or specify an existing one:

- To reserve a new IP address, select the **(create new)** option from the **Use the following address** drop-down list.
- To assign an existing IP address, select it from the **Use the following address** drop-down list.

For an IP address to be displayed in the list of available addresses, it must be allocated to the AWS Region specified at [step 3](#) of the wizard, as described in [AWS Documentation](#). Note that elastic IP addresses that are used by any other EC2 instances are not displayed in the list.

For more information on Elastic IP addresses, see [AWS Documentation](#).

NOTE

If you choose the **Private IP address** option, you must allow communication between the Veeam Backup & Replication server and the backup appliance. One possible solution is to establish an AWS Site-to-Site VPN (Site-to-Site VPN) connection between the VPC of the appliance and your on-premises network, as described in [Configuring Access to Backup Appliances in AWS](#).

New Veeam Backup for AWS Appliance

Connection type
Specify how the backup appliance should be accessed.

Deployment Mode

Account

EC2 Instance

Connection type

Networking

Guest OS

Apply

Summary

Public IP address (dynamic)
Dynamic IP addresses may change after each appliance reboot.

Public IP address (static)
Use the following Elastic IP address:
(create new)

Private IP address
The backup appliance will have no public IP address assigned.
See [this link](#) to learn more.

< Previous Next > Finish Cancel

Step 6. Specify Network Settings

At the **Networking** step of the wizard, do the following:

1. Choose an Amazon virtual private cloud (VPC) to which the backup appliance will be connected.

You can create a new VPC or specify an existing one:

- [Applies only if you have selected to assign a public IP address to the backup appliance at the **Specify Connection Type** step of the wizard] To create a new VPC, select the **(create new)** option from the **Amazon VPC** drop-down list. Veeam Backup & Replication will automatically create a virtual network with a set of predefined security group rules.
- To specify an existing VPC, select it from the **Amazon VPC** drop-down list. For a VPC to be displayed in the list of available networks, it must be created in AWS for the region specified at [step 3](#) of the wizard, as described in [AWS Documentation](#).

2. Choose a subnet in which the backup appliance will be launched.

You can create a new subnet or specify an existing one:

- [Applies only if you have selected to assign a public IP address to the backup appliance at the **Specify Connection Type** step of the wizard] To create a new subnet, select the **(create new)** option from the **Subnet** drop-down list. Veeam Backup & Replication will automatically create a subnet in the specified VPC.
- To specify an existing subnet, select it from the **Subnet** drop-down list. For a subnet to be displayed in the list of available subnets, it must be created in the specified VPC as described in [AWS Documentation](#).

IMPORTANT

Consider the following:

- The specified Amazon VPC and subnet must have the outbound internet access to AWS services listed in section [AWS Services](#).
- The specified Amazon VPC and subnet must allow inbound internet access from both the backup server and a local machine that you plan to use to work with Veeam Backup for AWS.

To learn how to enable internet access for Amazon VPCs and subnets, see [AWS Documentation](#).

3. Choose a security group that will be associated with the backup appliance.

You can create a new security group or specify an existing one:

- [Applies only if you have selected to assign a public IP address to the backup appliance at the **Specify Connection Type** step of the wizard] To create a new security group, select the **(create new)** option from the **Security group** drop-down list. Veeam Backup & Replication will automatically create a group.
- To specify an existing security group, select it from the **Security group** drop-down list. For a security group to be displayed in the list of available groups, it must be created in AWS as described in [AWS Documentation](#).

IMPORTANT

If you select an existing security group, consider that security group rules must allow inbound internet access from both the backup server and a local machine that you plan to use to work with Veeam Backup for AWS. To learn how to create security group rules, see [AWS Documentation](#).

4. [Applies only if you have selected to assign a public IP address to the backup appliance at the **Specify Connection Type** step of the wizard] In the **Backup server public IP address** field, specify an IP address or a range of IP addresses that will be allowed to access the backup appliance.
 - If you have chosen to create a new security group, Veeam Backup & Replication will create a security rule for the specified IP address ranges. Note that the backup server IP address must fall into the specified IP address range.
 - If you have chosen to specify an existing security group, Veeam Backup & Replication will verify whether the security group allows inbound HTTPS traffic (port **443**) from the specified IP addresses. If the security group restricts inbound HTTPS traffic, you will not be able to proceed with the wizard.
5. [Applies only if you have selected to assign a private IP address to the backup appliance at the **Specify Connection Type** step of the wizard] In the **Backup server IP address** field, specify an IP address or a range of IP addresses that will be allowed to access the backup appliance. Note that the backup server IP address must fall into the specified IP address range.

Veeam Backup & Replication will verify whether the specified security group allows inbound HTTPS traffic (port **443**) from the specified IP addresses. If the security group restricts inbound HTTPS traffic, you will not be able to proceed with the wizard.

TIP

The IPv4 address ranges must be specified in the CIDR notation (for example, `12.23.34.0/24`). To specify multiple IP addresses or multiple IP address ranges, use a comma-separated list.

The screenshot shows the 'Networking' step of the 'New Veeam Backup for AWS Appliance' wizard. The interface includes a sidebar with navigation options: Deployment Mode, Account, EC2 Instance, Connection type, **Networking**, Guest OS, Apply, and Summary. The main content area is titled 'Networking' and contains the following settings:

- Deployment Mode:** Amazon VPC (selected)
- Account:** vpc-536c7d34 (Default)
- EC2 Instance:** Subnet (selected)
- Connection type:** subnet-ddeae786 172.31.0.0/20 (ap-northeast-1c)
- Networking:** Security group: default
- Guest OS:** default
- Apply:** Backup server public IP address: 89.185.226.18/32
- Summary:** Specify public IP or IP range from which backup appliance will be accessed.

At the bottom of the wizard, there are four buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'.

Step 7. Specify User Credentials

At the **Guest OS** step of the wizard, do the following:

1. From the **Create the following administrator credentials** drop-down list, select a user whose credentials Veeam Backup & Replication will use to create the default user account on the backup appliance.

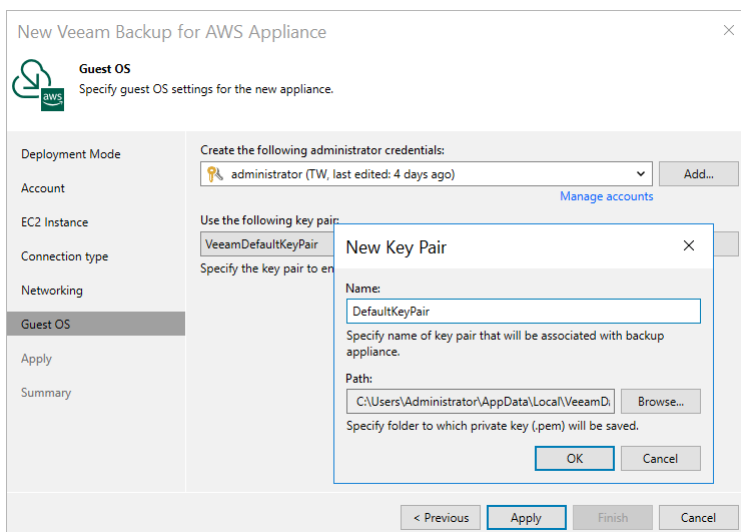
For a user to be displayed in the **Create the following administrator credentials** drop-down list, it must be added to the Credentials Manager as described in the Veeam Backup & Replication User Guide, section [Standard Accounts](#). If you have not added the necessary user to the Credential Manager beforehand, you can do it without closing the **New Veeam Backup for AWS Appliance** wizard. To do that, click either the **Manage accounts** link or the **Add** button, and specify the user name, password and description in the **Credentials** window.

IMPORTANT

The specified password must contain at least one special character, one lowercase and one uppercase letters, and must not contain monotonic sequence characters. The password length must be between 8 and 255 characters.


2. From the **Use the following key pair** drop-down list, select a key pair that will be used to authenticate against the backup appliance.

For a key pair to be displayed in the list of available keys, it must be created in AWS as described in [AWS Documentation](#). If you have not created the necessary key pair beforehand, you can do it without closing the **Guest OS** wizard. To do that, click **Add** and specify the private key name and folder path to the key in the **New Key Pair** window. Veeam Backup & Replication will create a key of the *ed25519* type.



Step 8. Track Progress

Veeam Backup & Replication will display the results of every step performed while deploying the backup appliance. At the **Apply** step of the wizard, wait for the process to complete and click **Next**.

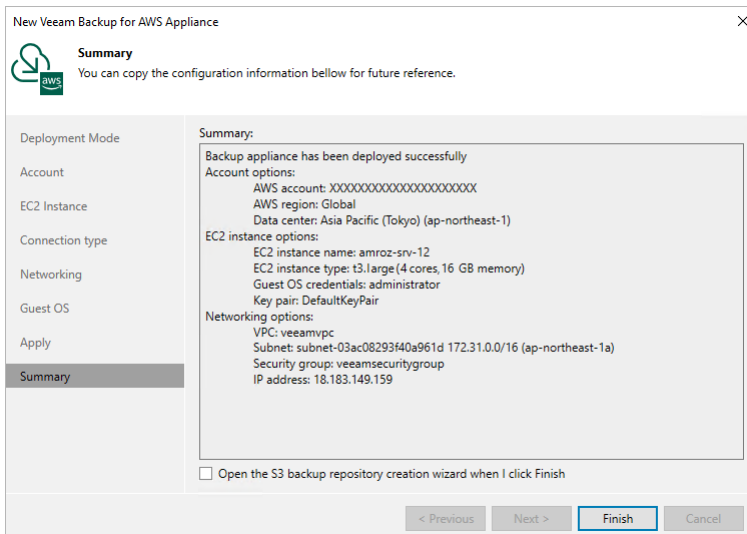


Step 9. Finish Working with Wizard

At the **Summary** step of the wizard, review summary information and click **Finish**. After the backup appliance is deployed, you will be able to configure its settings in the Veeam Backup for AWS Web UI as described in section [Configuring Veeam Backup for AWS](#).

TIP

If you want to add repositories immediately after the backup appliance is deployed, select the **Open the S3 backup repository creation wizard when I click Finish** check box and follow the instructions provided in section [Adding Backup Repositories Using Console](#).



Failure and Recovery

Taking into account that Veeam Backup for AWS deployed in the AWS environment works is protected using snapshots by default, we highly recommend to perform regular configuration backup. If a backup appliance goes down for some reason, you can quickly restore its configuration from the configuration backup. You can also use configuration backups to apply the configuration of one backup appliance to another one in the backup infrastructure. For more information, see [Performing Configuration Backup and Restore](#).

For application-related errors and issues, refer to the [Veeam Knowledge Base](#) or consider opening a [support ticket](#).

Licensing

Veeam Backup for AWS is licensed per protected instance. An instance is defined as a single AWS resource – EC2 instance, RDS resource, DynamoDB table, Redshift cluster, EFS file system, and so on. An instance is considered to be protected if it has a restore point (snapshot or backup) created by a backup policy during the past 31 days. Each protected instance consumes 1 license unit. However, if an instance has only manually created snapshots or backups, it does not consume any license units.

Product Editions

Veeam Backup for AWS is available in 3 editions:

- **Free** – allows you to protect up to 10 instances free of charge. This edition applies only to backup appliances that are no longer managed by Veeam Backup & Replication servers.

Note that this edition does not support indexing of EFS file systems, creating image-level backups of PostgreSQL DB instances, as well as protecting DynamoDB tables, Redshift clusters and FSx file systems.

- **BYOL (Bring Your Own License)** – allows you to protect the number of instances equivalent to the number of units specified in your license. This edition applies only to backup appliances that were previously deployed from the AWS Marketplace.

Veeam Backup for AWS *BYOL* edition can be licensed using the Veeam Universal License (VUL) installed on the Veeam Backup & Replication server. For more information on Veeam licensing terms and conditions, see [Veeam Licensing Policy](#).

IMPORTANT

If you plan to use the Veeam Universal License (VUL), consider that only the subscription license type is supported.

When the license expires, Veeam Backup for AWS offers a grace period to ensure a smooth license update and to provide sufficient time to install a new license file. The duration of the grace period is 31 days after the expiration of the license. During this period, you can perform all types of data protection and disaster recovery operations. After the grace period is over, Veeam Backup for AWS stops processing all instances and disables all scheduled backup policies. You must update your license before the end of the grace period.

To learn how to install the license on a backup appliance that was previously deployed from the AWS Marketplace, see [Appendix F. Installing and Removing License](#).

- **Paid** – allows you to protect an unlimited number of instances. This edition applies only to backup appliances that were previously deployed from the AWS Marketplace.

In the *Paid* edition of the product, you are charged by the number of instances that you actually protect. To track data protection operations on the backup appliance, Veeam Backup for AWS uses the [AWS Marketplace Metering Service](#). Every hour, the backup appliance sends information on the current number of protected instances to AWS. The billing for the protected instances is included into the monthly [AWS Cost and Usage report](#).

Limitations

Keep in mind the following limitations and considerations:

- If you use the *Veeam Cloud Connect service provider* license, the AWS Plug-in for Veeam Backup & Replication functionality is available from Veeam Service Provider Console only. For more information, see the Veeam Service Provider Console [Guide for Service Providers](#).
- If you have a *Perpetual* per-socket license installed on the backup server, and you want to add a backup appliance to the backup infrastructure, you must install an additional *Perpetual* per-instance license or a subscription license. When you install an additional license, the new license is automatically merged with the existing *Perpetual* per-socket license. For more information on the merging process, see the Veeam Backup & Replication User Guide, section [Merging Licenses](#).

If you do not install an additional *Perpetual* per-instance license or a subscription license, you will be able to use one free license instance per each socket (maximum 6 free instances per license). After you exceed the limit of free instances, Veeam Backup for AWS backup policies protecting resources that are not covered by the license will fail.

To obtain an additional license, contact a Veeam sales representative at [Sales Inquiry](#).

- If an instance has not been backed up within the past 31 days, Veeam Backup for AWS automatically revokes the license unit from the instance. If you need to manually revoke a license unit, follow the instructions provided in section [Revoking License Units](#).

Scenarios

Backup appliances managed by a Veeam Backup & Replication server use the same license that is installed on the backup server. To learn what types of licenses and licensing models are incorporated in Veeam solutions, see:

- The Veeam Backup & Replication User Guide, section [Licensing](#)
- The Veeam Backup & Replication Veeam Cloud Connect Guide, section [Licensing for Service Providers](#)

Licensing Scenarios

When you add a backup appliance to the backup infrastructure, the following scenarios are applied:

- If you [deploy a new backup appliance](#) from the Veeam Backup & Replication console, workloads start consuming license units from the license installed on the backup server after you create and run backup policies.

When you remove the backup appliance from the backup infrastructure, Veeam Backup & Replication stops counting backed-up workloads and Veeam Backup for AWS switches to the *Free* edition that allows you to protect up to 10 workloads free of charge.

- If you [connect to an existing backup appliance](#), the *BYOL* and *Paid* license installed on the appliance will be replaced with the license installed on the backup server. Protected instances start consuming license units from the license installed on the backup server only after the backup policy sessions run on the connected appliance.

When you remove the backup appliance from the backup infrastructure, Veeam Backup & Replication stop counting backed-up workloads. Veeam Backup for AWS continues using the license that had been used before you added the backup appliance to the backup infrastructure.

Licensing When Connection to Veeam Backup & Replication is Lost

Veeam Backup for AWS stores information on protected workloads licensed by Veeam Backup & Replication. This information allows you to back up workloads even if the connection between the backup appliance and backup server is lost. However, the following conditions must be met:

- The workload must have already been licensed by the backup server.
- The workload must be listed as licensed on the backup appliance side. For more information, see [Revoking License Units](#).
- The connection must be lost not more than 31 days ago.

Note that the loss of connection with Veeam Backup & Replication does not affect restore processes and creating of snapshots manually.

Viewing License Information


After you add a backup appliance to the backup infrastructure, you can view the number of protected workloads in the Veeam Backup & Replication console.

Viewing License Details in Veeam Backup & Replication Console

To view AWS Plug-in for Veeam Backup & Replication license details in the Veeam Backup & Replication console, open the main menu and select **License**.


The **License** tab of the **License Information** window provides general information on the currently installed AWS Plug-in for Veeam Backup & Replication license:

- **Status** – the license status. The status will depend on the license type, the number of days remaining until license expiration, the number of days remaining in the grace period (if any), and the number of workloads that exceeded the allowed increase limit (if any).
- **Type** – the license type (*Perpetual, Subscription, Rental, Evaluation, NFR, Free*).
- **Edition** – the license edition (*Community, Standard, Enterprise, Enterprise Plus*).
- **Support ID** – the ID of the contract (required for contacting Veeam Customer Support).
- **Licensed to** – the name of an organization to which the license was issued.
- **Package** – the software product for which the license was issued.
- **Instances** – the total number of license units included in the license file and the number of units consumed by protected workloads.
- **Support expiration date** – the date when the license will expire.



The **Instances** tab of the **License Information** window provides information on the currently protected workloads:

- **Type** – the type of protected workloads.
 - **Cloud VMs** – protected EC2 instances.
 - **Cloud Databases** – protected RDS resources including Aurora DB clusters.
 - **Cloud File Shares** – protected EFS file systems.
- **Count** – the number of protected workloads.
- **Multiplier** – the number of license units one protected workload consumes.
- **Instances** – the total number of the consumed license units.



Viewing License Details in Veeam Backup for AWS Web UI

To view details on the license that is currently installed on the backup appliance in the Veeam Backup for AWS Web UI, do the following:

1. Switch to the **Configuration** page.
2. Navigate to **Licensing > License Info**.

The licensing section provides general information on the Veeam Backup for AWS license:

- **Status** – the license status. The status depends on the license edition, the number of days remaining until license expiration and the number of days remaining in the grace period (if any).

- **Instances** – the total number of license units included in the license file and the number of units consumed by protected resources.

Each instance that has a restore point created in the past 31 days is considered to be protected and consumes 1 license unit. To view the list of instances that consume license units, switch to the **License Usage** tab.

- **Expiration Date** – the date when the license will expire.
- **License Type** – the license edition (*Free, Paid, Subscription*).

NOTE

Subscription is the name of the *BYOL* license in Veeam Backup for AWS.

- **License ID** – the unique identification number of the provided license file (required for contacting the Veeam Customer Support Team).
- **Licensed To** – the name of an organization to which the license was issued.
- **Support ID** – the unique identification number of the support contract (required for contacting the Veeam Customer Support Team).

The screenshot displays the Veeam Backup for AWS configuration page. The top navigation bar includes the Veeam logo, the title 'Veeam Backup for AWS', the server time 'Dec 2, 2020 11:51 AM', the user 'administrator Portal Administrator', and a 'Configuration' button. The left sidebar contains navigation options: 'Exit Configuration', 'Getting Started', 'Administration', 'Accounts', 'Repositories', 'Workers', 'Server settings', 'Settings', 'Licensing', and 'Support Information'. The main content area is divided into two tabs: 'License Info' (active) and 'License Usage'. Under 'License Info', there are buttons for 'Install License' and 'Remove License'. The license details are as follows:

Status:	Valid (180 days until expiration)
Instances:	100 (23 used)
Expiration Date:	06/01/2022 12:00:00 AM
License Type:	Subscription
License ID:	G1R1925L-GFH1-ZW12-3335-VDH2XX21391
Licensed To:	Veeam Software Group GmbH
Support ID:	02067762


Revoking License Units

By default, Veeam Backup for AWS automatically revokes a license unit from a protected instance if no new restore points have been created by the backup policy during the past 31 days. However, you can manually revoke license units from protected instances – this can be helpful, for example, if you remove a number of instances from a backup policy and do not want to protect them anymore.

Revoking License Units Using Veeam Backup & Replication Console

To revoke a license unit from a protected instance in the Veeam Backup & Replication console, do the following:

1. In the Veeam Backup & Replication console, open the main menu and select **License**.
2. In the **License Information** window, switch to the **Instances** tab and click **Manage**.
3. In the **Licensed Instances** window, select a protected workload and click **Revoke**. Veeam Backup & Replication will revoke a license unit from the selected workload.



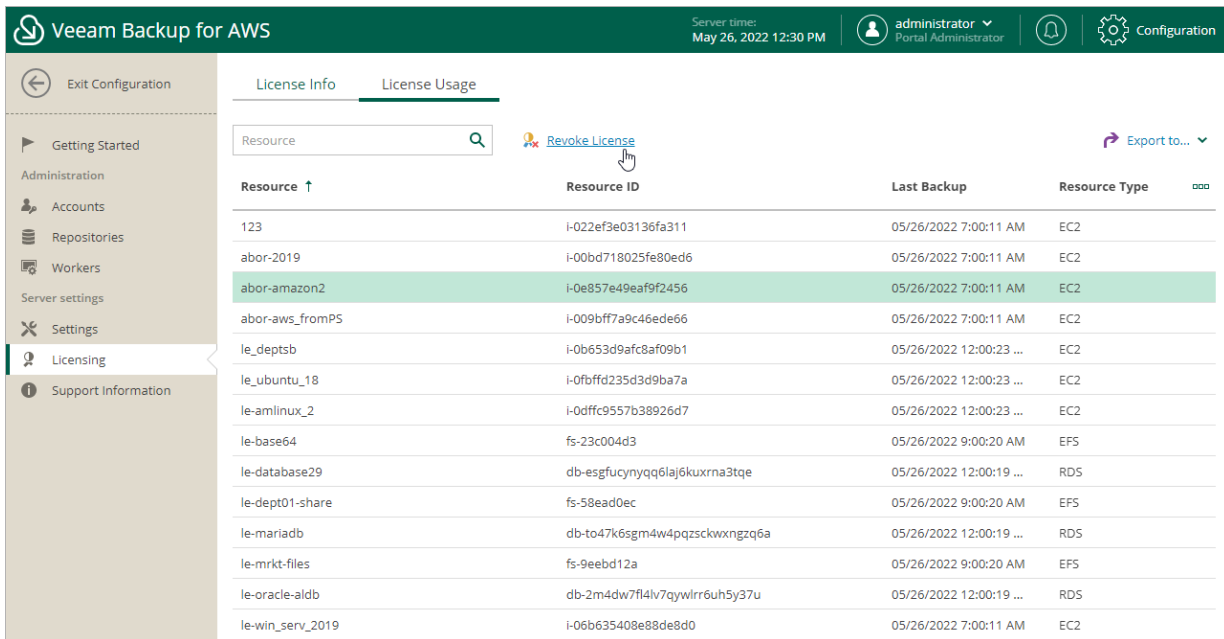
Revoking License Units Using Veeam Backup for AWS Web UI

To revoke a license unit from a protected instance in the Veeam Backup for AWS Web UI, do the following:

1. Switch to the **Configuration** page.
2. Navigate to **Licensing > License Usage**.
3. Select the instance that you no longer want to protect.

4. Click **Revoke License**.

5. In the **Revoke License** window, click **Yes** to confirm that you want to revoke the license unit.



The screenshot shows the Veeam Backup for AWS interface. The top navigation bar includes the Veeam logo, the text "Veeam Backup for AWS", the server time "May 26, 2022 12:30 PM", the user "administrator Portal Administrator", and a "Configuration" button. On the left, a sidebar menu lists "Getting Started", "Administration" (Accounts, Repositories, Workers), "Server settings" (Settings), "Licensing", and "Support Information". The main content area is titled "License Usage" and features a search bar, a "Revoke License" button, and an "Export to..." dropdown. Below this is a table with columns for Resource, Resource ID, Last Backup, and Resource Type. The row for "abor-amazon2" is highlighted in green.

Resource ↑	Resource ID	Last Backup	Resource Type
123	i-022ef3e03136fa311	05/26/2022 7:00:11 AM	EC2
abor-2019	i-00bd718025fe80ed6	05/26/2022 7:00:11 AM	EC2
abor-amazon2	i-0e857e49eaf9f2456	05/26/2022 7:00:11 AM	EC2
abor-aws_fromP5	i-009bff7a9c46ede66	05/26/2022 7:00:11 AM	EC2
le_deptsb	i-0b653d9afc8af09b1	05/26/2022 12:00:23 ...	EC2
le_ubuntu_18	i-0fbffd235d3d9ba7a	05/26/2022 12:00:23 ...	EC2
le-amlinux_2	i-0dff9557b38926d7	05/26/2022 12:00:23 ...	EC2
le-base64	fs-23c004d3	05/26/2022 9:00:20 AM	EFS
le-database29	db-esgfucyryqq6laj6kuxrna3tqe	05/26/2022 12:00:19 ...	RDS
le-dept01-share	fs-58ead0ec	05/26/2022 9:00:20 AM	EFS
le-mariadb	db-to47k6sgm4w4pqzscckwxngzq6a	05/26/2022 12:00:19 ...	RDS
le-mrkt-files	fs-9eebd12a	05/26/2022 9:00:20 AM	EFS
le-oracle-aldb	db-2m4dw7f4lv7qywlr6uh5y37u	05/26/2022 12:00:19 ...	RDS
le-win_serv_2019	i-06b635408e88de8d0	05/26/2022 7:00:11 AM	EC2

Accessing Veeam Backup for AWS

After you install Veeam Backup for AWS and add appliances to the backup infrastructure, you will be able to back up and restore AWS resources using both the Veeam Backup & Replication console and the Veeam Backup for AWS appliance Web UI.

Accessing Veeam Backup & Replication Console

The Veeam Backup & Replication console is a client-side component of the backup infrastructure that provides access to the backup server. The console allows you to log in to Veeam Backup & Replication and to perform data protection and disaster recovery operations on the server. To learn how to access the Veeam Backup & Replication console, see the Veeam Backup & Replication User Guide, section [Logging in to Veeam Backup & Replication](#).


By default, the Veeam Backup & Replication console is installed on the backup server automatically when you install Veeam Backup & Replication. However, in addition to the default console, you can install the Veeam Backup & Replication console on a dedicated machine to access the backup server remotely. To learn how to install Veeam Backup & Replication console, see the Veeam Backup & Replication User Guide, section [Installing Veeam Backup & Replication Console](#).

Accessing Web UI from Console

To access the Veeam Backup for AWS Web UI from the Veeam Backup & Replication console, do the following:

1. Open the **Backup Infrastructure** view.
2. Navigate to **Managed Servers**.
3. Select the backup appliance whose Web UI you want to open, and click **Open Console** on the ribbon.
Alternatively, you can right-click the appliance and select **Open console**.

Veeam Backup & Replication will open the Veeam Backup for AWS Web UI in your default web browser.



Accessing Web UI from Workstation

To access Veeam Backup for AWS, in a web browser, navigate to the Veeam Backup for AWS web address. The address consists of a public IPv4 address or DNS hostname of the backup appliance. Note that the website is available over HTTPS only.

IMPORTANT

Consider the following:

- If the backup appliance is deployed without a public IP address, you must establish a connection between the VPC of the appliance and your on-premises network to access Veeam Backup for AWS. For more information, see [Configuring Access to Backup Appliances in AWS](#).
- Internet Explorer is not supported. To access Veeam Backup for AWS, use Microsoft Edge (latest version), Mozilla Firefox (latest version) or Google Chrome (latest version).

You can access Veeam Backup for AWS using a local user account or a user account of an external identity provider. To learn how to add user accounts to Veeam Backup for AWS, see [Adding User Accounts](#).

NOTE

The web browser may display a warning notifying that the connection is untrusted. To eliminate the warning, you can replace the TLS certificate that is currently used to secure traffic between the browser and the backup appliance with a trusted TLS certificate. To learn how to replace certificates, see [Replacing Security Certificates](#).

Logging In Using Local User Account

To log in using credentials of a Veeam Backup for AWS user account, do the following:

1. In the **Username** and **Password** fields, specify credentials of the user account.

If you log in for the first time, use credentials of the default user that was created after the product installation. In future, you can add other user accounts to grant access to Veeam Backup for AWS. For more information, see [Managing User Accounts](#).

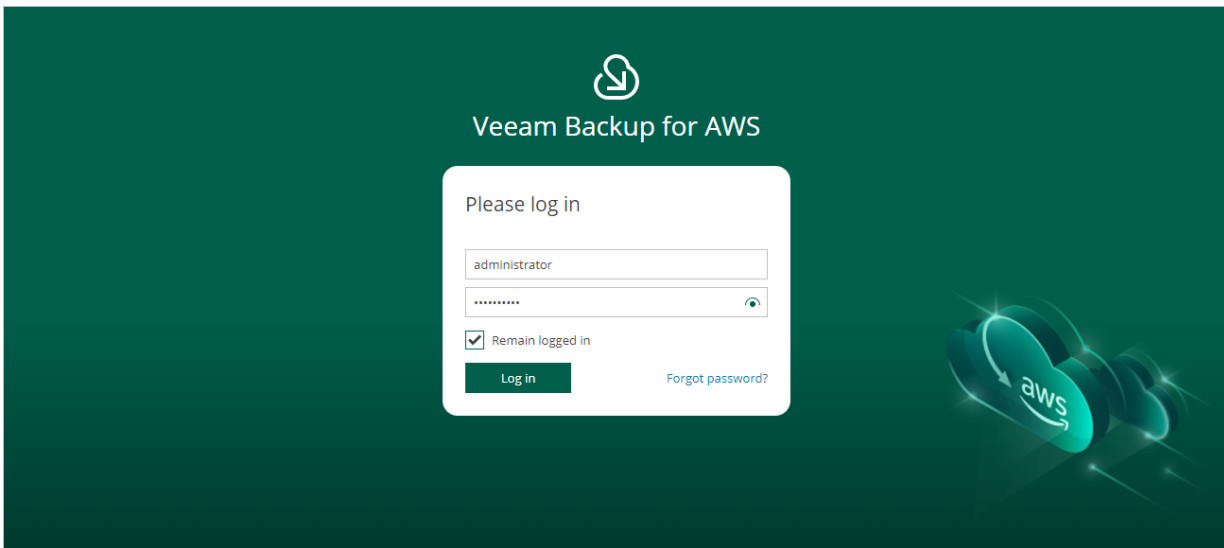
TIP

If you do not remember the user password, you can reset it. To do that, click the **Forgot password?** link and follow the instructions provided in [this Veeam KB article](#).

2. Select the **Remain logged in** check box to save the specified credentials in a persistent browser cookie so that you do not have to provide credentials every time you access Veeam Backup for AWS in a new browser session.

3. Click **Log in**.

If [multi-factor authentication \(MFA\) is enabled](#) for the user, Veeam Backup for AWS will prompt you to enter a code to verify the user identity. In the **Verification code** field, enter the temporary six-digit code generated by the authentication application running on your trusted device. Then, click **Log in**.



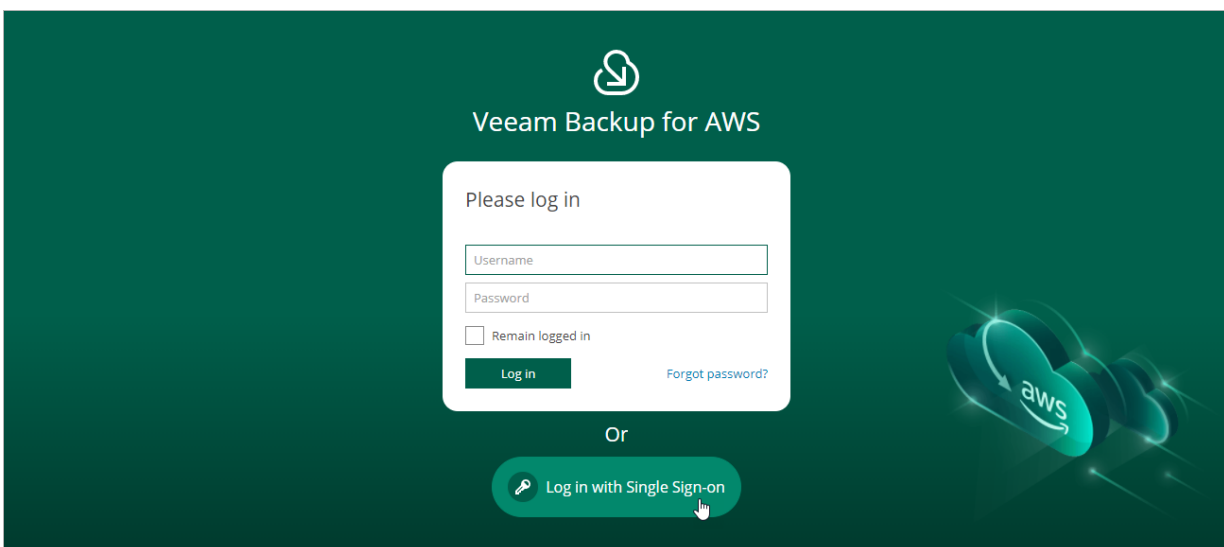
Logging In Using Identity Provider User Account

IMPORTANT

To access Veeam Backup for AWS under a user account of your identity provider, you must first [configure single sign-on settings](#) and then [add the identity provider user account](#) to Veeam Backup for AWS.

To log in using an identity provider, do the following:

1. Click **Log in with Single Sign-On**. You will be redirected to your identity provider portal.
2. If you have not logged in yet, log in to the identity provider portal. After that, you will be redirected to the **Veeam Backup for AWS Overview** page as an authorized user.



Logging Out

To log out, at the top right corner of the Veeam Backup for AWS window, click the user name and then click **Log out**.

Configuring Veeam Backup for AWS

To start working with Veeam Backup for AWS, perform a number of steps for its configuration:

1. [Add backup appliances to the backup infrastructure.](#)
2. [Add repositories that will be used to store backed-up data.](#)

This step applies if you plan to protect EC2 or DB instances with image-level backups, to perform EFS indexing operations, to back up Veeam Backup for AWS configuration and to keep additional copies of Amazon VPC configuration backups in Amazon S3.

3. Configure the added backup appliances:
 - a. [Add IAM roles to access AWS services and resources.](#)
 - b. [\[Optional\] Add users to control access to Veeam Backup for AWS.](#)
 - c. [Add database accounts to access databases of PostgreSQL DB instances.](#)
 - d. [\[Optional\] Configure worker instance settings.](#)

If you do not configure settings for worker instances, Veeam Backup for AWS will use the default settings of AWS Regions where worker instances will be deployed.

- e. [\[Optional\] Configure global retention, email notification and single-sign-on settings.](#)

NOTE

Even after you add IAM roles that manage your AWS resources and configure all the necessary settings, Veeam Backup for AWS will not populate [the list of resources on the Resources page](#) – unless you create backup policies and specify regions where the AWS resources belong, as described in section [Performing Backup](#).

Managing Backup Appliances

AWS Plug-in for Veeam Backup & Replication allows you to add backup appliances to the backup infrastructure, and to view and manage all the added appliances from the Veeam Backup & Replication console.

Adding Appliances

After you install AWS Plug-in for Veeam Backup & Replication, you must add backup appliances to the backup infrastructure. To do that, use either of the following options:

- [Deploy new Veeam Backup for AWS appliances](#) from the Veeam Backup & Replication console.
- [Connect to existing Veeam Backup for AWS appliances](#) if you have already deployed them.

NOTE

One backup appliance can be managed by one backup server only. If you add the appliance to the backup infrastructure of another backup server, the synchronization between the appliance and the previous backup server will be terminated, and appliance will be displayed as unavailable.

Connecting to Existing Appliances

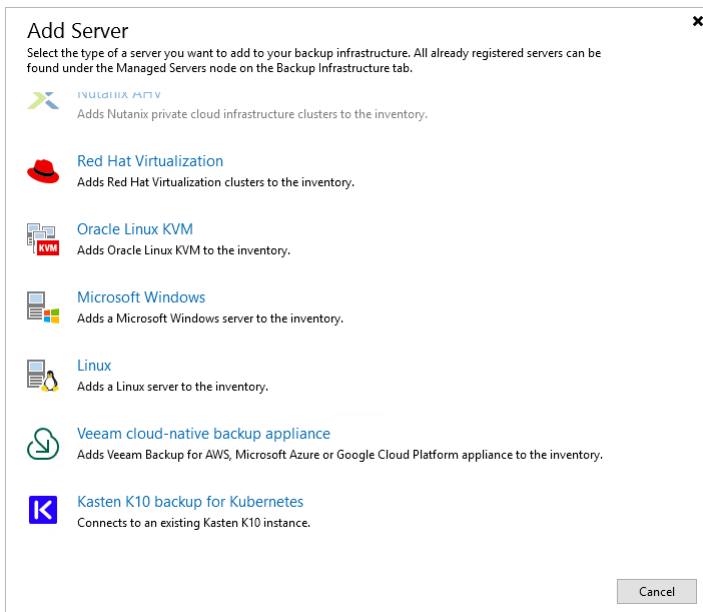
If you have already [deployed a backup appliance](#), you can add the appliance to the backup infrastructure:

1. [Launch the New Veeam Backup for AWS Appliance wizard](#).
2. [Choose a deployment mode](#).
3. [Specify an AWS account in which the appliance resides](#).
4. [Choose the appliance that you want to connect to](#).
5. [Specify the connection type](#).
6. [Specify a user whose credentials will be used to connect to the appliance](#).
7. [Configure repository settings](#).
8. [Wait for the appliance to be added to the backup infrastructure](#).
9. [Finish working with the wizard](#).

Step 1. Launch New Veeam Backup for AWS Appliance Wizard

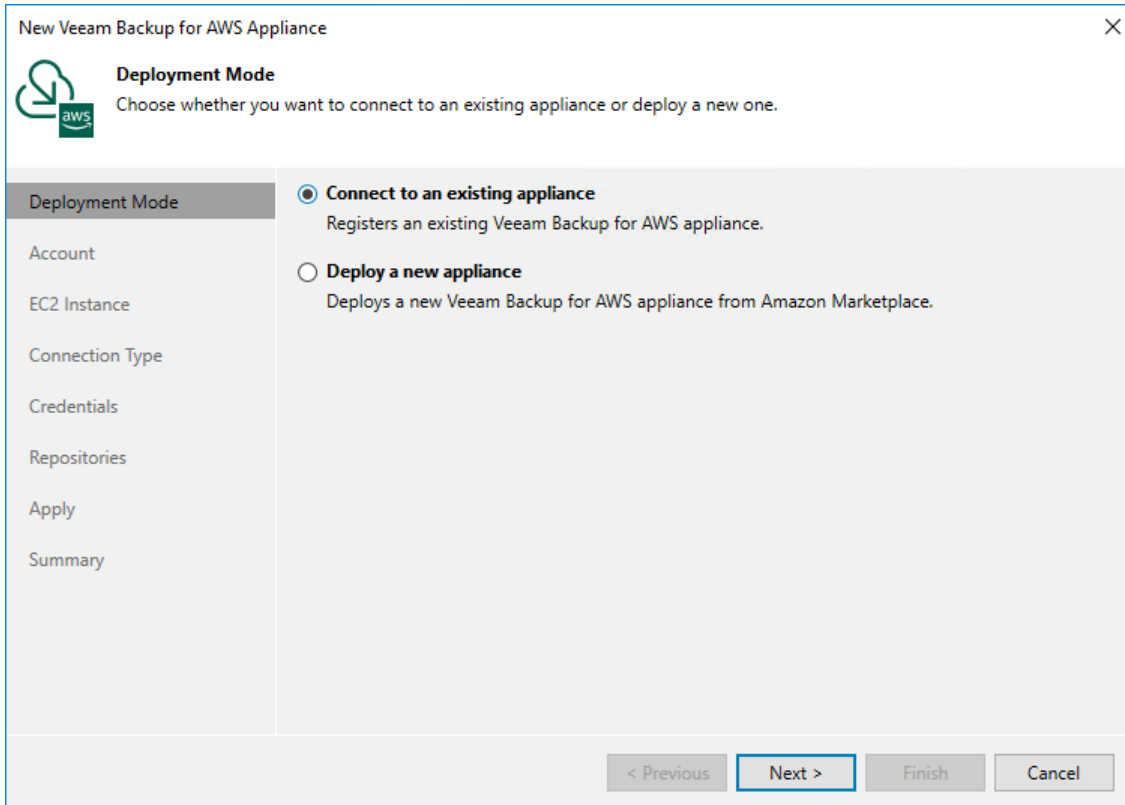
To launch the **New Veeam Backup for AWS Appliance** wizard, do one of the following:

1. In the Veeam Backup & Replication console, open the **Backup Infrastructure** view.
2. Navigate to **Managed Servers** and click **Add Server** on the ribbon.
Alternatively, you can right-click the **Managed Servers** node and select **Add Server**.
3. In the **Add Server** window:
 - a. [Applies only if you have several cloud plug-ins installed] Click **Veeam cloud-native backup appliance**.
 - b. Choose **Veeam Backup for AWS**.



Step 2. Choose Deployment Mode

At the **Deployment Mode** step of the wizard, select the **Connect to an existing appliance** option.



The screenshot shows a wizard window titled "New Veeam Backup for AWS Appliance" with a close button (X) in the top right corner. The window contains a sidebar on the left with the following steps: Deployment Mode (selected), Account, EC2 Instance, Connection Type, Credentials, Repositories, Apply, and Summary. The main area is titled "Deployment Mode" and includes the instruction: "Choose whether you want to connect to an existing appliance or deploy a new one." There are two radio button options: "Connect to an existing appliance" (which is selected) and "Deploy a new appliance". Below the options are four buttons: "< Previous", "Next >" (highlighted with a blue border), "Finish", and "Cancel".

Deployment Mode
Choose whether you want to connect to an existing appliance or deploy a new one.

Connect to an existing appliance
Registers an existing Veeam Backup for AWS appliance.

Deploy a new appliance
Deploys a new Veeam Backup for AWS appliance from Amazon Marketplace.

< Previous **Next >** Finish Cancel

Step 3. Specify AWS Account Settings

At the **Account** step of the wizard, do the following:

1. From the **AWS account** drop-down list, select access keys of an IAM user that belongs to an AWS account in which the backup appliance has been deployed. Veeam Backup & Replication will use permissions of the specified IAM user to connect to the backup appliance. For more information on the required permissions, see [Plug-in Permissions](#).

For access keys of an IAM user to be displayed in the **AWS account** drop-down list, they must be created in AWS and added to the Cloud Credentials Manager as described in the Veeam Backup & Replication User Guide, section [Access Keys for AWS Users](#). If you have not added the necessary keys to the Cloud Credentials Manager beforehand, you can do it without closing the wizard. To do that, click either the **Manage accounts** link or the **Add** button, and specify the access key and secret key in the **Credentials** window.

2. From the **AWS region** drop-down list, specify whether the backup appliance resides in the AWS Global or AWS GovCloud (US) region.

IMPORTANT

To check the availability of the region, Veeam Backup & Replication by default establishes a temporary test connection with the US East (N. Virginia) region using endpoints of the [AWS Security Token Service \(STS\)](#) and [Amazon Elastic Compute Cloud \(EC2\)](#) AWS services. That is why the backup server must have access to this AWS Region. If you want to change the default region for a test connection, open a [support case](#).

3. From the **Data center** drop-down list, select the AWS Region in which the backup appliance resides.

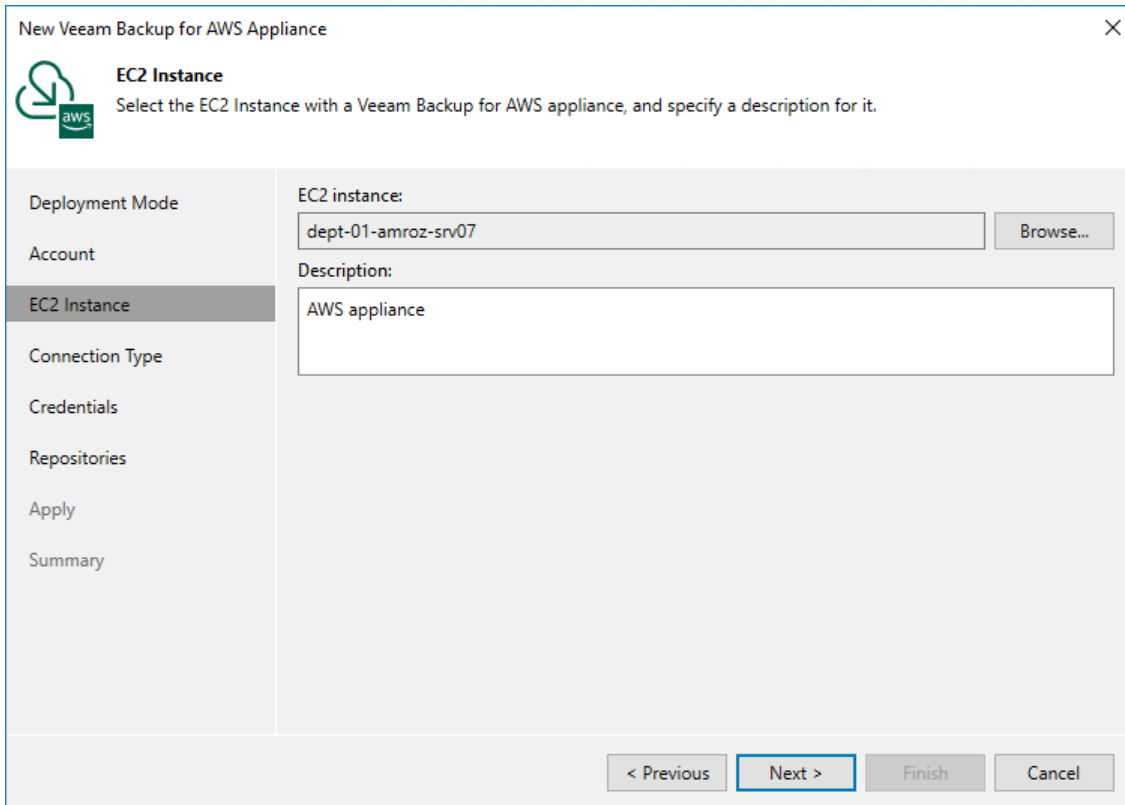
For more information on regions and availability zones, see [AWS Documentation](#).

The screenshot shows the 'New Veeam Backup for AWS Appliance' wizard window, specifically the 'Account' step. The window title is 'New Veeam Backup for AWS Appliance' with a close button (X) in the top right corner. The main heading is 'Account' with a sub-heading 'Specify AWS account and data center.' Below this, there is a sidebar on the left with navigation options: 'Deployment Mode', 'Account' (selected), 'EC2 Instance', 'Connection Type', 'Credentials', 'Repositories', 'Apply', and 'Summary'. The main content area contains three sections: 'AWS account:' with a dropdown menu showing 'XXXXXXXXXXXXXXXX (last edited: less than a day ago)' and an 'Add...' button; 'AWS region:' with a dropdown menu showing 'Global' and a 'Manage accounts' link; and 'Data center:' with a dropdown menu showing 'EU (Paris) (eu-west-3)'. At the bottom of the window, there are four buttons: '< Previous', 'Next >' (highlighted), 'Finish', and 'Cancel'.

Step 4. Specify Veeam Backup for AWS Appliance

At the **EC2 Instance** step of the wizard, choose the backup appliance that you want to add to the backup infrastructure:

1. Click **Browse**.
2. In the **EC2 Instance** window, select the necessary appliance and click **OK**.
3. In the **Description** field, specify a description for future reference.



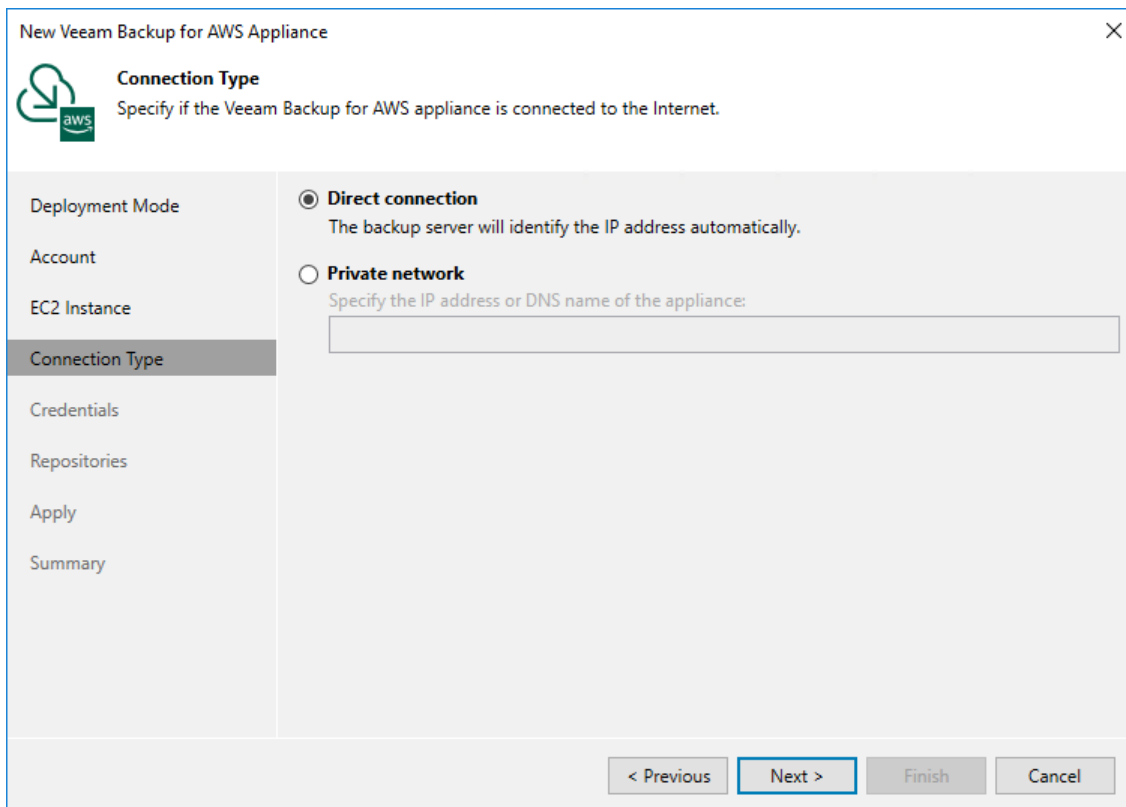
The screenshot shows a wizard window titled "New Veeam Backup for AWS Appliance" with a close button (X) in the top right corner. The window features a sidebar on the left with the following steps: Deployment Mode, Account, **EC2 Instance** (highlighted), Connection Type, Credentials, Repositories, Apply, and Summary. The main area is titled "EC2 Instance" and contains the instruction: "Select the EC2 Instance with a Veeam Backup for AWS appliance, and specify a description for it." Below this instruction, there are two input fields: "EC2 instance:" with the value "dept-01-amroz-srv07" and a "Browse..." button to its right; and "Description:" with the value "AWS appliance". At the bottom of the window, there are four buttons: "< Previous", "Next >" (highlighted with a blue border), "Finish", and "Cancel".

Step 5. Specify Connection Type

At the **Connection Type** step of the wizard, specify the way Veeam Backup & Replication will connect to the backup appliance:

- Select the **Direct connection** option if the backup appliance is connected to a VPC with the inbound internet access allowed and you want the backup server to connect to this Veeam Backup for AWS appliance over the internet. In this case, Veeam Backup & Replication will detect the public IP of the Veeam Backup for AWS appliance automatically.
- Select the **Private network** option if the backup appliance and the backup server are deployed within the same VPC, or if the backup appliance is deployed without a public IP address. In this case, you must specify the private IP address or DNS hostname of the backup appliance in the **Specify the IP address or DNS name of the appliance** field.

Note that you will have to establish connection between the VPC of the appliance deployed in a private environment and your on-premises network to allow a Veeam Backup & Replication server to communicate with the backup appliance. For more information, see [Backup Appliances in Private Environment](#).



The screenshot shows a wizard window titled "New Veeam Backup for AWS Appliance" with a close button (X) in the top right corner. The window has a sidebar on the left with the following items: Deployment Mode, Account, EC2 Instance, Connection Type (highlighted), Credentials, Repositories, Apply, and Summary. The main area is titled "Connection Type" and contains the instruction "Specify if the Veeam Backup for AWS appliance is connected to the Internet." There are two radio button options: "Direct connection" (selected) with the subtext "The backup server will identify the IP address automatically." and "Private network" with the subtext "Specify the IP address or DNS name of the appliance:" and an empty text input field below it. At the bottom of the window are four buttons: "< Previous", "Next >" (highlighted), "Finish", and "Cancel".

Step 6. Specify User Credentials

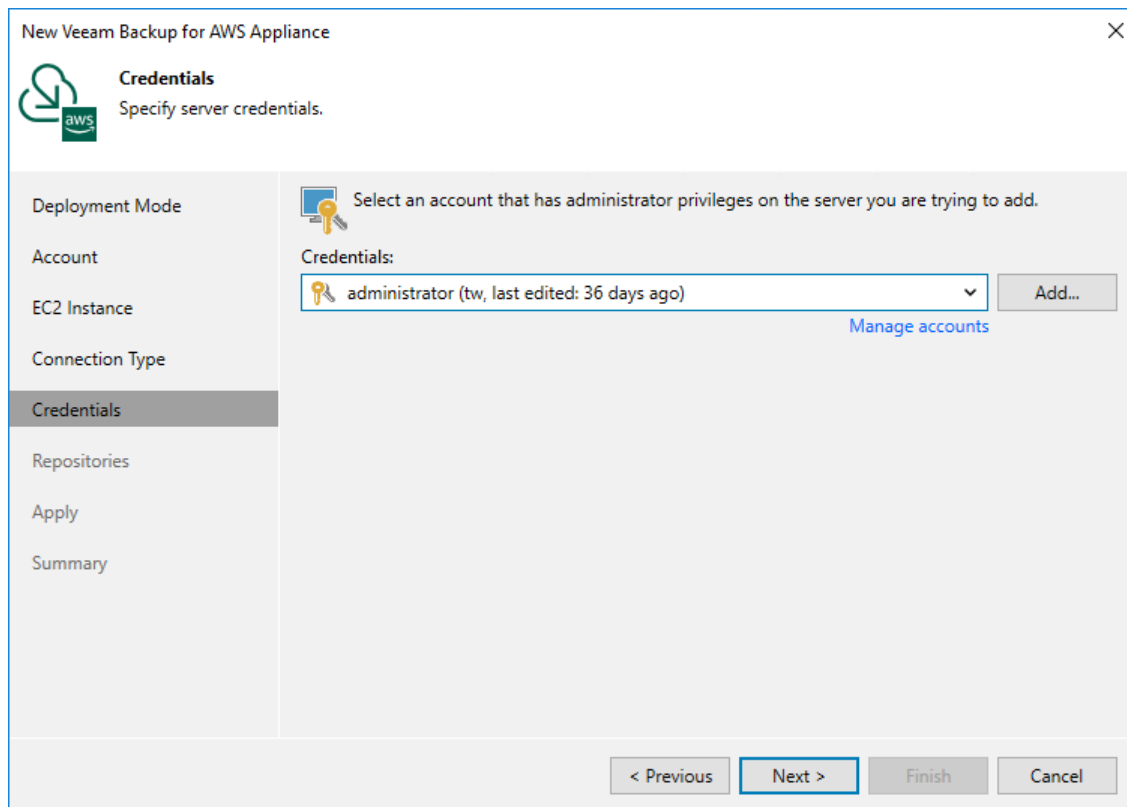
At the **Credentials** step of the wizard, specify a user whose credentials Veeam Backup & Replication will use to connect to the backup appliance.

For a user to be displayed in the **Credentials** list, it must be added to the Credentials Manager as described in the Veeam Backup & Replication User Guide, section [Standard Accounts](#). If you have not added the necessary user to the Credentials Manager beforehand, you can do it without closing the **New Veeam Backup for AWS Appliance** wizard. To do that, click either the **Manage accounts** link or the **Add** button, and specify the user name, password and description in the **Credentials** window.

IMPORTANT

Consider the following:

- The security group associated with the backup appliance must allow inbound HTTPS traffic (port **443**) from the backup server IP address. Otherwise, you will not be able to proceed with the wizard.
- The specified user must have multi-factor authentication (MFA) disabled and the *Portal Administrator* role assigned.



The screenshot shows the 'New Veeam Backup for AWS Appliance' wizard window, specifically the 'Credentials' step. The window title is 'New Veeam Backup for AWS Appliance' and it has a close button (X) in the top right corner. The main heading is 'Credentials' with the subtitle 'Specify server credentials.' Below this, there is a navigation pane on the left with options: Deployment Mode, Account, EC2 Instance, Connection Type, Credentials (selected), Repositories, Apply, and Summary. The main area contains a key icon and the instruction 'Select an account that has administrator privileges on the server you are trying to add.' Underneath, there is a 'Credentials:' label and a dropdown menu showing 'administrator (tw, last edited: 36 days ago)'. To the right of the dropdown is an 'Add...' button. Below the dropdown is a blue link labeled 'Manage accounts'. At the bottom of the window, there are four buttons: '< Previous', 'Next >' (highlighted), 'Finish', and 'Cancel'.

NOTE

As soon as you click **Next**, Veeam Backup & Replication will verify connection to the specified backup appliance. If the version of the appliance is not compatible with the Veeam Backup & Replication version or if the TLS certificate used to connect to the Veeam Backup for AWS Web UI is not trusted, you will receive a warning. To learn how to eliminate this warning, see [Eliminating Warnings](#).

Eliminating Warnings

If Veeam Backup & Replication encounters an issue while verifying the connection to the specified backup appliance, you may get one of the following warnings.

Version Compatibility Alert

If you try to add to the backup infrastructure an appliance that runs a version of Veeam Backup for AWS that is not [compatible with the version](#) of Veeam Backup & Replication, Veeam Backup & Replication will display a warning notifying that the appliance must be upgraded. To eliminate the warning, click **Yes**. Veeam Backup & Replication will automatically upgrade the appliance to the necessary version.

During [upgrade to version 8.0](#), Veeam Backup & Replication will verify whether the IAM user whose access keys are used to connect to the appliance has sufficient permissions to upgrade the appliance. If some permissions are missing, you will receive a warning.

You can manually grant missing permissions to the IAM user using AWS or instruct Veeam Backup & Replication to do it:

- If you want to grant the missing permissions manually, do the following:
 - a. Click **Copy permissions to Clipboard**.

Note that the list of copied permissions will contain all the permissions required to perform the upgrade operation, not the list of missing permissions.
 - b. In AWS, create an IAM policy with the missing permissions and attach the policy to the IAM user whose access key are used to connect to the appliance.

To learn how to create IAM policies, see [Appendix B. Creating IAM Policies in AWS](#).
 - c. Back to the Veeam Backup & Replication console, click **Proceed**.
- If you want to instruct Veeam Backup & Replication to grant the missing permissions automatically, click **Grant** and provide one-time access keys of an IAM user that is [authorized to grant IAM permissions](#) in the opened window. Note that the specified user must belong to the same AWS account in which the Veeam Backup for AWS appliance is deployed.

Veeam Backup & Replication will create an IAM policy with missing permissions and attach the policy to the IAM user whose permissions are used to connect to the appliance.

NOTE

Veeam Backup & Replication does not store the provided one-time access keys in the configuration database.

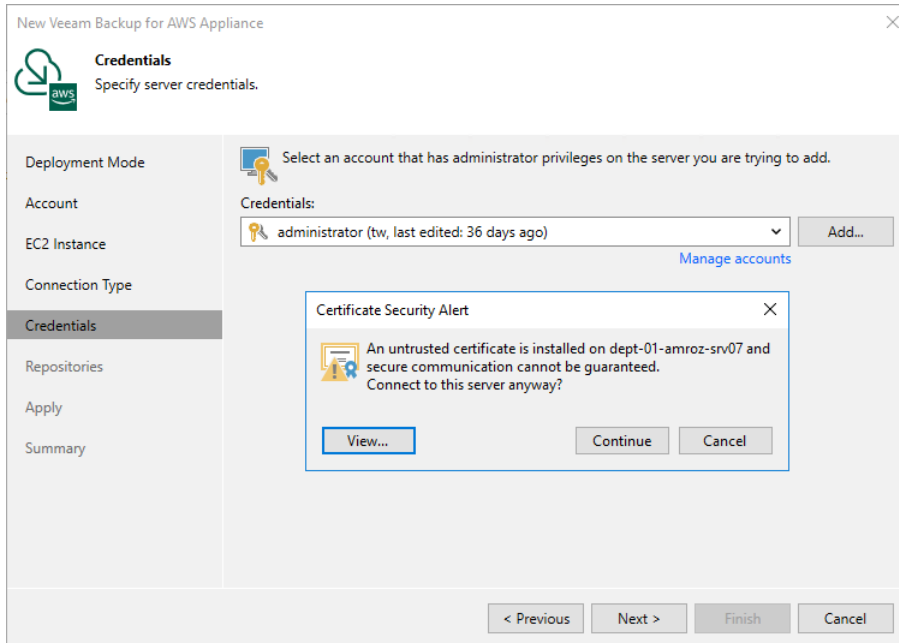
Certificate Security Alert

When you add a backup appliance to the backup infrastructure, Veeam Backup & Replication saves in the configuration database a thumbprint of the TLS certificate installed on the appliance. When Veeam Backup & Replication connects to the appliance, it uses the saved thumbprint to verify the appliance identity and to avoid the man-in-the-middle attack. For more information on managing TLS certificates, see [Replacing Security Certificates](#).

If the certificate installed on the backup appliance is not trusted, Veeam Backup & Replication will display a warning notifying that secure connection cannot be guaranteed. You can view the certificate and click **Continue** – in this case, Veeam Backup & Replication will remember the certificate thumbprint and will further trust the certificate when connecting to the appliance. Otherwise, you will not be able to proceed with the wizard.

NOTE

When you update a TLS certificate installed on a backup appliance, this appliance becomes unavailable in the Veeam Backup & Replication console. To make the appliance available again, acknowledge the new certificate at the **Credentials** step of the [Edit Veeam Backup for AWS Appliance wizard](#).



Step 7. Configure Repository Settings

At the **Repositories** step of the wizard, a list of all standard and archive backup repositories already configured on the selected backup appliance will be displayed. After you complete the wizard, Veeam Backup & Replication will automatically add these repositories to the backup infrastructure.

You can specify the following configuration settings for each repository whose restore points you want to use to recover backed-up data:

NOTE

The following procedure applies only to standard backup repositories. For archive backup repositories, there is no possibility to specify any configuration settings.

1. In the **Repositories** list, select the necessary repository and click **Edit**.
2. In the **Repository** window:
 - a. From the **Credentials** drop-down list, select access keys of an IAM user whose permissions will be used to access the repository. For more information on the required permissions, see [Plug-in Permissions](#).

For access keys of an IAM user to be displayed in the **Credentials** list, they must be added to the Cloud Credentials Manager as described in the Veeam Backup & Replication User Guide, section [Access Keys for AWS Users](#). If you have not added the necessary keys to the Cloud Credentials Manager beforehand, you can do it without closing the **Repository** window. To do that, click either the **Manage accounts** link or the **Add** button, and specify the access and secret key in the **Credentials** window.

NOTE

If you do not specify access keys of an IAM user for a standard backup repository, you will only be able to use the Veeam Backup & Replication console to perform [entire EC2 instance restore](#), [DB instance restore](#), [Aurora DB cluster restore](#) and [EFS file systems restore](#) from backups stored in this repository. Moreover, information on the repository displayed in the **Backup Infrastructure** view under the **External Repositories** node will not include statistics on the amount of storage space that is currently consumed by restore points created by Veeam Backup for AWS.

- b. From the **Use the following gateway server for the Internet access** drop-down list, select a gateway server that will be used to provide access to the repository.

For a gateway server to be displayed in the **Use the following gateway server for the Internet access** drop-down list, it must be added to the backup infrastructure. For more information on gateway servers, see [Solution Architecture](#).

- c. If encryption is enabled for the repository, the following scenarios may apply:
 - If data in the repository is encrypted using a password, select the **Use the following password for encrypted backups** check box. From the drop-down list, select the password that is used to encrypt data. Veeam Backup & Replication will use the specified password to decrypt backup files stored in this repository.

For a password to be displayed in the **Use the following password for encrypted backups** drop-down list, it must be added to the Password Manager as described in the Veeam Backup & Replication User Guide, section [Creating Passwords](#). If you have not added the necessary password beforehand, you can do it without closing the **Repository** window. To do that, click either the **Manage accounts** link or the **Add** button, and specify the password and hint in the **Password** window.

NOTE

If you do not specify a password for a standard backup repository with encryption enabled, you will have to decrypt data stored in this repository manually as described in section [Managing Backed-Up Data Using Console](#).


- If data in the repository is encrypted with a KMS key, Veeam Backup & Replication will show the used KMS key in the **Perform AWS encryption with the following KMS key** drop-down list but will not allow to change it.

For Veeam Backup & Replication to be able to decrypt data stored in the repository, the IAM user whose permissions will be used to access the repository must also have permissions to access KMS keys. For more information on the required permissions, see [Plug-in Permissions](#).

After you finish working with the wizard, all the added repositories will be displayed in the **Backup Infrastructure** view under the **External Repositories** node.

NOTE

If some of the repositories are already added to the backup infrastructure of another backup server, you will be prompted to claim the ownership of these repositories. To learn how to claim the ownership, see the Veeam Backup & Replication User Guide, section [Ownership](#).



Step 8. Track Progress

Veeam Backup & Replication will display the results of every step performed while connecting the backup appliance. At the **Apply** step of the wizard, wait for the process to complete and click **Next**.

NOTE

When adding an existing appliance to the backup infrastructure, Veeam Backup & Replication collects session results only for the past 48 hours, as well as information on all snapshots, backups and policies.

New Veeam Backup for AWS Appliance

Apply
Please wait while required operations are being performed. This may take a few minutes...


Message	Duration
✓ Backup appliance has been connected successfully	0:00:05
✓ Backup appliance configuration has been collected successfully	0:00:06
✓ External repositories connected	0:00:27
✓ External repository backup-dept05 has been connected succe...	0:00:24
✓ External repository backup-dept06 has been connected succe...	0:00:03

< Previous **Next >** Finish Cancel

Step 9. Finish Working with Wizard

At the **Summary** step of the wizard, review summary information and click **Finish**.

After the backup appliance is added to the backup infrastructure, you can configure its settings in the Veeam Backup for AWS Web UI as described in section [Configuring Veeam Backup for AWS](#). If you want Veeam Backup & Replication to open the Web UI of the added backup appliance immediately, click the **backup appliance console** link.



Editing Appliance Settings

For each backup appliance managed by the backup server, you can modify the settings configured while adding the appliance to the backup infrastructure:

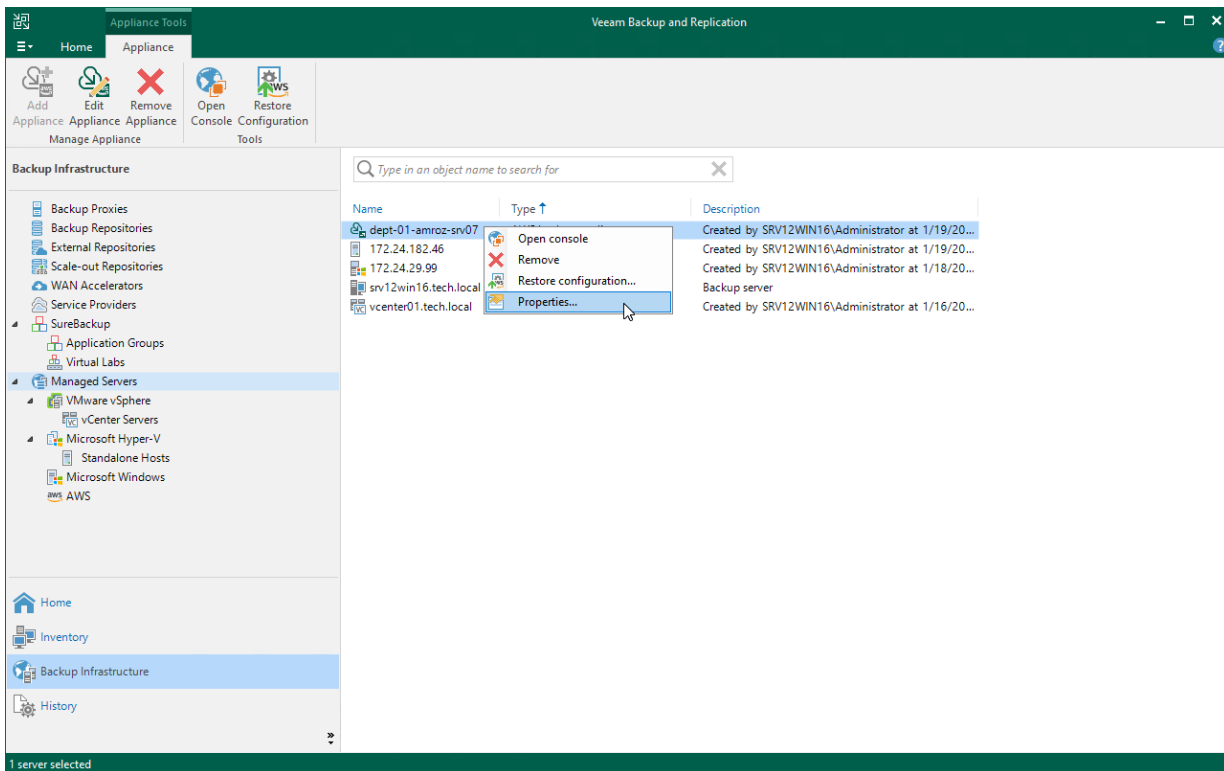
1. In the Veeam Backup & Replication console, open the **Backup Infrastructure** view.
2. Navigate to **Managed Servers**.
3. Select the necessary backup appliance and click **Edit Appliance** on the ribbon.
Alternatively, you can right-click the appliance and select **Properties**.
4. Complete the **Edit Veeam Backup for Veeam Backup for AWS Appliance** wizard:
 - a. To change the access keys of an IAM user that are used to connect to the backup appliance, follow the instructions provided in section [Connecting to Existing Veeam Backup for AWS Appliances](#) (step 1).
 - b. To provide a new description for the backup appliance, follow the instructions provided in section [Connecting to Existing Veeam Backup for AWS Appliances](#) (step 4).
 - c. To change the way Veeam Backup & Replication connects to the backup appliance, follow the instructions provided in section [Connecting to Existing Veeam Backup for AWS Appliances](#) (step 5).

IMPORTANT

You cannot change the way Veeam Backup & Replication connects to a backup appliance deployed in a private environment.

- d. To change the user whose credentials Veeam Backup & Replication uses to connect to the backup appliance, follow the instructions provided in section [Connecting to Existing Veeam Backup for AWS Appliances](#) (step 6).
- e. To accept the newly created self-signed certificate, follow the instructions provided in section [Connecting to Existing Appliances](#) (step 6).
- f. To edit settings of the backup appliance repositories added to the backup infrastructure, follow the instructions provided in section [Connecting to Existing Veeam Backup for AWS Appliances](#) (step 7).
- g. At the **Apply** step of the wizard, wait for the changes to be applied and click **Next**.

h. At the **Summary** step of the wizard, review summary information and click **Finish** to confirm the changes.



Rescanning Appliances

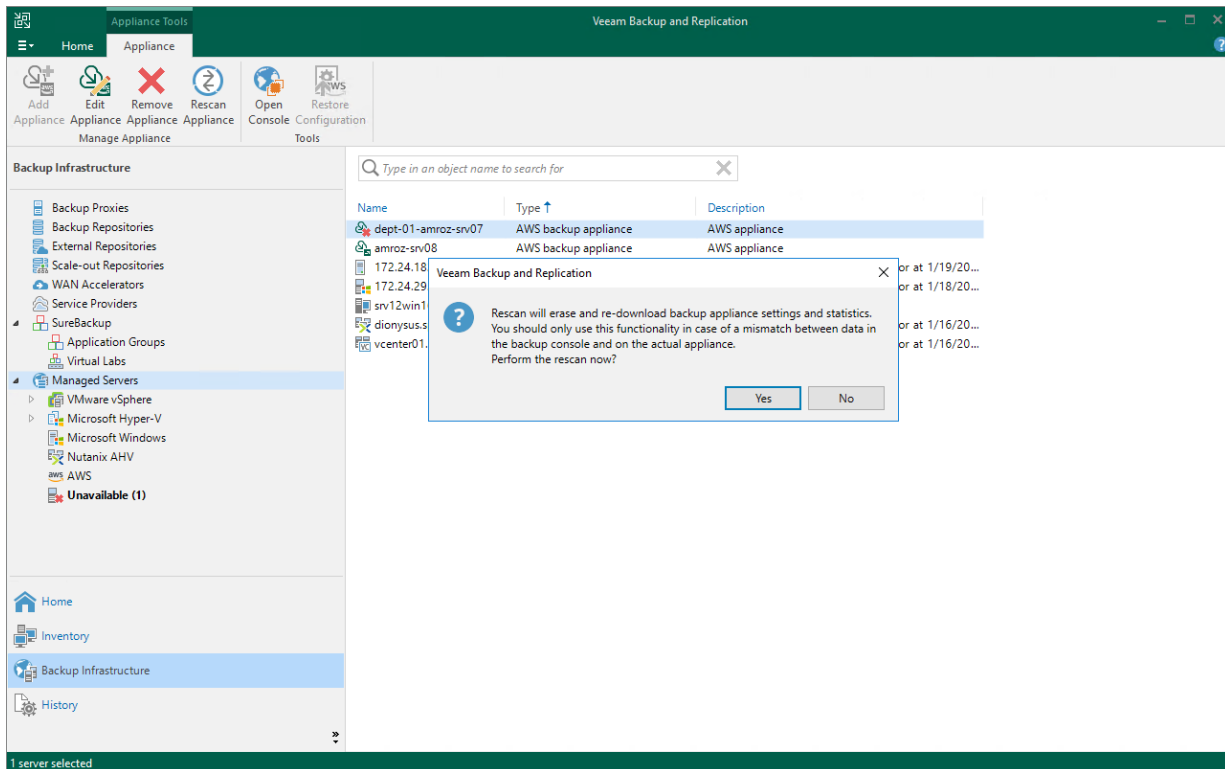
If a backup appliance becomes unavailable, for example, due to connectivity problems, you can rescan the appliance:

1. In the Veeam Backup & Replication console, open the **Backup Infrastructure** view.
2. Navigate to **Managed Servers**.
3. Select the necessary backup appliance and click **Rescan appliance** on the ribbon.
Alternatively, you can right-click the appliance and select **Rescan**.
4. In the opened window, click **Yes**.

Veeam Backup & Replication will remove all data collected from the appliance configuration database. Then, Veeam Backup & Replication will recollect session results for the past 48 hours, as well as information on all snapshots, backups and policies.

NOTE

The rescan operation cannot be performed for available backup appliances and appliances that require upgrade. To learn how to upgrade backup appliances, see [Upgrading Appliances Using Console](#).



Removing Appliances

AWS Plug-in for Veeam Backup & Replication allows you to permanently remove backup appliances from the backup infrastructure.

NOTE

After you remove a backup appliance, the following limitations will apply:

- Repositories for which you have not specified access keys of IAM users will be removed automatically from the backup infrastructure.
- Repositories for which you have specified access keys of IAM users will remain in the backup infrastructure. However, you will have to rescan the repositories to collect information on all newly created and recently deleted (both manually and by retention) restore points.
- You will not be able to manage backup policies created on the appliance.
- You will not be able to restore EC2 instances from snapshots.
- Restore to AWS from image-level backups will start working as described in the Veeam Backup & Replication User Guide, section [How Restore to Amazon EC2 Works](#).

Also, the restore process will start taking more time to complete causing data transfer costs to increase as Veeam Backup & Replication will not be able to use native AWS capabilities and will have to process more data.

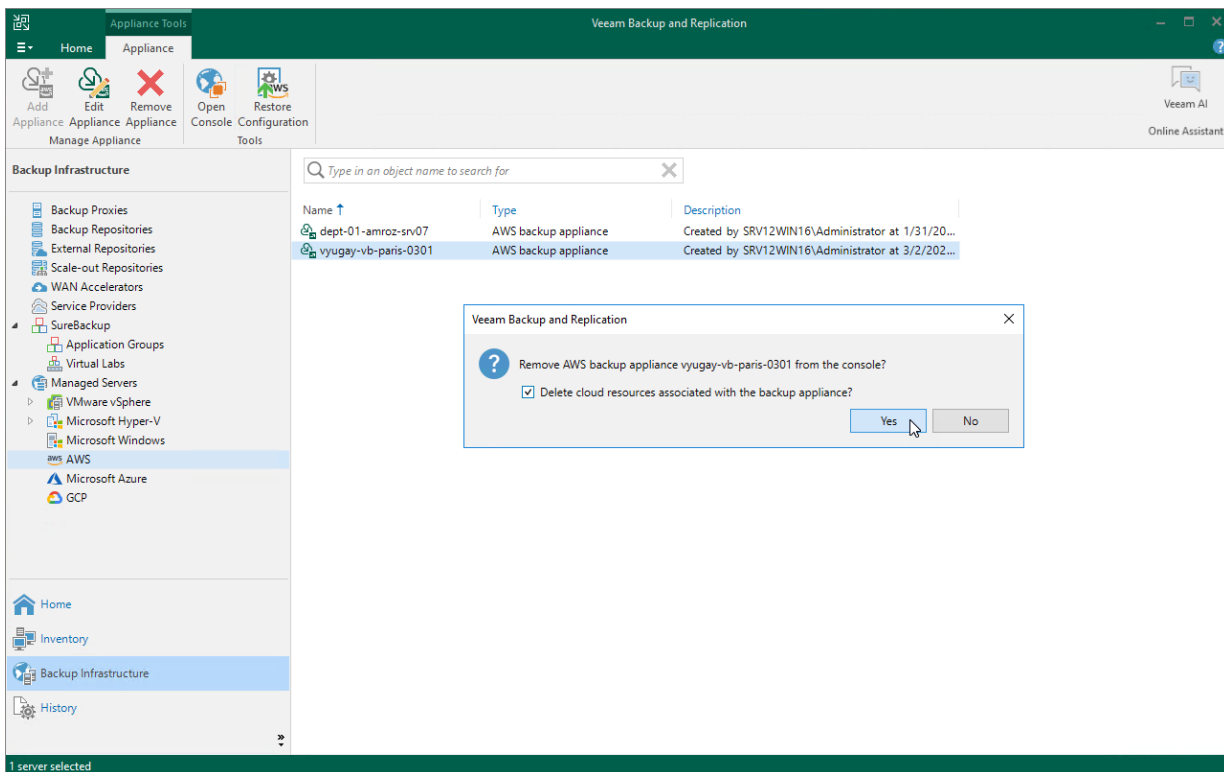
To remove a backup appliance, do the following:

1. In the Veeam Backup & Replication console, open the **Backup Infrastructure** view.
2. Navigate to **Managed Servers**.
3. Select the necessary backup appliance and click **Remove Appliance** on the ribbon.
Alternatively, you can right-click the appliance and select **Remove**.
4. In the **Veeam Backup & Replication** window, click **Yes** to acknowledge the operation.

TIP

If you want to remove an appliance from both the backup infrastructure and AWS, select the **Delete cloud resources associated with the backup appliance?** check box in the opened window. Veeam Backup for AWS will remove all resources associated with this appliance in AWS.

However, if an appliance has been deployed from the AWS Marketplace or is running Veeam Backup for AWS version 3.x (or earlier), to remove resources from AWS, you must follow the instructions provided in section [Appendix E. Uninstalling Backup Appliances Deployed from AWS Marketplace](#).



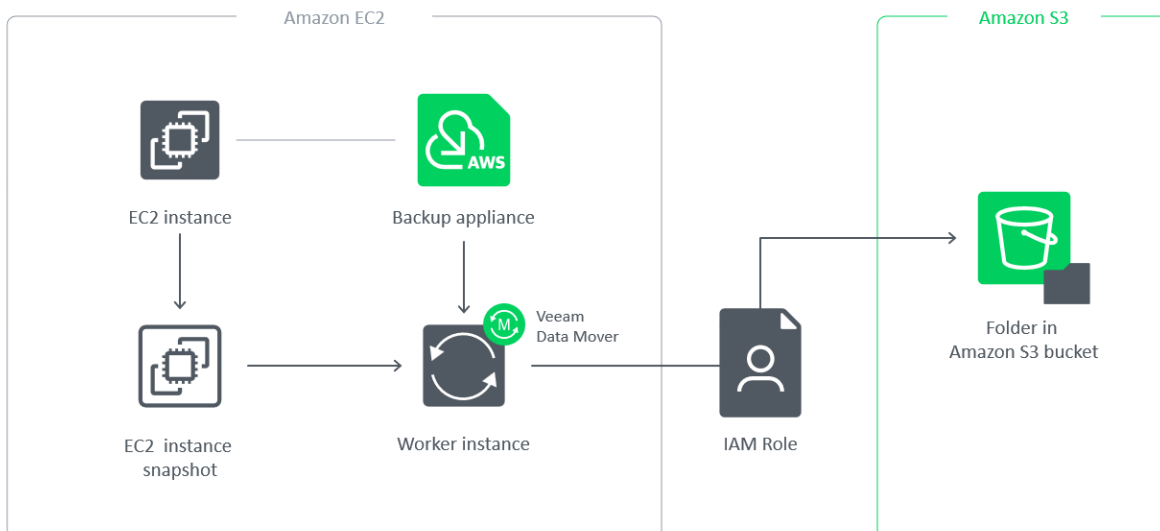
Managing Backup Repositories

Veeam Backup for AWS uses Amazon S3 buckets as target locations for EC2 and RDS image-level backups, additional copies of Amazon VPC backups, indexes of EFS file systems and Veeam Backup for AWS configuration backups. To store backups in Amazon S3 buckets, configure backup repositories. A repository is a specific folder created by Veeam Backup for AWS in an Amazon S3 bucket.

IMPORTANT

A backup repository must not be managed by multiple backup appliances simultaneously – retention sessions running on different backup appliances may corrupt backups stored in the repository, which may result in unpredictable data loss.

To communicate with the backup repository, Veeam Backup for AWS uses Veeam Data Mover – the service running on a worker instance that is responsible for data processing and transfer. When a backup policy addresses the backup repository, Veeam Data Mover establishes a connection with the repository enabling data transfer. To let Veeam Data Mover access the target Amazon S3 bucket, Veeam Backup for AWS uses permissions of an IAM role specified in [backup repository settings](#).



Adding Backup Repositories Using Console

Depending on whether you want to store backups in a high-performance, high-cost and short-term storage, or a secure, low-cost and long-term storage, you can configure repositories of the following storage classes:

- **Standard repositories**

Use repositories of the S3 Standard storage class to store data that you plan to access frequently. Backups stored in these repositories are shown under the **External Repository** node.

To store backups in a standard repository, first add it to the backup infrastructure and then enable image-level backups, VPC backup copy or EFS indexing in the backup policies settings. For more information, see sections [Creating EC2 Backup Policies](#), [Creating RDS Backup Policies](#), [Editing VPC Configuration Backup Policy](#) and [Creating EFS Backup Policies](#).

- **[Applies only to EC2 and RDS backups] Archive repositories**

Use repositories of the S3 Glacier Flexible Retrieval storage class to store data that you plan to access infrequently, and S3 Glacier Deep Archive storage class to store data that you plan to access once or twice a year. Backups stored in these repositories are shown under the **External Repository (Archive)** node.

To store backups in archive repository, first add it to the backup infrastructure and then enable backup archiving for any backup policy that will store backups in this repository. For more information, see [Creating EC2 Backup Policies](#).

To learn how backup archiving works, see [Enabling Backup Archiving](#).

IMPORTANT

Note that you can perform a limited scope of operations with archive repositories from the Veeam Backup & Replication console:

- You cannot edit and rescan archive repositories.
- You can only restore [entire EC2 instances](#) from backups stored in archive repositories. However, you can perform volume-level and file-level recovery operations from these backups using the Veeam Backup for AWS appliance Web UI. For more information, see sections [Performing Volume-Level Restore](#) or [Performing File-Level Recovery](#).
- You can restore specific databases of PostgreSQL DB instances using the Veeam Backup for AWS appliance Web UI only. For more information, see [Restoring RDS Databases](#).

For more information on Amazon S3 storage classes, see [AWS Documentation](#).

How to Add Backup Repositories

After you add a backup appliance to the backup infrastructure, you can configure repositories that will be used to store backups. To do that, use either of the following options:

- [Create new repositories](#).
- [Add existing repositories to the backup infrastructure](#) if you have already configured them on the backup appliance.

Creating New Repositories

To add a new repository, do the following:

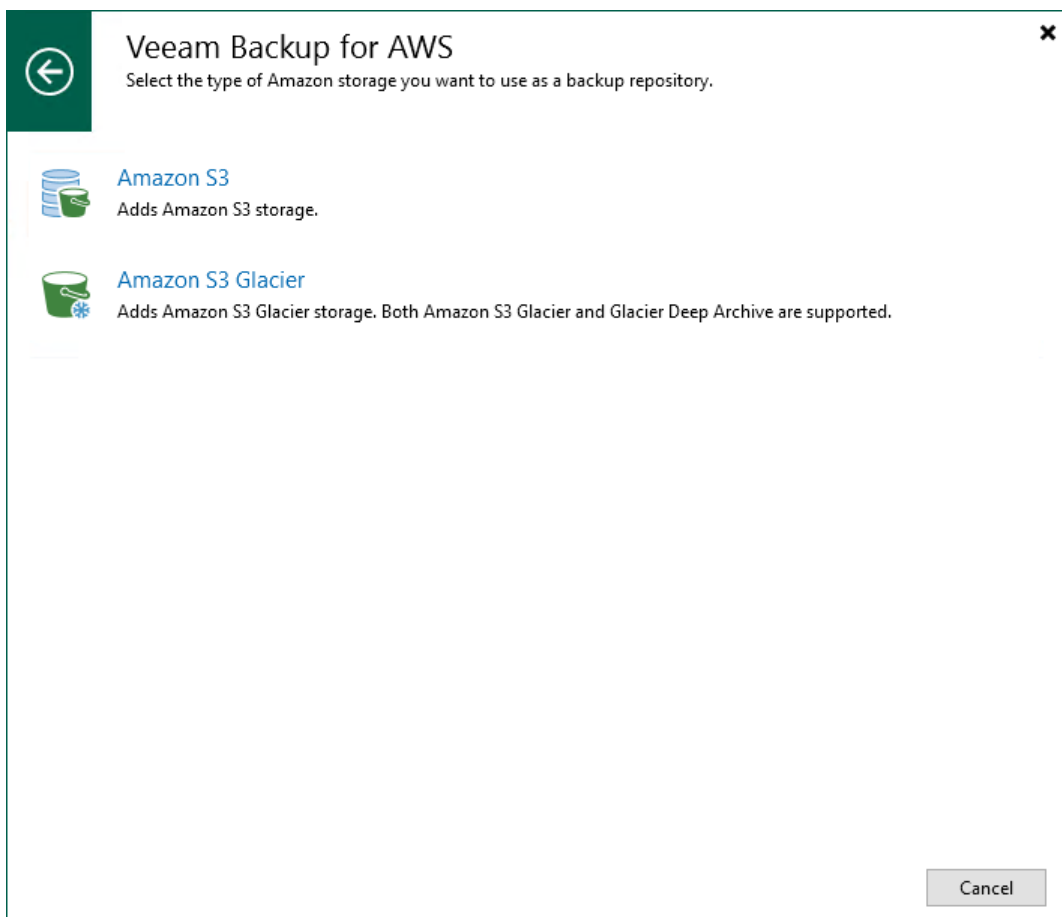
1. [Launch the Add External Repository wizard](#).

2. Specify an appliance, and provide a repository name and description.
3. Specify AWS account settings.
4. Specify an IAM role.
5. Specify an Amazon S3 bucket.
6. Enable data encryption.
7. Specify an S3 interface endpoint.
8. Wait for the repository to be added to the backup infrastructure.
9. Finish working with the wizard.

Step 1. Launch Add External Repository Wizard

To launch the **Add External Repository** wizard, do the following:

1. In the Veeam Backup & Replication console, open the **Backup Infrastructure** view.
2. Navigate to **External Repositories** and click **Add Repository** on the ribbon.
Alternatively, you can right-click the **External Repositories** node and select **Add**.
3. In the **Add External Repository** window:
 - a. [Applies only if you have several cloud plug-ins installed] Click **Veeam Backup for AWS**.
 - b. Choose whether you want to create a standard or an archive backup repository:
 - Select the **Amazon S3** option if you want to create a repository with the S3 Standard storage class assigned.
 - Select the **Amazon S3 Glacier** option if you want to create a repository with the S3 Glacier Flexible Retrieval or S3 Glacier Deep Archive storage class assigned.



Step 2. Specify Repository Details

At the **Veeam Backup for AWS** step of the wizard, do the following:

1. From the **Appliance** drop-down list, select a backup appliance that will manage the repository.

For an appliance to be displayed in the **Appliance** drop-down list, it must be added to the backup infrastructure as described in section [Deploying Backup Appliance](#) or [Connecting to Existing Appliances](#).

2. Use the **Repository name** and **Description** fields to enter a name for the new repository. The maximum length of the name is 125 characters; the following characters are not supported: \ / " ' [] : | < > + = ; , ? * @ & _ .

Veeam Backup & Replication will create a folder with the specified name in the storage bucket that you will specify at the [step 5](#) of the wizard. This folder will be used to store backed-up data.

Add External Repository [Close]

Veeam Backup for AWS
Specify the Veeam Backup for AWS appliance to create the S3 backup repository for.

Veeam Backup for AWS | Account | Bucket | Encryption | Settings | Apply | Summary

Appliance: aws-vy-12

Repository name: backup-repository

Description: standard backup repository

< Previous | **Next >** | Finish | Cancel

Step 3. Specify AWS Account Settings

At the **Account** step of the wizard, do the following:

1. From the **AWS account** drop-down list, select access keys of an IAM user whose permissions Veeam Backup & Replication will use to access the repository. For more information on the required permissions that must be assigned to the IAM user, see [Plug-In Permissions](#).

For access keys of an IAM user to be displayed in the **AWS account** drop-down list, they must be created in AWS and added to the Cloud Credentials Manager as described in the Veeam Backup & Replication User Guide, section [Access Keys for AWS Users](#). If you have not added the keys to the Cloud Credentials Manager beforehand, you can do it without closing the wizard. To do that, click either the **Manage cloud accounts** link or the **Add** button, and specify the access key and secret key in the **Credentials** window.

2. From the **AWS region** drop-downlist, specify whether the repository will be located in an AWS Global or AWS GovCloud (US) region.

IMPORTANT

To check the availability of the region, Veeam Backup & Replication by default establishes a temporary test connection with the US East (N. Virginia) region using endpoints of the [AWS Security Token Service \(STS\)](#) and [Amazon Elastic Compute Cloud \(EC2\)](#) AWS services. That is why the backup server must have access to this AWS Region. If you want to change the default region for a test connection, open a [support case](#).

3. [Applies only if you choose to create a standard backup repository] From the **Gateway server** drop-down list, select a gateway server that will be used to access the repository.

For a server to be displayed in the **Gateway server** list, it must be added to the backup infrastructure. For more information on gateway servers, see [Solution Architecture](#).

The screenshot shows the 'Add External Repository' wizard window, specifically the 'Account' step. The window title is 'Add External Repository' with a close button in the top right corner. On the left side, there is a navigation pane with the following items: 'Veeam Backup for AWS', 'Account' (which is selected and highlighted), 'IAM Role', 'Bucket', 'Encryption', 'Settings', 'Apply', and 'Summary'. The main content area is titled 'Account' and contains the following fields and instructions:

- AWS account:** A dropdown menu showing 'XXXXXXXXXXXXXXXXXX (VY, last edited: 1 day ago)' with a key icon on the left and an 'Add...' button on the right. Below the dropdown is a link that says 'Manage cloud accounts'.
- AWS region:** A dropdown menu showing 'Global'. Below it is the instruction: 'Select an AWS region based on your regulatory and compliance requirements.'
- Gateway server:** A dropdown menu showing 'xxx-xxxxxxx-VBR.xxxxxxxx (Backup server)'. Below it is the instruction: 'Select a gateway server to proxy access to Amazon S3 bucket with backup files. The server will store a cache of backup metadata for enhanced performance.'

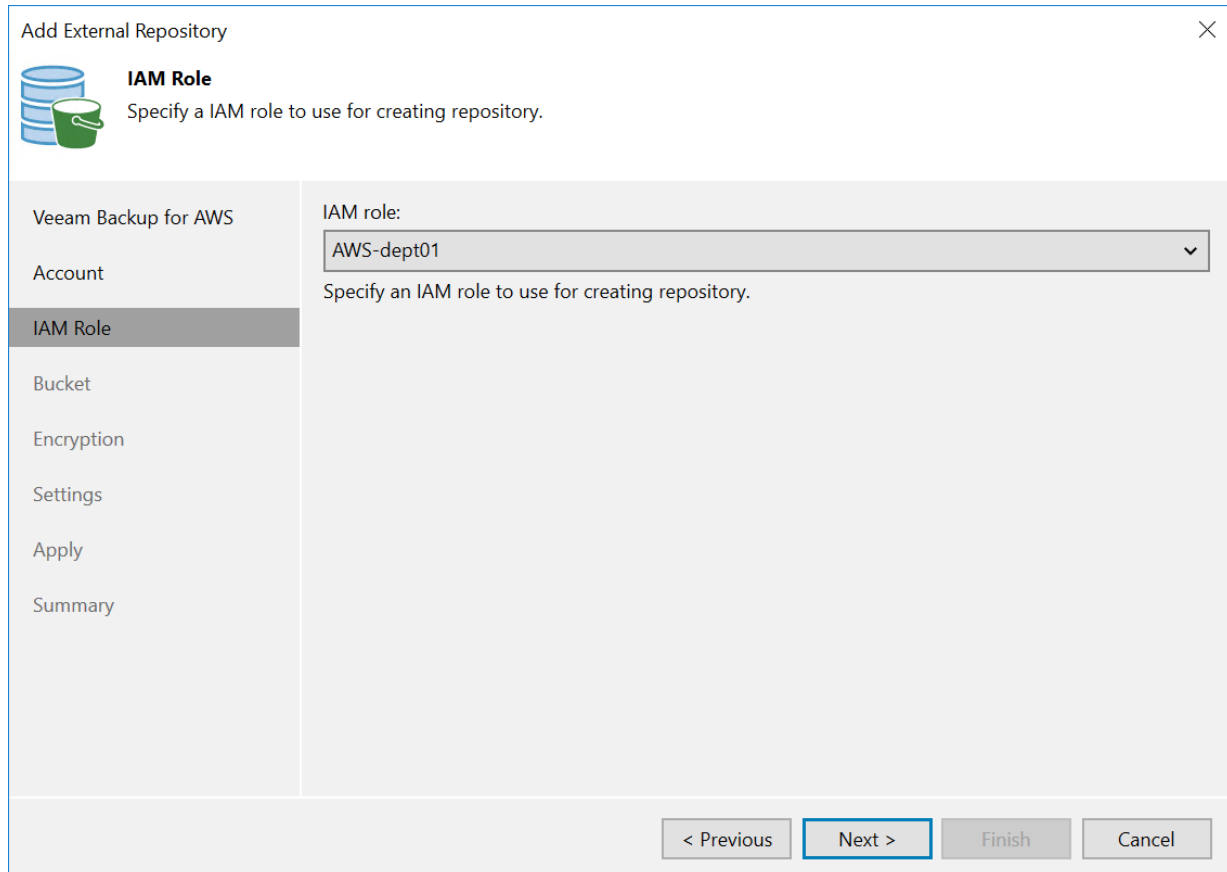
At the bottom of the window, there are four buttons: '< Previous', 'Next >' (which is highlighted with a blue border), 'Finish', and 'Cancel'.

Step 4. Specify IAM Role

[This step applies only if you have added to the backup appliance multiple IAM roles belonging to the same AWS account]

At the **IAM Identity** step of the wizard, select an IAM role whose permissions will be used to create the repository and to access the target Amazon S3 bucket. For more information on the required permissions that must be assigned to the IAM role, see [Restore IAM Permissions](#).

For an IAM role to be displayed in the **IAM role** drop-down list, it must be added to the backup appliance as described in section [Adding IAM Roles](#), and must belong to the same AWS account to which the IAM user specified at [step 3](#) of the wizard belongs.



The screenshot shows a wizard window titled "Add External Repository" with a close button (X) in the top right corner. Below the title is an icon of a database and a bucket, followed by the heading "IAM Role" and the instruction "Specify a IAM role to use for creating repository." A left-hand navigation pane lists the following steps: "Veeam Backup for AWS", "Account", "IAM Role" (which is highlighted), "Bucket", "Encryption", "Settings", "Apply", and "Summary". The main area of the wizard contains the label "IAM role:" above a dropdown menu that currently displays "AWS-dept01". Below the dropdown is the instruction "Specify an IAM role to use for creating repository." At the bottom of the window, there are four buttons: "< Previous", "Next >" (which is highlighted with a blue border), "Finish", and "Cancel".

Step 5. Specify Amazon S3 Bucket

At the **Bucket** step of the wizard, do the following:

1. From the **Data center** drop-down list, select an AWS Region where the repository will be located.
2. Choose whether you want to use an existing bucket or to create a new one as the target location for image-level backups of EC2 instances and RDS resources, additional copies of Amazon VPC backups and indexes of EFS file systems:

- To specify an existing bucket, in the **Bucket** field, enter the name of an Amazon S3 bucket where the repository will be created.

Alternatively, click **Browse** and select the necessary bucket in the **Select Bucket** window. For a bucket to be displayed in the **Bucket** list, it must be created in AWS as described in [AWS Documentation](#).

IMPORTANT

Consider the following:

- If you have any S3 Lifecycle configuration associated with the selected Amazon S3 bucket, it is recommended that you limit the scope of lifecycle rules applied to backup files created by the backup appliance. Otherwise, the backup files may be unexpectedly deleted or transitioned to another storage class, and the backup appliance will not be able to access the files. For more information on managing S3 Lifecycle configurations, see [AWS Documentation](#).
- If you plan to enable immutability settings for the created repository, S3 Versioning and Object Lock must be enabled for the specified Amazon S3 bucket, and no default retention period must be configured for the bucket. For more information on Amazon S3 immutability features, see [AWS Documentation](#).

- To create a new bucket, click **Browse**. In the **Select Bucket** window, click **New Bucket** and enter a name for the bucket. Veeam Backup & Replication will automatically create a bucket in the specified AWS Region. Note that the bucket name must meet the requirements described in [AWS Documentation](#).

If you want to enable immutability settings for the bucket, select the **Enable immutability** check box in the **New Bucket** window. Veeam Backup & Replication will automatically create a bucket with the S3 Versioning and Object Lock options enabled in the specified AWS Region. For more information on Amazon S3 immutability features, see [AWS Documentation](#).

3. [Applies only if you have selected or created a bucket with immutability settings enabled] If you want to protect backups stored in the repository from being lost as a result of malware, ransomware or any other malicious actions, you can enable immutability at the repository level. To do that, select the **Make backups immutable for the entire duration of their retention policy** check box. For more information on immutability, see [Immutability](#).

IMPORTANT


Consider the following:

- You cannot create standard backup repositories with the disabled immutability settings in Amazon S3 buckets with the S3 Versioning and Object Lock options enabled.
 - You cannot edit the configured immutability settings after the repository is created.
4. [Applies only if you choose to create an archive backup repository] When you create an archive backup repository, backups are stored in a secure, durable and low-cost S3 Glacier Flexible Retrieval storage class by default. To store backups in the lowest-cost S3 Glacier Deep Archive storage class that you plan to access once or twice a year, select the **Use the Deep Archive storage class** check box. Note that after the repository is created, you will be unable to change the selected storage class.

NOTE

When you create an archive backup repository, Veeam Backup for AWS does not create any S3 Glacier vaults in Amazon S3. Instead, it assigns the selected storage class (S3 Glacier Flexible Retrieval or S3 Glacier Deep Archive) to backups stored in the repository. That is why the archived backups remain in Amazon S3 and cannot be accessed directly through the Amazon S3 Glacier service.

Add External Repository ✕

 **Bucket**
Specify Amazon S3 bucket to connect to.

Veeam Backup for AWS	Data center:
Account	EU (Frankfurt) (eu-central-1) ▼
IAM Role	Bucket:
Bucket	<input type="text" value="XXXXXXXXXXXXXXXXXXXX"/> Browse...
Encryption	<input checked="" type="checkbox"/> Make backups immutable for the entire duration of their retention policy Protects backups from modification or deletion by ransomware, hackers or malicious insiders using native object storage capabilities.
Settings	
Apply	
Summary	

< Previous Next > Finish Cancel

Step 6. Enable Data Encryption

At the **Encryption** step of the wizard, choose whether you want to encrypt backups stored in the created repository.

IMPORTANT

After you create a repository with encryption enabled, you can no longer disable encryption for this repository. However, you will be able to change the encryption settings as described in section [Editing Backup Repository Settings](#).

If you select the **Enable backup file encryption** check box, also choose whether you want to use a password or an AWS Key Management Service (KMS) key to encrypt the backed-up data:

- To encrypt data using an AWS KMS key, select the **Perform AWS encryption with the following KMS key** option and choose the necessary KMS key from the drop-down list.

For a KMS key to be displayed in the list of available encryption keys, it must be [created in the AWS Region](#) where the selected Amazon S3 bucket is located, and the IAM role specified to access the bucket must have permissions to access the key. For more information on permissions required for the IAM role, see [Repository IAM Permissions](#).

NOTE

For Veeam Backup & Replication to be able to decrypt data stored in the repository, the IAM user specified at [step 3](#) of the wizard must have permissions to access KMS keys. For more information on the required permissions, see [Plug-in Permissions](#).

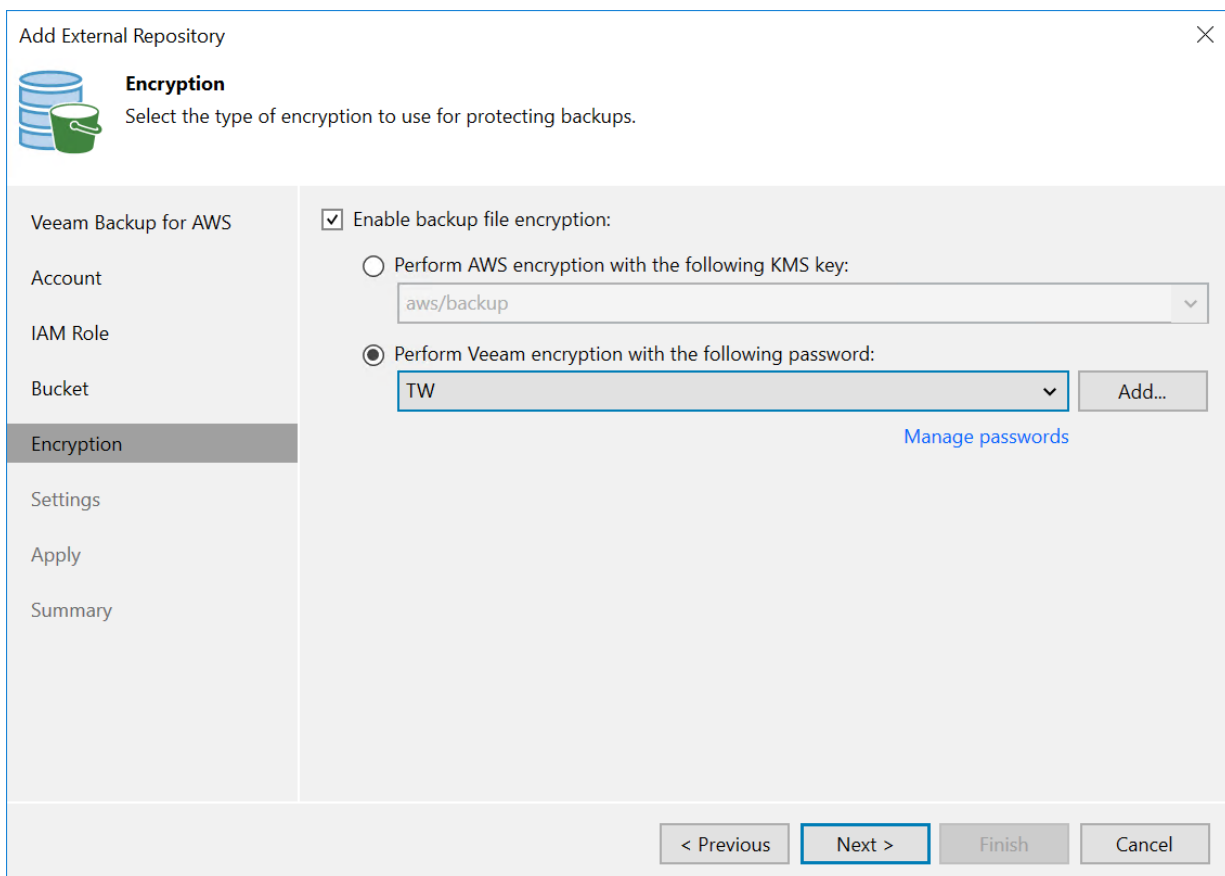
- To encrypt data using a password, select the **Perform Veeam encryption with the following password** option and choose the necessary password from the drop-down list.

For a password to be displayed in the list of available passwords, it must be added to Veeam Backup & Replication as described in the Veeam Backup & Replication User Guide, section [Creating Passwords](#). If you have not added the necessary password beforehand, you can do it without closing the wizard. To do that, click either the **Manage passwords** link or the **Add** button, and specify the password and hint in the **Password** window.

IMPORTANT

If you select the **Perform AWS encryption with the following KMS key** option, consider the following:

- Only symmetric KMS keys are supported.
- Do not disable the KMS key specified in the repository settings. Otherwise, the backup appliance will not be able to encrypt data, and backup policies that use the repository as the backup target will fail to complete successfully.
- Do not delete the KMS key specified in the repository settings. Otherwise, the backup appliance will not be able to decrypt data stored in the repository.



The screenshot shows the 'Add External Repository' dialog box with the 'Encryption' tab selected. The dialog has a sidebar on the left with the following items: Veeam Backup for AWS, Account, IAM Role, Bucket, Encryption (highlighted), Settings, Apply, and Summary. The main area contains the following options:

- Enable backup file encryption:
 - Perform AWS encryption with the following KMS key:
aws/backup
 - Perform Veeam encryption with the following password:
TW

There is an 'Add...' button next to the password field and a 'Manage passwords' link below it. At the bottom of the dialog are four buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'.

Step 7. Specify VPC Interface Endpoint

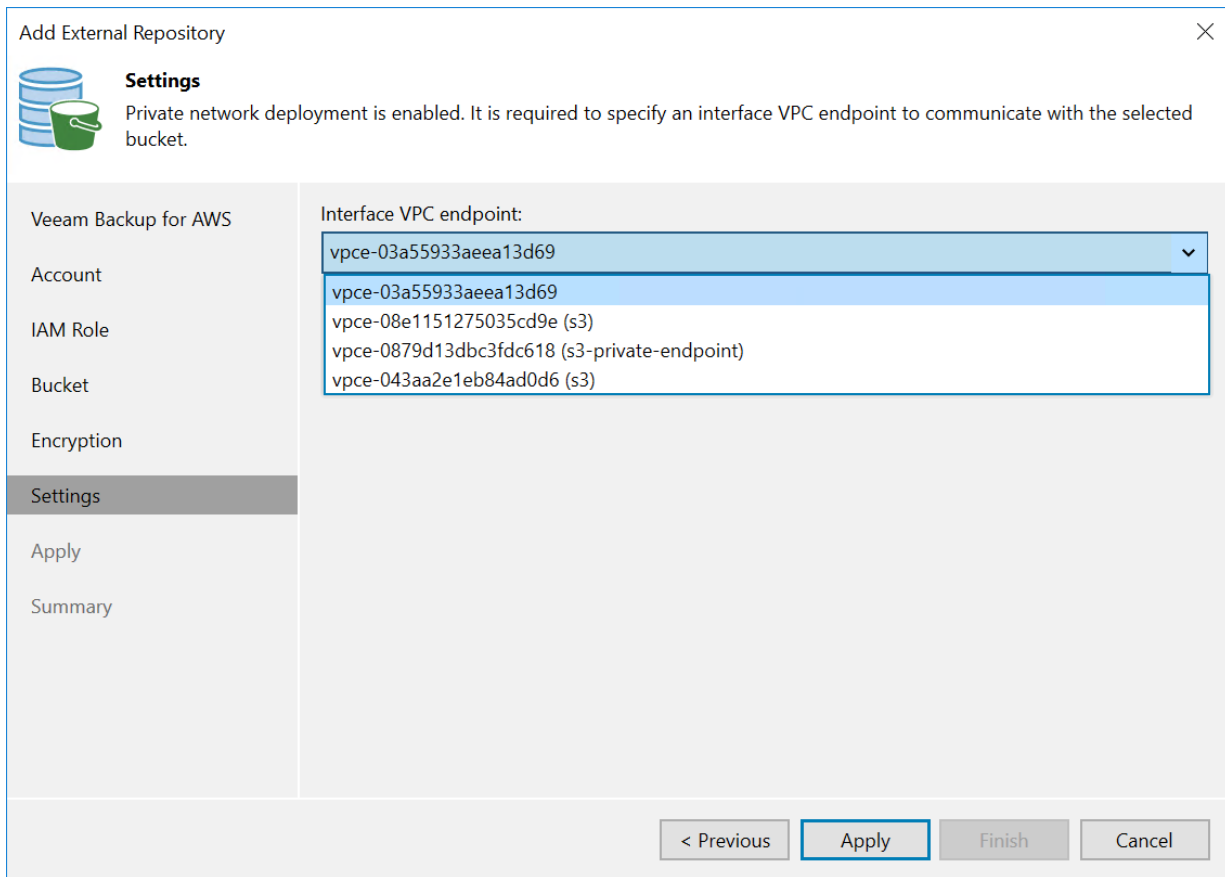
[This step applies only if you have enabled the [private network deployment](#) functionality]

At the **Settings** step of the wizard, specify an S3 interface endpoint that will be used to communicate with the Amazon S3 service.

For an S3 interface endpoint to be displayed in the **Interface VPC endpoint** list, it must be created in the Amazon VPC console for all subnets to which the worker instances will be connected, as described in section [Configuring Private Networks](#).

IMPORTANT

S3 gateway endpoints are not supported when using the private network deployment functionality.



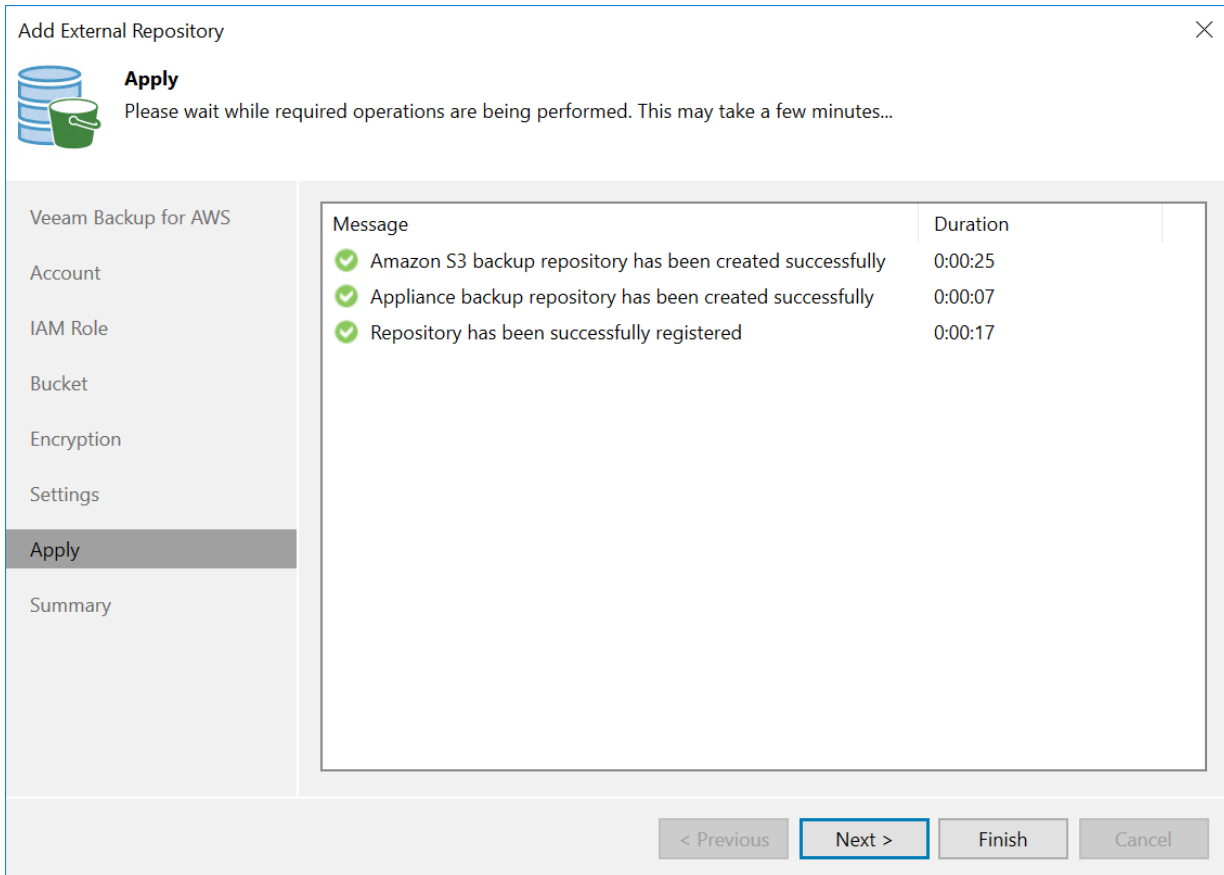
The screenshot shows the 'Add External Repository' wizard in the 'Settings' step. A message states: 'Private network deployment is enabled. It is required to specify an interface VPC endpoint to communicate with the selected bucket.' The 'Interface VPC endpoint' dropdown menu is open, displaying a list of endpoints:

- vpce-03a55933aeea13d69
- vpce-03a55933aeea13d69
- vpce-08e1151275035cd9e (s3)
- vpce-0879d13dbc3fdc618 (s3-private-endpoint)
- vpce-043aa2e1eb84ad0d6 (s3)


The 'Apply' button is highlighted in blue.

Step 8. Track Progress

Veeam Backup & Replication will display the results of every step performed while creating the repository. At the **Apply** step of the wizard, wait for the process to complete and click **Next**.



Add External Repository ✕


 **Apply**
Please wait while required operations are being performed. This may take a few minutes...

Message	Duration
✓ Amazon S3 backup repository has been created successfully	0:00:25
✓ Appliance backup repository has been created successfully	0:00:07
✓ Repository has been successfully registered	0:00:17

< Previous **Next >** Finish Cancel

Step 9. Finish Working with Wizard

At the **Summary** step of the wizard, review summary information and click **Finish**.



Connecting to Existing Repositories

When you connect to a backup appliance, all repositories that have already been configured on the appliance are automatically added to the backup infrastructure.

If an existing repository is not displayed under the **External Repositories** node or if you have recently configured a new repository on the backup appliance that is already connected to the backup server, do the following:

1. In the Veeam Backup & Replication console, open the **Backup Infrastructure** view.
2. Navigate to **Managed Servers**.
3. Select a backup appliance that manages the necessary repository and click **Edit Appliance** on the ribbon. Alternatively, you can right-click the backup appliance and select **Properties**.
4. In the **Edit Veeam Backup for AWS Appliance** wizard, do the following:
 - a. Navigate to the **Repositories** step of the wizard and complete the step as described in section [Connecting to Existing Veeam Backup for AWS Appliances](#) (step 7).
 - b. Complete the **Edit Veeam Backup for AWS Appliance** wizard as described in section [Connecting to Existing Veeam Backup for AWS Appliances](#) (steps 8-9).

Open the **Backup Infrastructure** view to verify that the repository is displayed under the **External Repositories** node.

NOTE

If you do not specify access keys of an IAM user for a standard backup repository, you will only be able to use the Veeam Backup & Replication console to perform [entire EC2 instance restore](#) from backups stored in this repository. Moreover, information on the repository displayed in the **Backup Infrastructure** view under the **External Repositories** node will not include statistics on the amount of storage space that is currently consumed by restore points created by Veeam Backup for AWS.

Adding Backup Repositories Using Web UI

IMPORTANT

If your backup appliance is managed by a Veeam Backup & Replication server and you add a new backup repository using the Veeam Backup for AWS Web UI, Veeam Backup for AWS will not propagate these settings to the Veeam Backup & Replication server automatically. To discover new backup repositories created on the backup appliance, follow the instructions provided in section [Connecting to Existing Repositories](#).

To add a backup repository, do the following:

1. [Check prerequisites and limitations](#).
2. [Launch the Add Repository wizard](#).
3. [Specify a backup repository name and description](#).
4. [Configure backup repository settings](#).
5. [Enable data encryption for the backup repository](#).
6. [Specify an S3 interface endpoint](#).
7. [Finish working with the wizard](#).

Limitations and Considerations

When adding a backup repository to Veeam Backup for AWS, keep in mind the following limitations and considerations.

Amazon S3 Bucket

Before you add a backup repository, check the following prerequisites:

- An Amazon S3 bucket must be created in AWS beforehand as described in [AWS Documentation](#).
- If you have any S3 Lifecycle configuration associated with the selected Amazon S3 bucket, it is recommended that you limit the scope of lifecycle rules applied to Amazon S3 objects in the bucket so that no rules are applied to backup files created by Veeam Backup for AWS. Otherwise, the files may be unexpectedly deleted or transitioned to another storage class, and Veeam Backup for AWS may not be able to access the files. For more information on managing S3 Lifecycle configurations, see [AWS Documentation](#).

IMPORTANT

To maintain the security of your data, you should never use a public S3 bucket as a repository for Veeam Backup for AWS. For more information on creating buckets, see [AWS Documentation](#).

Repository Folder

If you plan to select an existing folder for storing backup files, consider the following:

- The folder must not be specified as a backup repository on multiple backup appliances simultaneously. Retention sessions running on different backup appliances may corrupt backup files stored in the folder, which may result in unpredictable data loss.
- The created backup repository will have the storage class that has been specified when creating the folder. You cannot change the storage class for the repository.
- If encryption at the repository level is enabled for the selected folder, it will be required to provide a password or an encryption key for this folder at [step 4](#) of the wizard.
- If the selected folder already contains backups created by the Veeam backup service, Veeam Backup for AWS will import the backed-up data to the configuration database. You can then use this data to perform all disaster recovery operations described in section [Performing Restore](#).

By default, Veeam Backup for AWS applies retention settings saved in the backup metadata to the imported backups. However, if the selected folder contains backups of resources that you plan to protect by a backup policy with the created repository specified as a backup target, Veeam Backup for AWS will rewrite the saved retention settings and will apply to the imported backups new retention settings configured for that backup policy.

Immutability

If you plan to add a repository with immutability enabled, keep in mind the following limitations:

- S3 Object Lock and S3 Versioning must be enabled for an Amazon S3 bucket in which the repository will be located. The default retention period must not be configured in the Object Lock settings. For more information on the S3 Versioning and S3 Object Lock features, see [AWS Documentation](#).
- Veeam Backup for AWS does not support enabling immutability settings in the AWS Management Console for buckets that are already used as target locations for image-level backups. However, you can enable immutability settings in the AWS Management Console for existing buckets when creating new backup repositories, and then specify these repositories in the settings of new backup policies.
- An IAM role that you plan to specify to create the repository and further to access the repository when performing data protection and recovery tasks must be assigned permissions to collect immutability settings of Amazon S3 buckets and to create immutable backups. For more information on the required permissions, see [Repository IAM Role Permissions](#).
- You cannot store indexes of EFS file systems and backups of the appliance configuration database in the repository with immutability enabled.
- You cannot remove immutable data manually using the Veeam Backup for AWS Web UI, as described in sections [Removing EC2 Backups and Snapshots](#), [Removing RDS Backups and Snapshots](#) and [Removing VPC Configuration Backups](#).
- You can neither remove immutable data from AWS using any cloud service provider tools nor request the technical support department to do it for you. Since Veeam Backup for AWS uses S3 Object Lock in the compliance mode, none of the protected objects can be overwritten or deleted by any user, including the root user in your AWS account. For more information on S3 Object Lock retention modes, see [AWS Documentation](#).

Encryption

If you plan to enable encryption for a backup repository, consider the following:

- After you create a repository with encryption enabled, you will not be able to disable encryption for this repository. However, you will still be able to change the encryption settings as described in section [Editing Backup Repository Settings](#).
- If you enable encryption for a repository where EC2 image-level backups are stored when editing the repository, this will affect the creation of an existing backup chain – the next sequence of backups will be a full backup instead of an incremental backup. After creating the full backup, Veeam Backup for AWS will continue to copy only those data blocks that have changed since the previous backup session.
- If you choose to encrypt data using an AWS KMS key, keep in mind that:
 - AWS managed keys cannot be used to encrypt data stored in repositories due to [AWS limitations](#).
 - Only symmetric KMS keys are supported.
 - Do not disable KMS keys specified in the repository settings. Otherwise, Veeam Backup for AWS will not be able to encrypt data, and backup policies that store backups in these repositories will fail to complete successfully.
 - Do not delete KMS keys specified in the repository settings. Otherwise, Veeam Backup for AWS will not be able to decrypt data stored in these repositories.


If a KMS key is scheduled for deletion, it will acquire the Pending deletion state. In this case, Veeam Backup for AWS will raise the warning notifying that you must either change the encryption settings for the backup repository in Veeam Backup for AWS or cancel the key deletion during the following 7 days.

For more information on managing AWS KMS keys, see [AWS Documentation](#).

Step 1. Launch Add Repository Wizard

To launch the **Add Repository** wizard, do the following:

1. Switch to the **Configuration** page.
2. Navigate to **Repositories**.
3. Click **Add**.



Step 2. Specify Repository Name and Description

At the **Info** step of the wizard, specify a name and description for the new backup repository. The name must be unique in Veeam Backup for AWS; the maximum length of the name is 125 characters; the maximum length of the description is 1024 characters.

The screenshot shows the 'Add Repository' wizard in Veeam Backup for AWS. The interface has a dark green header with the Veeam logo, 'Veeam Backup for AWS', server time 'Nov 30, 2023 5:05 PM', and user 'administrator Portal Administrator'. A left sidebar contains navigation tabs: 'Info' (selected), 'Bucket', 'Settings', and 'Summary'. The main area is titled 'Specify repository name and description' and includes the instruction 'Enter a name and description for the repository.' There are two input fields: 'Name:' with the text 'Immutable Repository' and 'Description:' with the text 'Repository for storing immutable data'. At the bottom, there are 'Next' and 'Cancel' buttons.

Step 3. Configure Repository Settings

At the **Bucket** step of the wizard, specify an IAM role that will be used to access the created repository, choose an Amazon S3 bucket in which the repository will be created, and review immutability settings for the repository.

Specifying IAM Role

In the **IAM role** section, specify an IAM role whose permissions Veeam Backup for AWS will use to create the new repository in the target Amazon S3 bucket and further to access the repository when performing data protection and recovery tasks. It is recommended that you check whether the selected IAM role has all the permissions required to perform the operation. To do that, click **Check Permissions** and follow the instructions provided in section [Checking IAM Role Permissions](#). For more information on permissions required for the IAM role, see [Repository IAM Role Permissions](#).

For an IAM role to be displayed in the **IAM role** list, it must be added to Veeam Backup for AWS with the *Repository role* selected as described in section [Adding IAM Roles](#). If you have not added the necessary IAM role to Veeam Backup for AWS beforehand, you can do it without closing the **Add Repository** wizard. To add an IAM role, click **Add** and complete the [Add IAM Role](#) wizard.

Choosing Repository Location

In the **Location** section, do the following:

1. Specify an Amazon S3 bucket where you want to store backups.
 - a. Click the **Choose bucket** link.
 - b. In the **Choose bucket** window, select the Amazon S3 bucket that will be used as a target location for backups, and click **Apply**.

For an Amazon S3 bucket to be displayed in the **Bucket** list, it must be created within an AWS account to which the specified IAM role belongs. To learn how to create Amazon S3 buckets, see [AWS Documentation](#).

It may take some time for Veeam Backup for AWS to retrieve information about existing Amazon S3 buckets from AWS.

2. Choose whether you want to use an existing folder inside the selected Amazon S3 bucket or to create a new one to group backup files stored in the bucket.
 - To use an existing folder, select the **Use existing folder** option and click **Choose folder**. In the **Choose folder** window, select the necessary folder and click **Apply**. Keep in mind [limitations and considerations](#) for existing repository folders.

For a folder to be displayed in the **Folder** list, it must have been created by any backup appliance as a repository (either existing or already removed from the backup infrastructure) in the selected Amazon S3 bucket.
 - To create a new folder, select the **Create new folder** option and specify a name for the new folder. The maximum length of the name is 125 characters; the slash (/) character is not supported.
3. [This step applies only if you have selected the **Create new folder** option] From the **Storage class** drop-down list, select a storage class for the backup repository:
 - To store backups in the S3 Standard storage class – a high-availability and high-performance storage that you plan to access frequently, select *S3 Standard*.

- To store backups in the S3 Glacier Flexible Retrieval storage class – a secure, durable and low-cost archive storage that you plan to access infrequently, select *S3 Glacier Flexible Retrieval*.
- To store backups in the S3 Glacier Deep Archive storage class – the lowest-cost archive storage that you plan to access once or twice a year, select *S3 Glacier Deep Archive*.

For more information on Amazon S3 storage classes, see [AWS Documentation](#).

NOTE

When you select the **S3 Glacier Flexible Retrieval** or **S3 Glacier Deep Archive** option for a backup repository, Veeam Backup for AWS does not create any S3 Glacier vaults in your AWS environment – it assigns the selected storage class to backups stored in the repository. That is why the archived backups remain in Amazon S3 and cannot be accessed directly through the Amazon S3 Glacier service.

Reviewing Immutability Settings

Veeam Backup for AWS allows you to protect backups stored in the repository from being lost as a result of malware, ransomware or any other malicious actions. To do that, you can create repositories with [immutability](#) enabled. For more information on requirements and limitations, see [Limitations and Considerations](#).

NOTE

For security reasons, it is recommended that you store immutable backup files in a dedicated AWS account. To do that, specify an IAM role that belongs to the necessary account as described in section [Specifying IAM Role](#), and then choose an Amazon S3 bucket that meets the immutability requirements.

As soon as you choose an Amazon S3 bucket, Veeam Backup for AWS verifies the immutability settings configured at the bucket level, and displays the following information in the **Immutability** section:

- If both S3 Versioning and S3 Object Lock are enabled for the specified bucket, and the default retention period is not configured in the Object Lock settings, Veeam Backup for AWS automatically selects the **Backups stored in this repository will be immutable** check box. In this case, the repository will be created with immutability enabled. For more information, see [Immutability](#).
- If S3 Object Lock is disabled and S3 Versioning is disabled (or suspended) for the specified bucket, Veeam Backup for AWS automatically clears the **Backups stored in this repository will be immutable** check box. In this case, the repository will be created with immutability disabled.
- If none of the cases apply, Veeam Backup for AWS raises an error notifying that the bucket cannot be used to create the repository. In this case, either choose another Amazon S3 bucket or reconfigure the bucket settings in the AWS Management Console.

IMPORTANT

It is recommended that S3 Object Lock and S3 Versioning are either both enabled or both disabled for the selected bucket. Otherwise, enabling S3 Versioning alone will significantly increase the amount of space consumed by backups stored in the bucket.

For more information on the S3 Versioning and S3 Object Lock features, see [AWS Documentation](#).

The screenshot shows the 'Add Repository' wizard in Veeam Backup for AWS, specifically the 'Bucket' step. The interface includes a top navigation bar with the Veeam logo, server time (Nov 30, 2023 5:06 PM), and user information (administrator, Portal Administrator). The main content area is titled 'Configure general settings' and contains the following sections:

- Info:** 'Configure general settings' - Select an IAM role to be used to access the repository and an Amazon S3 bucket where backup files will be stored.
- Bucket:** 'IAM role' - A dropdown menu is set to 'Repository role (role to access repository)'. There are '+ Add' and 'Check Permissions' buttons.
- Location:** 'Bucket:' is set to 'am-bckt-immutable'. Below this, there are two options: 'Use existing folder:' (unchecked) and 'Create new folder:' (checked) with the text 'dept_01_immutable' in the input field.
- Storage class:** A dropdown menu is set to 'S3 Standard'.
- Immutability settings:** A checkbox 'Backups stored in this repository will be immutable' is checked. Below it, an information icon states: 'Immutability settings cannot be changed. Immutability will be enabled for the entire duration of the retention policy.'

At the bottom of the wizard, there are three buttons: 'Previous', 'Next', and 'Cancel'.

NOTE

As soon as you click **Next**, Veeam Backup for AWS will check the repository ownership. If the backup repository is already managed by another backup appliance, you will receive a warning. To learn how to eliminate this warning, see [Eliminating Repository Ownership Warning](#).

Eliminating Repository Ownership Warning

Due to technical limitations, a backup repository that is already added to a backup appliance cannot be added to other backup appliances. Otherwise, retention sessions running on different appliances may corrupt backup files stored in this repository, which may result in unpredictable data loss. For this reason, Veeam Backup for AWS checks whether a repository is managed by any appliances as soon as you click **Next** at the **Bucket** step of the wizard.

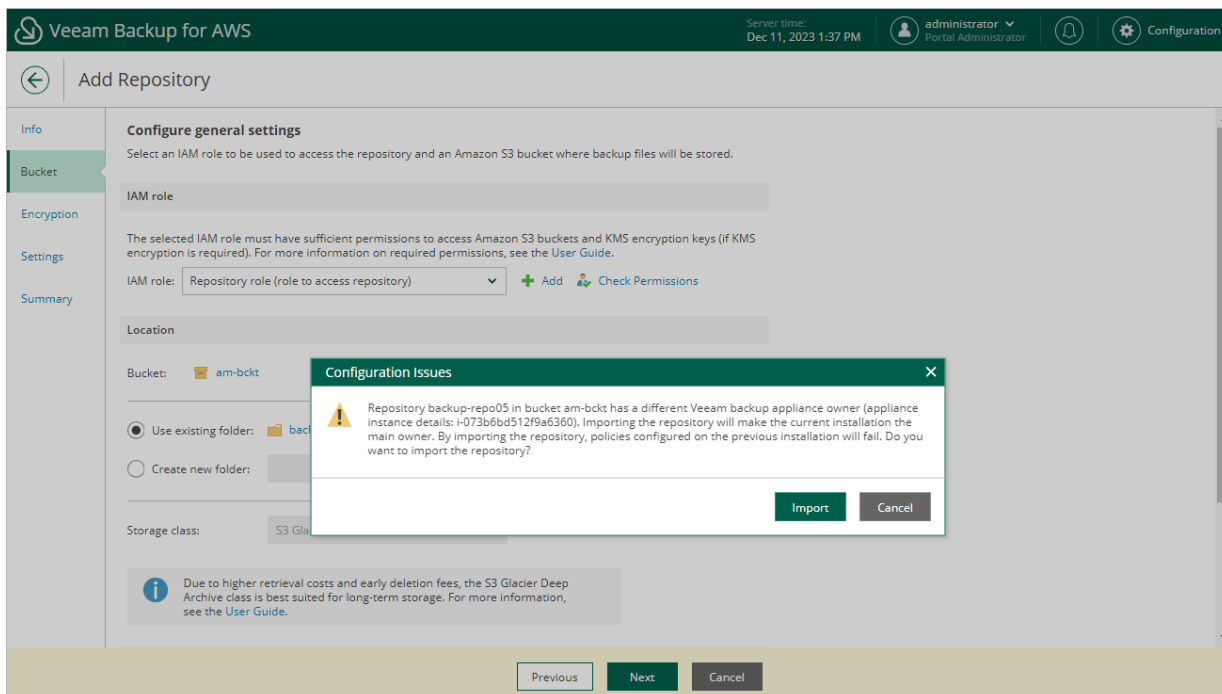
NOTE

The repository ownership check is supported only for those backup repositories that are managed by backup appliances running Veeam Backup for AWS version 7.0 or later.

If the backup repository is already added to another backup appliance, Veeam Backup for AWS will display a warning notifying that the repository has an owner. If you want the current appliance to become the owner of the repository, click **Import**. If you want to preserve the repository ownership, click **Cancel** and choose another folder as a target backup repository.

IMPORTANT

If you chose to import the backup repository, the previous backup appliance will lose the repository ownership and you will have to reconfigure the repository settings on that appliance. Otherwise, backup policies configured on the appliance will start failing.



Step 4. Enable Data Encryption

[This step applies only if you have selected the **Create new folder** option at the **Bucket** step of the wizard, or if you have selected an existing folder with encryption enabled at the repository level]

At the **Encryption** step of the wizard, do either of the following:

- If you have selected an existing folder at the **Bucket** step of the wizard, you must provide the currently used password or an encryption key that was used to encrypt data stored in this folder to let Veeam Backup for AWS access the folder and add it as a backup repository. You cannot change these settings while adding the repository – however, you will be able to [edit the repository settings](#) later.

IMPORTANT

Make sure that the encryption key that was used to encrypt data stored in the existing folder is available in AWS, and the IAM role specified at [step 3](#) of the wizard has permissions to access the key. Otherwise, Veeam Backup for AWS will not be able to add the repository to the backup infrastructure.

- If you have selected the **Create new folder** option at the **Bucket** step of the wizard, choose whether you want to encrypt backup files stored in the selected Amazon S3 bucket folder. Before you enable encryption at the repository level, check the limitations described in section [Limitations and Considerations](#).

To enable encryption:

- a. Set the **Enable encryption** toggle to *On*.
- b. Choose whether you want to use a password or an [AWS Key Management Service \(KMS\) key](#) to encrypt the backed-up data. For more information on encryption algorithms, see [Backup Repository Encryption](#).
 - To encrypt data using a password, select the **Use password encryption** option and specify the password and a password hint.

- To encrypt data using an AWS KMS key, select the **Use KMS encryption key** option and choose the necessary KMS key from the **Encryption key** drop-down list.

For a KMS key to be displayed in the list of available encryption keys, it must be created in the AWS Region where the selected Amazon S3 bucket is located, and the IAM role specified to access the bucket must have permissions to the key. For more information on permissions required for the IAM role, see [Repository IAM Role Permissions](#).

Veeam Backup for AWS Server time: Nov 30, 2023 5:07 PM administrator Portal Administrator Configuration

Add Repository

Configure encryption settings
Choose whether you want to enable encryption at the repository level.

Enable encryption: On

Use password encryption

Password:

Repeat password:

Password hint:

Use KMS encryption key

Encryption key:

i The selected IAM role must have permissions to access the encryption key. For more information, see the [User Guide](#).

Previous Next Cancel

Step 5. Specify VPC Interface Endpoint


[This step applies only if you have enabled the [private network deployment](#) functionality]

At the **Settings** step of the wizard, specify an S3 interface endpoint that will be used to communicate with the Amazon S3 service.

For an S3 interface endpoint to be displayed in the **Interface VPC endpoint** list, it must be created in the Amazon VPC console for all subnets to which the worker instances will be connected, as described in section [Configuring Private Networks](#).

IMPORTANT

S3 gateway endpoints are not supported when using the private network deployment functionality.



Step 6. Finish Working with Wizard

At the **Summary** step of the wizard, review summary information and check whether the specified IAM role has all the required permissions – to do that, click **Check Permissions**. Veeam Backup for AWS will display the **Permission check** window where you can track the progress and view the results of the performed check. If some permissions of the IAM role are missing, the check will complete with errors, and the list of permissions that must be granted to the IAM role will be displayed in the **Missing Permissions** column.

You can grant the missing permissions to the IAM role [using the AWS Management Console](#) or instruct Veeam Backup for AWS to do it:

1. In the **Permission check** window, click **Grant**.
2. In the **Grant permissions** window, provide one-time access keys of an IAM user that is authorized to update permissions of IAM roles, and then click **Apply**.

The IAM user must have the following permissions:

```
"iam:AttachRolePolicy",  
"iam:CreatePolicy",  
"iam:CreatePolicyVersion",  
"iam:CreateRole",  
"iam:GetAccountSummary",  
"iam:GetPolicy",  
"iam:GetPolicyVersion",  
"iam:GetRole",  
"iam>ListAttachedRolePolicies",  
"iam>ListPolicyVersions",  
"iam:SimulatePrincipalPolicy",  
"iam:UpdateAssumeRolePolicy"
```

NOTE

Veeam Backup for AWS does not store one-time access keys in the configuration database.

- After the required permissions are granted, close the **Permission check** window and review configuration information. Then, choose whether you want to proceed to the [Session Log tab](#) to track the progress of repository creation and click **Finish**.

The screenshot shows the 'Add Repository' wizard in Veeam Backup for AWS. The left sidebar contains navigation tabs: Info, Bucket, Encryption, Settings, and Summary. The main area is divided into two panes. The left pane, titled 'Review configured settings', shows configuration details for an immutable repository. The right pane, titled 'Permission check', displays a success message and a table of permissions.

Review configured settings
 Review the repository settings, and click Finish to exit the wizard.

To verify whether the selected IAM role has all the necessary permissions, click **Check Permissions** or **Copy to Clipboard**.

Info
 Name: Immutable Repository
 Description: Repository for storing immutable data

Bucket
 IAM role: Repository role
 Storage class: S3 Standard
 Region: Europe (Paris)
 Bucket: am-bckt-immutable
 Folder: dept_01_immutable
 Immutability: Enabled

Encryption
 Encryption: Enabled
 Type: Password
 Password hint: hint

After you click Finish, the repository will be created. To view the session progress, switch to the Session Logs page.

Go to Sessions

Permission check

Your account meets the required permissions.

Grant Recheck Export Missing Permissions

Type	Status	Missing Permissions
KMS permissions	Passed	—
S3 permissions	Passed	—
EC2 permissions	Passed	—
IAM permissions	Passed	—
Trust relationships	Passed	—

Close

Editing Backup Repository Settings

The settings that you can modify for a backup repository depend on whether the repository has been added to the backup infrastructure using the Veeam Backup & Replication console or the Veeam Backup for AWS Web UI.

Editing Backup Repository Settings Using Veeam Backup & Replication Console

For each standard backup repository, you can modify settings configured while adding the repository to the backup infrastructure:


1. In the Veeam Backup & Replication console, open the **Backup Infrastructure** view.
2. Navigate to **External Repositories**.
3. Select the necessary repository and click **Edit Repository** on the ribbon.
Alternatively, you can right-click the repository and select **Properties**.
4. Complete the **Edit External Repository** wizard:
 - a. To specify a new name and description for the repository, follow the instructions provided in section [Creating New Repositories](#) (step 2).
 - b. To change the access keys of the IAM user and the gateway server used to access the repository, follow the instructions provided in section [Creating New Repositories](#) (step 3).
 - c. To enable encryption or change the encryption settings of the repository, follow the instructions provided in section [Creating New Repositories](#) (step 6).

IMPORTANT

If you change the encryption settings of the repository from the Veeam Backup & Replication console, Veeam Backup & Replication will not propagate these settings to the backup appliance automatically. Consider updating the settings manually as described in [Editing Backup Repository Settings Using Veeam Backup for AWS Web UI](#).

- d. At the **Apply** step of the wizard, wait for the changes to be applied and click **Next**.

- e. At the **Summary** step of the wizard, review summary information and click **Finish** to confirm the changes.



Editing Backup Repository Settings Using Veeam Backup for AWS Web UI

For each backup repository, you can modify settings configured while adding the repository to Veeam Backup for AWS:

1. Switch to the **Configuration** page.
2. Navigate to **Repositories**.
3. Select the check box next to the backup repository and click **Edit**.
4. Complete the **Edit Repository** wizard.
 - a. To provide a new name and description for the backup repository, follow the instructions provided in section [Adding Backup Repositories Using Web UI](#) (step 2).
 - b. To change the IAM role whose permissions Veeam Backup for AWS uses to access the repository, follow the instructions provided in section [Adding Backup Repositories Using Web UI](#) (step 3).
 - c. [Applies only to repositories managed by another backup appliance] To change the owner of the backup repository, navigate to the **Bucket** step and click **Next**. Then, follow the instructions provided in section [Adding Backup Repositories Using Web UI](#) (step 3).
 - d. To enable data encryption or change the configured encryption settings, follow the instructions provided in section [Adding Backup Repositories Using Web UI](#) (step 4).
 - e. To specify the S3 interface endpoint that will be used to communicate with the Amazon S3 service in private deployment mode, follow the instructions provided in section [Adding Backup Repositories Using Web UI](#) (step 5).

- f. At the **Summary** step of the wizard, review summary information, choose whether you want to proceed to the [Session Log tab](#) to track the progress of modifying the backup repository settings, and click **Finish** to confirm the changes.

Manage your backup repositories used to store backed-up data created by Veeam Backup for AWS. If required, encryption can be enabled at the repository level.

Repository Filter (None)

+ Add Edit View Info Remove Export to... ▼

Repository	Description	Bucket	Folder	Storage Class	Immutability
<input type="checkbox"/> backup-repo-02	archive backup repository	am-bckt	backup-repo-02	S3 Glacier Flexible R...	Disabled
<input checked="" type="checkbox"/> backup-repo-03	primary backup repository	am-bckt	backup-repo-03	S3 Standard	Disabled
<input type="checkbox"/> backup-repo04	Created by administrator at...	am-bckt	am-folder02	S3 Standard	Disabled

Selected: 1 of 3

Rescanning Backup Repositories

Veeam Backup & Replication periodically rescans standard backup repositories for newly created restore points and metadata – the results of every rescan session are displayed in the **History** view under the **System** node. A rescan operation is launched automatically every 24 hours or in the following cases:


- After you add a repository to the backup infrastructure.
- After a backup chain stored in the repository is modified (for example, if a restore point is added or deleted from the chain).

However, you can perform a rescan operation for a repository manually:

1. In the Veeam Backup & Replication console, open the **Backup Infrastructure** view.
2. Navigate to **External Repositories**.
3. Select the necessary repository and click **Rescan** on the ribbon.

Alternatively, you can right-click the repository and select **Rescan**.

If multiple repositories are present in the backup infrastructure, you can perform the rescan operation for all repositories simultaneously. To do that, right-click the **External Repositories** node and select **Rescan**.



Removing Backup Repositories

The consequences of actions performed with a backup repository depend on whether the repository has been added to the backup infrastructure using the Veeam Backup & Replication console or the Veeam Backup for AWS Web UI.


Removing Backup Repository Using Veeam Backup & Replication Console

AWS Plug-in for Veeam Backup & Replication allows you to permanently remove repositories from the backup infrastructure:

1. In the Veeam Backup & Replication console, open the **Backup Infrastructure** view.
2. Navigate to **External Repositories**.
3. Select the necessary repository and click **Remove Repository** on the ribbon.

Alternatively, you can right-click the repository and select **Remove**.

Note that the repository will not be removed from the backup appliance. To learn how to remove repositories from backup appliances, see [Removing Backup Repository Using Veeam Backup for AWS Web UI](#).



Removing Backup Repository Using Veeam Backup for AWS Web UI

You can remove backup repositories from Veeam Backup for AWS. When you remove a repository, Veeam Backup for AWS unassigns the repository role from the folder in the Amazon S3 bucket so that this folder is no longer used as a backup repository.

NOTE

Even though the Amazon S3 bucket is no longer used as a backup repository, Veeam Backup for AWS preserves all backup files previously stored in the repository and keeps these files in Amazon S3. You can assign the Amazon S3 bucket to a new backup repository so that Veeam Backup for AWS imports the backed-up data to the configuration database. In this case, you will be able to perform all disaster recovery operations described in section [Performing Restore](#).


If you no longer need the backed-up data, either delete it as described in section [Managing Backed-Up Data](#) before you remove the repository from Veeam Backup for AWS, or [use the AWS Management Console](#) to delete the data if the repository has already been removed.

To remove a backup repository, do the following:

1. Switch to the **Configuration** page.
2. Navigate to **Repositories**.
3. Select the check box next to the backup repository and click **Remove**.
4. In the **Remove Repository** window, click **Remove** to acknowledge the operation.

IMPORTANT

You cannot remove a backup repository that is used by any backup policy or by a scheduled configuration backup. Modify the settings of all the related policies to remove references to the repository, and then try removing the repository again. To learn how to modify the backup policy settings, see [Performing Backup](#).



Managing IAM Roles

NOTE

This section assumes that you have a good understanding of [IAM Roles](#), [Creating IAM Policies](#) and [Adding and Removing IAM Identity Permissions](#).

Veeam Backup for AWS uses permissions of IAM roles to access AWS services and resources, and to perform the backup and restore operations. For example, Veeam Backup for AWS requires access to the following AWS resources:

- **EC2 resources** – to display the list of EC2 instances in backup policy settings, to create cloud-native snapshots, snapshot replicas, to deploy worker instances and to restore backed-up data.
- **S3 resources** – to store backed-up data in backup repositories, to perform transform operations with backup chains, and to copy backed-up data from backup repositories to worker instances during restore.

For each data protection and disaster recovery operation performed by Veeam Backup for AWS, you must specify an IAM role. By design, Veeam Backup for AWS comes with the *Default Backup Restore* IAM role. This role is added to the configuration database upon product installation and is automatically assigned all the permissions required to perform data protection tasks within the initial AWS account in which the backup appliance resides.

If you want to back up and restore resources in other AWS accounts, or if you want to specify custom IAM roles with granular permissions to perform specific operations, [add IAM roles to Veeam Backup for AWS](#). You can add IAM roles that already exist in your AWS accounts, or instruct Veeam Backup for AWS to [create and add IAM roles](#) with predefined permission sets. To learn how to create IAM roles in the AWS Management Console, see [Appendix A. Creating IAM Roles in AWS](#).


Adding IAM Roles

To add an IAM role to Veeam Backup for AWS, do the following:

1. [Launch the Add IAM Role wizard.](#)
2. [Specify a name and description for the IAM role.](#)
3. [Specify IAM role settings.](#)
4. [Specify IAM role permissions.](#)
5. [Finish working with the wizard.](#)

Before You Begin

When you deploy a backup appliance, Veeam Backup for AWS automatically creates a specific IAM role named *Impersonation* role – and attaches this role to the backup appliance. The *Impersonation* IAM role is then used to assume other IAM roles added to Veeam Backup for AWS to perform operations in your infrastructure, and is automatically assigned the `sts:AssumeRole` permission required to assume these roles.



Before you start adding an IAM role to Veeam Backup for AWS, you must check the following prerequisites:

- The *Impersonation* IAM role must have the following permissions:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "sts:AssumeRole"
      ],
      "Resource": "*"
    }
  ]
}
```

To obtain the ARN of the *Impersonation* IAM role, you can look it up on the **Instances** page in the AWS Management Console.

- Trust relationships must be configured for the IAM role you want to add, and the following statement must be included into the trust policy:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "sts:AssumeRole",
      "Principal": {
        "AWS": "<Role ARN>"
      }
    }
  ]
}
```

Where `<Role ARN>` is either the ARN of the *Impersonation* IAM role or the ARN of the AWS account to which the backup appliance belongs.

Configuring Trust Relationships

To allow Veeam Backup for AWS to use an IAM role to perform operations in your infrastructure, you must configure trust relationships for the role you want to add:

1. Open the [EC2 console](#) and do the following:
 - a. Navigate to **Instances**.
 - b. In the **Instances** section, locate the EC2 instance running the backup appliance.
 - c. On the **Summary** page, switch to the **Security** tab and click the link in the **IAM Role** field. The **IAM** console will open.
2. In the [IAM console](#), do the following:
 - a. Copy the value displayed in the **ARN** field – you will need it later.

- b. Navigate to **Roles** and locate the IAM role for which you want to configure trust relationships.
- c. On the **Summary** page, switch to the **Trust relationships** tab and click **Edit trust policy**.
- d. In the **Edit trust policy** field, add the following statement:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "sts:AssumeRole",
      "Principal": {
        "AWS": "<Role ARN>"
      }
    }
  ]
}
```

Where `<Role ARN>` is the ARN either of the *Impersonation* IAM role that you have copied at step 2a, or the ARN of the AWS account to which the backup appliance belongs.

- e. Click **Update policy**. Note that it may take up to 5 minutes for AWS to update the trust policy.

Step 1. Launch Add IAM Role Wizard

To launch the **Add IAM Role** wizard, do the following:

1. Switch to the **Configuration** page.
2. Navigate to **Accounts > IAM Roles**.
3. Click **Add**.

The screenshot displays the Veeam Backup for AWS configuration interface. The top navigation bar includes the Veeam logo, the text 'Veeam Backup for AWS', the server time 'Jun 13, 2024 7:41 AM', and the user 'administrator Portal Administrator'. The left sidebar contains navigation options: Exit Configuration, Getting Started, Administration, Accounts, Repositories, Workers, Server settings, General, Configuration Backup, Licensing, and Support Information. The main content area is titled 'IAM Roles' and contains a descriptive paragraph about IAM roles. Below this is an information icon and a note about the Default Backup Restore IAM role. A search bar and a '+ Add' button are visible. A table lists existing IAM roles with columns for IAM Role, AWS Account, Last Edited, and Description. The 'Add' button is highlighted with a mouse cursor.

<input type="checkbox"/>	IAM Role	AWS Account	Last Edited	Description ↑	⋮
Selected: 0 of 7					
<input type="checkbox"/>	Default Backup Restore	611610170176 (veeam-tw)	08/15/2023 10:32:56 AM	Default Backup Restore	
<input type="checkbox"/>	Replication role	611610170176 (veeam-tw)	06/29/2023 3:48:15 PM	role for cloud-native snapshots replication	
<input type="checkbox"/>	Repository role	611610170176 (veeam-tw)	06/30/2023 9:00:15 AM	role to access repository	
<input type="checkbox"/>	Indexing worker role	611610170176 (veeam-tw)	07/13/2023 12:40:45 PM	role to attach to and communicate with worker in...	

Step 2. Specify IAM Role Name and Description

At the **Info** step of the wizard, specify a name and description for the IAM role. The specified name and description will be displayed on the **IAM Roles** tab.

Consider the following limitations:

- The specified name must be unique in Veeam Backup for AWS.
- The length of the name must not exceed 127 characters.
- The length of the specified description must not exceed 255 characters.

The screenshot shows the 'Add IAM Role' wizard in the Veeam Backup for AWS console. The interface is in a dark theme. At the top, the header bar displays 'Veeam Backup for AWS', the server time 'Jun 13, 2024 7:43 AM', and the user 'administrator Portal Administrator'. The main content area is titled 'Add IAM Role' and is currently on the 'Info' step. A sidebar on the left contains tabs for 'Info', 'Type', 'Permissions', and 'Summary'. The 'Info' tab is active, showing the title 'Specify IAM role name and description' and the instruction 'Enter a name and description for the IAM role.' Below this, there are two input fields: 'Name:' with the value 'Production worker role' and 'Description:' with the value 'role to launch worker instances in production accounts'. At the bottom right of the form, there are 'Next' and 'Cancel' buttons.

Step 3. Specify IAM Role Settings

At the **Type** step of the wizard, select one of the following options:

- **IAM role from current account** – select this option if you want to add an existing IAM role from the AWS account to which the backup appliance belongs.
- **IAM role from another account** – select this option if you want to add an existing IAM role from an AWS account other than the account to which the backup appliance belongs.
- **Create new IAM role** – select this option if you want Veeam Backup for AWS to create a new IAM role in AWS automatically.
- **Create template** – select this option if you want Veeam Backup for AWS to create a CloudFormation template or a JSON policy document that you can then use to create an IAM role in AWS manually.

Specifying Settings for IAM Role from Initial Account

[This step applies only if you have selected the **IAM role from current account** option]

At the **Type** step of the wizard, use the **AWS role name** field to enter the IAM role name as specified in AWS. If the IAM role was created with a path, you must specify the complete path and the name of the IAM role. For example, */dept_1/backup_role*.

IMPORTANT

To allow the backup appliance to assume the IAM role, you must configure trust relationships for the role as described in section [Before You Begin](#).

The screenshot shows the 'Add IAM Role' wizard in Veeam Backup for AWS. The 'Type' step is active, displaying four radio button options. The first option, 'IAM role from current account', is selected. Below it, there is a text input field for 'AWS role name' containing the text 'production_worker_role'. The other three options are 'IAM role from another account', 'Create new IAM role', and 'Create template'. At the bottom of the wizard, there are 'Previous', 'Next', and 'Cancel' buttons.

Specifying Settings for IAM Role from Another Account

[This step applies only if you have selected the **IAM role from another account** option]

At the **Type** step of the wizard, specify the following settings:

1. In the **Account ID** field, enter the 12-digit number of the AWS account to which the IAM role you want to add belongs.

2. In the **AWS role name** field, enter the IAM role name as specified in AWS.

If the IAM role was created with a path, you must specify the complete path and the name of the IAM role. For example, `/dept_1/backup_role`.

3. [Optional] In the **External ID** field, enter the external ID – the property in the trust policy of the IAM role from another account used for enhanced security. For more information, see [AWS Documentation](#).

IMPORTANT

To allow the backup appliance to assume the IAM role, you must configure trust relationships for the role as described in section [Before You Begin](#).

The screenshot shows the 'Add IAM Role' wizard in Veeam Backup for AWS. The 'Type' step is active, and the 'IAM role from another account' option is selected. The following fields are filled in:

- Account ID: 3945681109
- AWS role name: production_worker_role
- External ID (optional): ad39458dept01

At the bottom of the wizard, there are 'Previous', 'Next', and 'Cancel' buttons.

Specifying Settings for New IAM Role

[This step applies only if you have selected the **Create new IAM role** option]

At the **Type** step of the wizard, specify the following settings:

1. In the **AWS role name** field, specify a name that will be used to create the IAM role in AWS.

Consider the following limitations:

- The specified name must be unique within one AWS account.
- The following characters are not supported: \ / " ' [] : | < > ; ? * & .
- The length of the name must not exceed 63 characters.

For more information on IAM name limitations, see [AWS Documentation](#).

2. Provide one-time access keys of an IAM user that is authorized to create IAM roles in the AWS account.

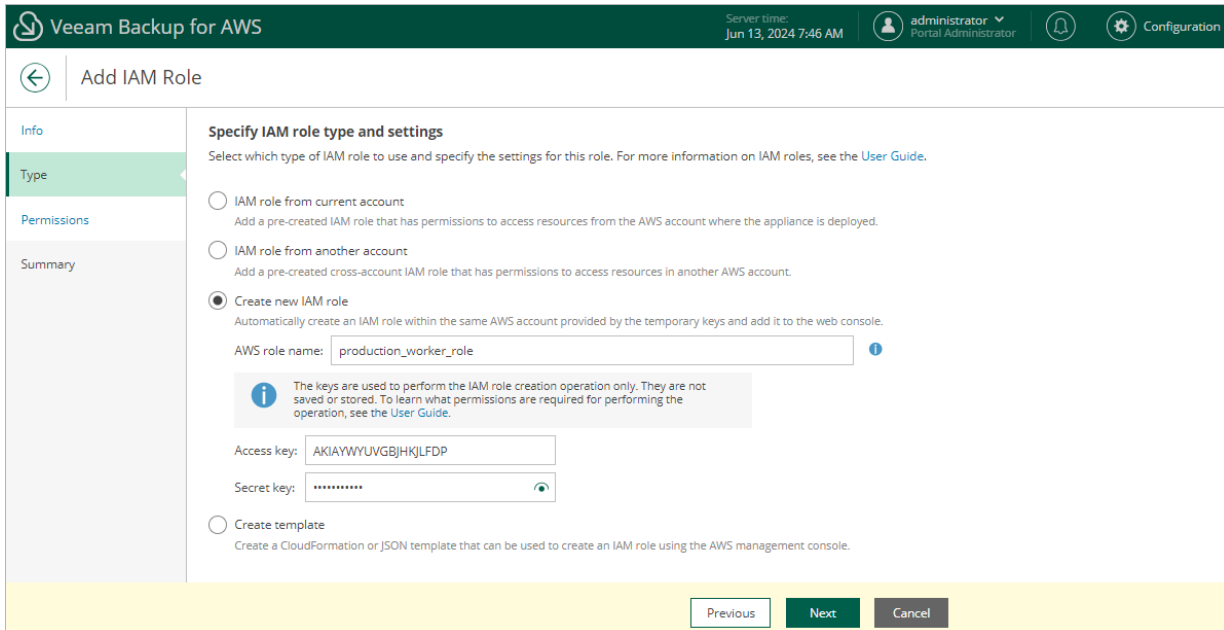
The specified access keys determine in which AWS account the role will be created. For example, if you specify access keys of an IAM user from the initial AWS account, the IAM role will be created in the initial AWS account.

The IAM user must have the following permissions:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "Statement1",
      "Effect": "Allow",
      "Action": [
        "iam:AddRoleToInstanceProfile",
        "iam:AttachRolePolicy",
        "iam:CreateInstanceProfile",
        "iam:CreatePolicy",
        "iam:CreatePolicyVersion",
        "iam:CreateRole",
        "iam>DeletePolicyVersion",
        "iam:GetAccountSummary",
        "iam:GetInstanceProfile",
        "iam:GetPolicy",
        "iam:GetPolicyVersion",
        "iam:GetRole",
        "iam:ListAttachedRolePolicies",
        "iam:ListInstanceProfilesForRole",
        "iam:ListPolicyVersions",
        "iam:PassRole",
        "iam:SimulatePrincipalPolicy",
        "iam:UpdateAssumeRolePolicy"
      ],
      "Resource": [
        "*"
      ]
    }
  ]
}
```

NOTE

Veeam Backup for AWS does not store one-time access keys in the configuration database.



The screenshot shows the 'Add IAM Role' wizard in Veeam Backup for AWS. The 'Type' step is active, and the 'Create new IAM role' option is selected. The 'AWS role name' field is filled with 'production_worker_role'. There are also fields for 'Access key' and 'Secret key'. The 'Next' button is highlighted in green.

Specifying Settings for Template

[This step applies only if you have selected the **Create template** option]

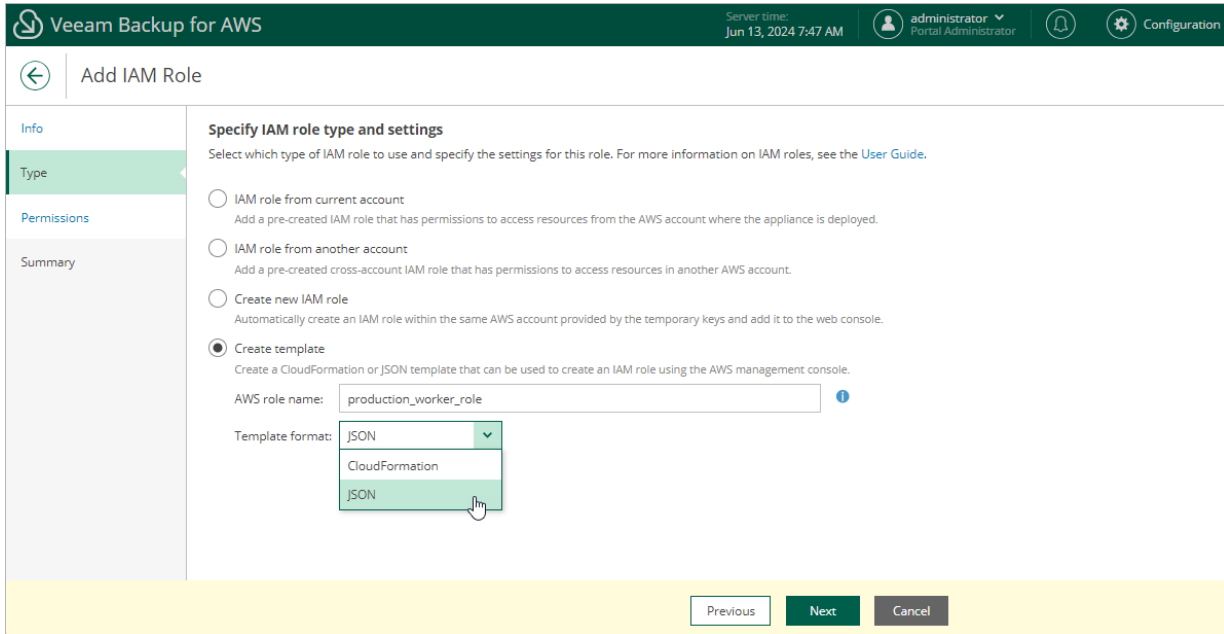
At the **Type** step of the wizard, specify the following settings:

1. In the **AWS role name** field, specify a name that will be assigned to the IAM role in AWS.
2. Use the **Template format** drop-down list to choose whether you want to generate a CloudFormation template or a JSON policy document that will be used to create the IAM role in the AWS Management Console:
 - Select *CloudFormation* if you want to create a CloudFormation template and export it to a .CFORM file. You can further upload the file to the CloudFormation service and use it to create the IAM role automatically.

To learn how to upload templates to the CloudFormation service, see [AWS Documentation](#).

- Select *JSON* if you want to create a policy document and export it to a .JSON file. You can further use the file to create an IAM policy using the IAM service and attach the policy to the IAM role manually.

To learn how to create an IAM role in the AWS Management Console, see [Appendix A. Creating IAM Roles in AWS](#). To learn how to attach IAM policies to IAM roles, see [Appendix B. Creating IAM Policies in AWS](#).



Step 4. Specify IAM Role Permissions

At the **Permissions** step of the wizard, you can define specific operations that Veeam Backup for AWS will be able to perform using the permissions of the created IAM role. Depending on the option that you have selected at the **Type** step of the wizard, Veeam Backup for AWS will do either of the following:

- If you have selected the **IAM role from current account** or the **IAM role from another account** option, Veeam Backup for AWS will become able to filter IAM roles and check their permissions in backup and restore settings – but it will not assign any permissions to the role.

In this case, you can grant the permissions to the role manually [using the AWS Management Console](#) or instruct Veeam Backup for AWS to do it, as described in section [Checking IAM Role Permissions](#).

- If you have selected the **Create new IAM role** option, Veeam Backup for AWS will become able to filter IAM roles and check their permissions in backup and restore settings – and will also assign the specified permissions to the role.
- If you have selected the **Create template** option, Veeam Backup for AWS will add the specified permissions to the created CloudFormation template or JSON policy document.

To specify permissions granularly, do the following:

1. In the **Specify IAM role permissions** section, set the **Specify granular permissions** toggle to *On*.
2. In the **Veeam management roles** section, choose actions that will be performed using the IAM role:
 - **Worker deployment role** – will be used to deploy worker instances in the [backup account](#). If you choose this action for an IAM role, you will be able to select it [when adding worker configurations](#).
 - **Production worker role** – will be used to communicate with worker instances in [production accounts](#). If you choose this action for an IAM role, you will be able to select it [when enabling indexing for EFS policies](#), [creating EC2 backup policies](#), [creating RDS backup policies](#), [performing entire EC2 instance restore](#), [performing EC2 volume-level restore](#) or [performing RDS database restore](#).
 - **Repository role** – will be used to create new repositories in Amazon S3 buckets and to further access the repositories during data protection and disaster recovery operations. If you choose this action for an IAM role, you will be able to select it [when configuring repository settings](#).

IMPORTANT

For Veeam Backup for AWS to perform the selected actions using the IAM role, it must be assigned the permissions listed in sections [Worker Deployment Role Permissions in Backup Account](#), [Worker Deployment Role Permissions in Production Accounts](#) and [Repository IAM Permissions](#).

3. In the **Workload permissions** section, choose resources that will be protected using the IAM role, and operations that will be performed with these resources:
 - **Backup** – Veeam Backup for AWS will protect EC2, Redshift, DynamoDB, EFS, FSx and VPC resources. If you select this operation for an IAM role, you will be able to select it in the [EC2 backup](#), [Redshift backup](#), [DynamoDB backup](#), [EFS backup](#), [FSx backup](#) and [VPC configuration backup](#) settings.
Note that the list of permissions for this role will also contain additional permissions required to deploy worker instances in production accounts during EFS indexing and EC2 backup operations.
 - **Replication** – Veeam Backup for AWS will replicate cloud-native snapshots of EC2 and RDS resources. If you select this operation for an IAM role, you will be able to select it in the [EC2 backup](#) and [RDS backup](#) settings.

- **Snapshot** – Veeam Backup for AWS will create cloud-native snapshots of RDS resources. If you select this operation for an IAM role, you will be able to select it in the [RDS backup](#) settings.

Note that the list of permissions for this role will also contain additional permissions required to deploy worker instances in production accounts during RDS backup operations.

- **Restore** – Veeam Backup for AWS will restore EC2, RDS, Redshift, DynamoDB, EFS, FSx and VPC resources. If you select this operation for an IAM role, you will be able to select it when performing [entire EC2 instance restore](#), [EC2 volume-level restore](#), [EC2 file-level recovery](#), [RDS restore](#), [Redshift restore](#), [DynamoDB restore](#), [EFS restore](#), [FSx restore](#), [entire VPC configuration restore](#) and [selected VPC items restore](#).

Note that the list of permissions for this role will also contain additional permissions required to deploy worker instances in production accounts during EC2 and RDS restore operations.

IMPORTANT

For Veeam Backup for AWS to perform the selected operations using the IAM role, it must be assigned the permissions listed in sections [Backup IAM Permissions](#) and [Restore IAM Permissions](#).

Note that if you do not specify any management roles and resource permissions for the IAM role at this step, all the listed actions and resource operations will be selected for the role automatically.

The screenshot displays the 'Add IAM Role' configuration interface in Veeam Backup for AWS. The interface is divided into several sections:

- Header:** Veeam Backup for AWS, Server time: Jun 13, 2024 7:49 AM, administrator Portal Administrator, Configuration.
- Navigation:** Back arrow, Add IAM Role.
- Info Panel (Left):** Info, Type, Permissions (selected), Summary.
- Specify IAM role permission:** To perform a permissions check, you must select the required permissions. An information icon indicates: 'By selecting which workload you are protecting, you can configure the permissions for the IAM role.' Below this, there are buttons for 'Select All', 'Clear All', and 'Reset'. A toggle for 'Specify granular permissions' is currently turned on.
- Veeam management permissions:** A list of permissions for various AWS services: Amazon EC2, Amazon RDS, Amazon EFS, Amazon VPC, Amazon DynamoDB, Amazon FSx, and Amazon Redshift. An 'Edit Permissions' link is provided.
- Veeam management roles:** A section for selecting roles to manage backup repositories. It includes:
 - Worker deployment role
 - Production worker role
 - Repository role
- Workload permissions:** A section for selecting workloads to protect and the actions the IAM role should perform. It includes:
 - Amazon EC2:
 - Backup
 - Replication
 - Restore
 - Amazon RDS:
 - Snapshot
 - Replication
 - Restore
 - Amazon EFS:
 - Backup
 - Restore
 - Amazon VPC
 - Amazon DynamoDB
 - Amazon FSx
 - Amazon Redshift

Step 5. Finish Working with Wizard

At the **Summary** step of the wizard, review configuration information and click **Finish**.

TIPS

- You can view the configured IAM role permissions at the IAM Roles tab. To do that, select the necessary IAM role and click **View Permissions**.
- After you add the IAM role to Veeam Backup for AWS, it is recommended that you verify whether the IAM role has all the permissions required to perform operations with the selected workloads. That is why make sure that the **Perform permission check when I click finish** check box is selected – in this case, Veeam Backup for AWS will display the **Permission check** window where you can track the progress and view the results of the check.

The screenshot shows the 'Add IAM Role' wizard in the Veeam Backup for AWS console. The 'Summary' step is active, displaying the following configuration details:

Info	
Name:	Production worker role
Description:	role to launch worker instances in production accounts

Type	
Type	IAM role from the current account

Permissions	
Veeam management permissions:	Production worker role
Amazon EC2:	Backup, Restore
Amazon RDS:	No permissions configured
Amazon EFS:	Backup, Restore
Amazon VPC:	No permissions configured
Amazon DynamoDB:	No permissions configured
Amazon FSx:	No permissions configured
Amazon Redshift:	No permissions configured

After you complete the wizard the IAM role will be added. It is recommended to perform a permission check to assure everything is configured correctly.

Perform permission check when I click Finish

Navigation buttons: Previous, Finish, Cancel

Editing IAM Role Settings

For each IAM role added to Veeam Backup for AWS, you can modify the IAM role settings:

1. Switch to the **Configuration** page.
2. Navigate to **Accounts > IAM Roles**.
3. Select the check box next to an IAM role whose settings you want to edit.
4. Click **Edit**.
5. Complete the **Edit IAM Role** wizard.
 - a. To provide a new name and description for the IAM role, follow the instructions provided in section [Adding IAM Roles](#) (step 2).
 - b. To edit the IAM role settings, follow the instructions provided in section [Adding IAM Roles](#) (step 3).
 - c. To edit the IAM role permissions, follow the instructions provided in section [Adding IAM Roles](#) (step 4).

When you edit the workload permissions, Veeam Backup for AWS does not automatically update the permissions already assigned to the IAM role. If you want to update these permissions, you must manually modify the IAM role in AWS Management Console as described in [AWS Documentation](#).
 - d. At the **Permission check** step of the wizard, Veeam Backup for AWS will verify whether the IAM role has all the permissions required to perform operations with the selected workloads.

If some of the required permissions are missing, the check will complete with errors, and the **Missing Permissions** column will display the list of permissions that must be granted to the IAM role. You can grant the missing permissions to the IAM role [using the AWS Management Console](#) or instruct Veeam Backup for AWS to do it, as described in section [Checking IAM Role Permissions](#).
 - e. At the **Summary** step of the wizard, review summary information and click **Finish** to confirm the changes.

IMPORTANT

After you upgrade Veeam Backup for AWS from a version prior to 7.0, the **IAM Roles** page will also display roles that previously existed on the backup appliance, with all actions and workload permissions available. If you want to modify the list of actions and permissions assigned to the roles, follow the instructions provided in section [Adding IAM Roles](#) (step 4).

The screenshot shows the Veeam Backup for AWS configuration interface. The top navigation bar includes the Veeam logo, the text 'Veeam Backup for AWS', the server time 'Jun 13, 2024 7:51 AM', and the user 'administrator Portal Administrator'. The left sidebar contains navigation options: Exit Configuration, Getting Started, Administration, Accounts (selected), Repositories, Workers, Server settings, General, Configuration Backup, Licensing, and Support Information. The main content area is titled 'IAM Roles' and includes a sub-header with 'Database Accounts', 'SMTP Accounts', and 'Portal Users'. A descriptive paragraph explains that Veeam Backup for AWS uses IAM roles for data protection and disaster recovery. Below this is an information icon and a note: 'The Default Backup Restore IAM role is preconfigured and has all the required permissions to protect instances within the initial AWS account. For large and secure environments with multiple accounts, we recommend to configure multiple IAM roles specific to each account.' A toolbar contains a search box for 'IAM Role', and buttons for '+ Add', 'Edit', 'Remove', 'Check Permissions', 'View Permissions', and 'Export to...'. A table lists the IAM roles with columns for 'IAM Role', 'AWS Account', 'Last Edited', and 'Description'. The 'Repository role' is selected with a checkmark.

<input type="checkbox"/>	IAM Role	AWS Account	Last Edited	Description ↑
<input type="checkbox"/>	Default Backup Restore	611610171076 (veeam-tw)	08/15/2023 10:32:56 AM	Default Backup Restore
<input type="checkbox"/>	Replication role	611610171076 (veeam-tw)	06/29/2023 3:48:15 PM	role for cloud-native snapshots replication
<input checked="" type="checkbox"/>	Repository role	611610171076 (veeam-tw)	06/30/2023 9:00:15 AM	role to access repository
<input type="checkbox"/>	Indexing worker role	611610171076 (veeam-tw)	07/13/2023 12:40:45 PM	role to attach to and communicate with worker in...
<input type="checkbox"/>	Production worker role	611610171076 (veeam-tw)	07/03/2023 3:09:41 PM	role to launch worker instances in production acc...

Checking IAM Role Permissions

It is recommended that you check whether IAM roles specified to perform operations in Veeam Backup for AWS have all the required permissions – otherwise, the operations may fail to complete successfully. The check must be performed not only when you specify a new IAM role to perform an operation, but also after you make any changes in your AWS account and want to ensure that the permissions granted to the existing IAM roles remain sufficient.

You can verify IAM role permissions either using the built-in wizard permission check that is available when specifying roles for operations, or using the permission check at the **IAM Roles** tab or in the **Edit IAM Role** wizard.

IMPORTANT

If your organization uses service control policies (SCPs) to manage permissions in its accounts, and some of the permissions required for an operation are forbidden by these SCPs, Veeam Backup for AWS will not be able to perform the operation even if you grant the permissions to the selected IAM role. For more information on SCPs, see [AWS Documentation](#).

Checking IAM Role Permissions Using Wizard Functionality

To check permissions of an IAM role specified to perform an operation, navigate to the step of the wizard at which you have selected the role, and click **Check Permissions**. Veeam Backup for AWS will display the **Permission check** window where you can track the progress and view the results of the check. If some permissions of the IAM role are missing, the check will complete with errors, and the **Missing Permissions** column will display the list of permissions that must be granted to the IAM role. You can grant the missing permissions to the role [using the AWS Management Console](#) or instruct Veeam Backup for AWS to do it.

TIP

To download the full list of missing permissions as a single JSON policy document that you can use to grant the permissions to the role in the AWS Management Console, click **Export Missing Permissions**.

To let Veeam Backup for AWS grant the missing permissions:

1. In the **Permission check** window, click **Grant**.
2. In the **Grant Permissions** window, provide [one-time access keys of an IAM user](#) that is authorized to update permissions of IAM roles, and then click **Apply**.

The IAM user must have the following permissions:

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "Statement1",
      "Effect": "Allow",
      "Action": [
        "iam:AttachRolePolicy",
        "iam:CreatePolicy",
        "iam:CreatePolicyVersion",
        "iam:CreateRole",
        "iam:GetAccountSummary",
        "iam:GetPolicy",
        "iam:GetPolicyVersion",
        "iam:GetRole",
        "iam:ListAttachedRolePolicies",
        "iam:ListPolicyVersions",
        "iam:SimulatePrincipalPolicy",
        "iam:UpdateAssumeRolePolicy",
        "iam:GetInstanceProfile",
        "iam:CreateInstanceProfile",
        "iam:AddRoleToInstanceProfile",
        "iam:PassRole",
        "iam:ListInstanceProfilesForRole"
      ],
      "Resource": [
        "*"
      ]
    }
  ]
}

```

NOTE

Veeam Backup for AWS does not store one-time access keys in the configuration database.

3. To make sure that the missing permissions have been granted successfully, click **Recheck**.

Checking IAM role Permissions Using IAM Role Tab

If you are not sure whether an IAM role is currently used to perform any operations and if you want to check permission for this IAM role, you can use the permission check at the **IAM Roles** tab. The permission check verifies whether the IAM role has all the permissions required to perform operations with the workloads selected at the **Permissions** step of the **Add IAM Role** wizard.

To run the permission check for an IAM role, do the following:

1. Switch to the **Configuration** page.
2. Navigate to **Accounts > IAM Roles**.
3. Select the necessary IAM role and click **Check Permissions**.

You can track the progress and view the results of the permission check in the **AWS Permission Check** window. If some of the IAM role permissions are missing, the check will complete with errors, and the **Missing Permissions** column will display the list of permissions that must be granted to the IAM role. You can grant the missing permissions to the IAM role [using the AWS Management Console](#) or instruct Veeam Backup for AWS to do it.

TIPS

To download the full list of missing permissions as a single JSON policy document that you can use to grant the permissions to the role in the AWS Management Console, click **Export Missing Permissions**.

To view the configured IAM role permissions at the **IAM Roles** tab, select the necessary IAM role and click **View Permissions**.

To let Veeam Backup for AWS grant the missing permissions:

1. In the **Permission check** window, click **Grant**.
2. In the **Grant Permissions** window, provide [one-time access keys of an IAM user](#) that is authorized to update permissions of IAM roles, and then click **Apply**.

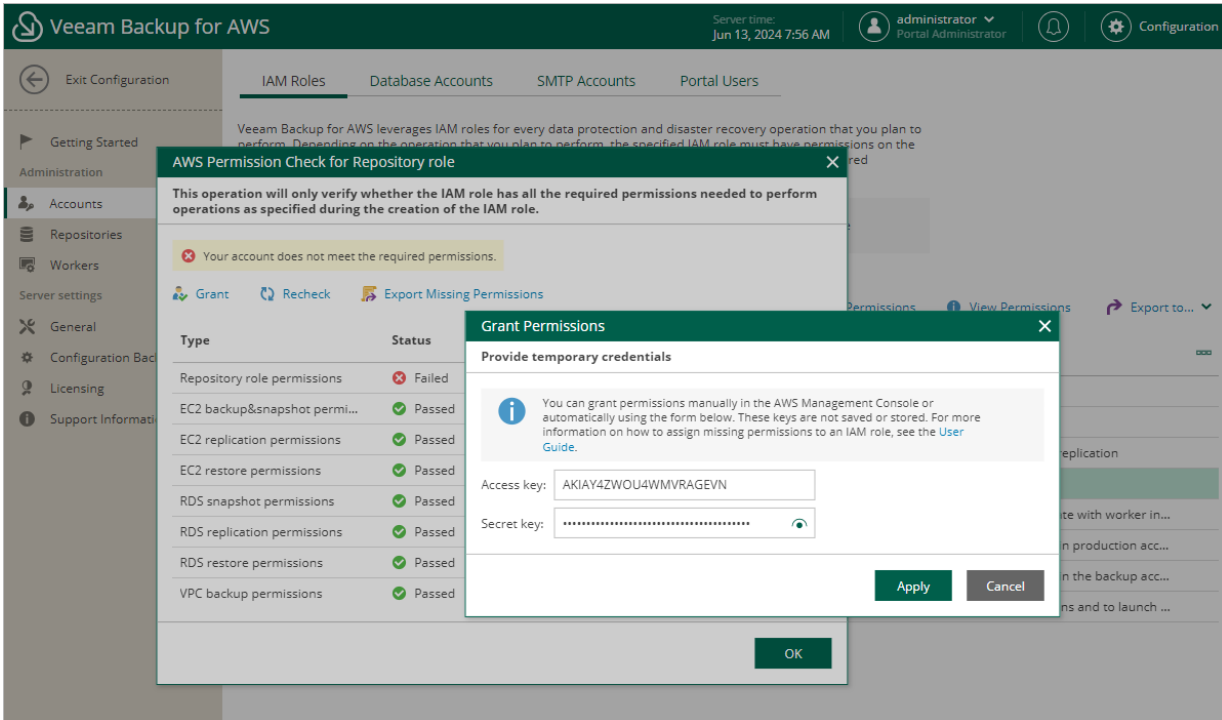
The IAM user must have the following permissions:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "Statement1",
      "Effect": "Allow",
      "Action": [
        "iam:AttachRolePolicy",
        "iam:CreatePolicy",
        "iam:CreatePolicyVersion",
        "iam:CreateRole",
        "iam:GetAccountSummary",
        "iam:GetPolicy",
        "iam:GetPolicyVersion",
        "iam:GetRole",
        "iam:ListAttachedRolePolicies",
        "iam:ListPolicyVersions",
        "iam:SimulatePrincipalPolicy",
        "iam:UpdateAssumeRolePolicy",
        "iam:GetInstanceProfile",
        "iam:CreateInstanceProfile",
        "iam:AddRoleToInstanceProfile",
        "iam:PassRole",
        "iam:ListInstanceProfilesForRole"
      ],
      "Resource": [
        "*"
      ]
    }
  ]
}
```


NOTE

Veeam Backup for AWS does not store one-time access keys in the configuration database.

3. To make sure that the missing permissions have been granted successfully, click **Recheck**.



The screenshot shows the Veeam Backup for AWS configuration interface. The main window displays the 'IAM Roles' tab. A dialog box titled 'AWS Permission Check for Repository role' is open, showing a table of permissions and their status. The 'Repository role permissions' row is marked as 'Failed', while all other permissions are marked as 'Passed'. Below the table, there are buttons for 'Grant', 'Recheck', and 'Export Missing Permissions'. A sub-dialog box titled 'Grant Permissions' is also open, prompting for temporary credentials. It includes an information message, an 'Access key' field with the value 'AKIAY4ZWOU4WMVRAGEVN', and a 'Secret key' field with masked characters. There are 'Apply' and 'Cancel' buttons at the bottom of the sub-dialog.

Type	Status
Repository role permissions	Failed
EC2 backup&snapshot permi...	Passed
EC2 replication permissions	Passed
EC2 restore permissions	Passed
RDS snapshot permissions	Passed
RDS replication permissions	Passed
RDS restore permissions	Passed
VPC backup permissions	Passed

Removing IAM Roles

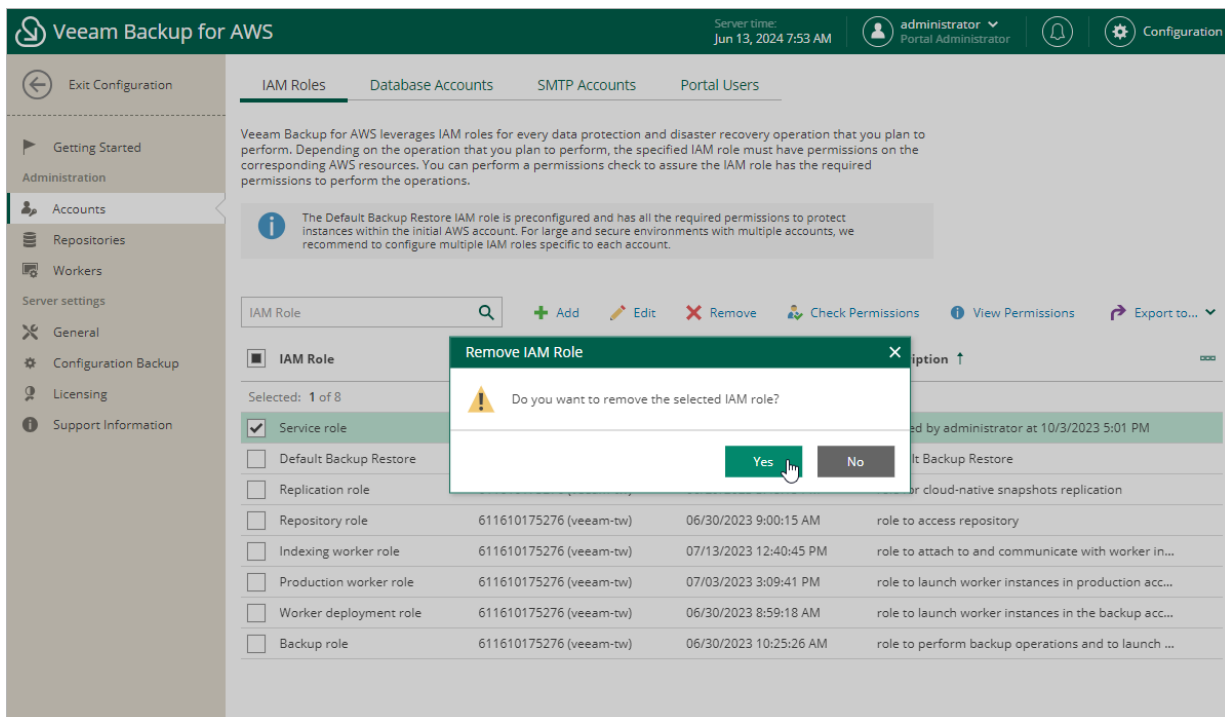
You can remove an IAM role from Veeam Backup for AWS if it is no longer used to perform data protection and disaster recovery operations.

IMPORTANT

You cannot remove an IAM role that is used to access backup repositories or is specified in the settings of any configured backup policy.

To remove an IAM role, do the following:

1. Switch to the **Configuration** page.
2. Navigate to **Accounts > IAM Roles**.
3. Select the IAM role and click **Remove**.
3. In the **Remove IAM Role** window, click **Yes** to acknowledge the operation.



Managing User Accounts

Veeam Backup for AWS controls access to its functionality with the help of user roles. A role defines what operations users can perform and what range of data is available to them in Veeam Backup for AWS.

There are 3 user roles that you can assign to users working with Veeam Backup for AWS. Actions a user can perform depend on the role.

- **Portal Administrator** – can perform all configuration actions and can also act as a Portal Operator and Restore Operator.
- **Portal Operator** – can create and manage backup policies, manage the protected data, perform all restore operations and view session statistics.
- **Restore Operator** – can only perform restore operations and view session statistics.

IMPORTANT

The list of portal users may display user accounts with the *Company Administrator* role assigned – these accounts are intended to be used for the integration of Veeam Backup for AWS and Veeam Service Provider Console, and are created using the [Veeam Service Provider Console plug-in](#). It is not recommended that you perform any actions with these users.

The following table describes the functionality available to users with different roles in the Veeam Backup for AWS UI.

Tab	Functionality	Portal Administrator	Portal Operator	Restore Operator
Overview	Dashboard	Full	Full	N/A
Resources	Infrastructure	Full	Full	N/A
Policies	Backup policies	Full	Full	N/A
Protected Data	Restore	Full	Full	Full
	File-level recovery	Full	Full	Full
	Remove	Full	Full	N/A
Session Log	Session log	Full	Full	Full
	Stop session execution	Full	Full	N/A
Configuration				

Tab	Functionality	Portal Administrator	Portal Operator	Restore Operator
Accounts	IAM roles, SMTP accounts, Portal Users	Full	N/A	N/A
Repositories	Backup repositories	Full	N/A	N/A
Workers	Worker instances	Full	N/A	N/A
Settings	General settings	Full	N/A	N/A
Licensing	Licensing	Full	N/A	N/A
Support Information	Updates and logs	Full	N/A	N/A

Adding User Accounts

To manage access to Veeam Backup for AWS, you can create local user accounts or add user accounts of your identity provider.

To add a new user account, do the following:

1. [Launch the Add Account wizard.](#)
2. [Choose an account type.](#)
3. [Specify an account name and description.](#)
4. [Specify general settings.](#)
5. [Finish working with the wizard.](#)

Step 1. Launch Add Account Wizard

To launch the **Add Account** wizard, do the following:

1. Switch to the **Configuration** page.
2. Navigate to **Accounts > Portal Users**.
3. Click **Add**.

The screenshot shows the Veeam Backup for AWS Configuration page. The top navigation bar includes the Veeam logo, the text 'Veeam Backup for AWS', the server time 'Jun 24, 2024 12:18 PM', and the user 'administrator Portal Administrator'. The left sidebar contains navigation options: 'Exit Configuration', 'Getting Started', 'Administration', 'Accounts', 'Repositories', 'Workers', 'Server settings', 'General', 'Configuration Backup', 'Licensing', and 'Support Information'. The main content area is titled 'Portal Users' and contains a sub-header: 'To control access to Veeam Backup for AWS, you can create portal users. Depending on the assigned role, portal users can perform specific activities — configure Veeam Backup for AWS settings, create instance backups, restore backed-up data and so on.' Below this is a search bar for 'Account' and a 'Filter (None)' button. Action buttons include '+ Add', 'Edit', 'Remove', 'Enable MFA', 'Disable MFA', 'Change Password', and 'Export to...'. A table lists the existing portal users:

Account ↑	Role	Type	Description	MFA Enabled
administrator	Portal Administrator	Veeam Backup for AWS	Created by bd at 5/14/202...	No
backup-operator	Portal Operator	Veeam Backup for AWS	Created by bd at 5/6/2024 ...	No
bd	Portal Administrator	Veeam Backup for AWS	'Veeam Backup for AWS m...	No
bd1	Portal Administrator	Veeam Backup for AWS		No
restore	Restore Operator	Veeam Backup for AWS	Created by bd at 5/6/2024 ...	No

Step 2. Choose Account Type

At the **Account Type** step of the wizard, choose whether you want to create a new Veeam Backup for AWS user or to retrieve a user identity from your identity provider. To retrieve user identities from the identity provider, you must first [configure single sign-on settings](#).

The screenshot shows the 'Add Account' wizard in the Veeam Backup for AWS interface. The top navigation bar includes the Veeam logo, the title 'Veeam Backup for AWS', the server time 'Jun 24, 2024 12:18 PM', the user 'administrator Portal Administrator', and a 'Configuration' icon. The main content area is titled 'Add Account' and features a sidebar with navigation options: 'Account Type' (selected), 'Account Info', 'General Settings', and 'Summary'. The 'Specify account type' section contains the instruction: 'Specify whether you want to create an account within the Veeam Backup appliance or leverage an identity from an Identity Provider.' Below this, there are two radio button options: 'Veeam Backup account' (which is selected) and 'Identity Provider account'. At the bottom of the wizard, there are 'Next' and 'Cancel' buttons.

Step 3. Specify Account Name and Description

At the **Account Info** step of the wizard, enter a name for the new user account and provide a description for future reference.

You cannot use *admin* as the account name. The maximum length of the name is 32 characters for the Veeam Backup for AWS user and 125 characters for the user identity from your identity provider. The following characters are supported: lowercase Latin letters, numeric characters, underscores and dashes; the dollar sign (\$) is supported but only if it the last character of the name.

IMPORTANT

If you have selected the **Identity Provider account** option at step 1, the name specified for a user account must match the value of an attribute that the identity provider will send to Veeam Backup for AWS to authenticate the user. For more information, see [Configuring SSO Settings](#).

The screenshot shows the 'Add Account' wizard in Veeam Backup for AWS. The interface is divided into a top navigation bar, a left sidebar, and a main content area. The top bar includes the Veeam logo, 'Veeam Backup for AWS', server time 'Jun 24, 2024 12:19 PM', and user information 'administrator Portal Administrator'. The sidebar contains navigation options: 'Account Type', 'Account Info' (selected), 'General Settings', and 'Summary'. The main content area is titled 'Specify account name and description' and includes the instruction 'Enter a name and description for the user account.' There are two input fields: 'Name:' with the value 'donna_ortiz' and 'Description:' with the value 'Created by administrator at 6/24/2024 12:19 PM'. At the bottom, there are three buttons: 'Previous', 'Next', and 'Cancel'.

Step 4. Specify General Settings

At the **General Settings** step of the wizard, select a role for the user account. For more information on user roles, see [Managing User Accounts](#).

If you have selected the **Veeam Backup for AWS account** option at step 2, specify a password for the new Veeam Backup for AWS user account.

The screenshot shows the 'Add Account' wizard in the Veeam Backup for AWS interface. The top navigation bar includes the Veeam logo, 'Veeam Backup for AWS', the server time 'Jun 24, 2024 12:19 PM', and the user 'administrator Portal Administrator'. The main content area is titled 'Add Account' and has a left sidebar with navigation options: 'Account Type', 'Account Info', 'General Settings' (which is highlighted), and 'Summary'. The 'Specify account settings' section contains a 'Role' section with a dropdown menu set to 'Portal Operator'. Below this is a 'Password' section with two input fields: 'Password' and 'Repeat password', both masked with dots. A blue information icon is present next to the password fields. A tooltip at the bottom of the password section states: 'Password should be 8 characters minimum with one digit, one uppercase and one lowercase. Monotonic sequences such as 1234 are not allowed.' At the bottom of the wizard, there are three buttons: 'Previous', 'Next' (highlighted in green), and 'Cancel'.

Step 5. Finish Working with Wizard

At the **Summary** step of the wizard, review summary information and click **Finish**.

The screenshot shows the 'Add Account' wizard in the Veeam Backup for AWS interface. The top navigation bar includes the Veeam logo, the product name 'Veeam Backup for AWS', the server time 'Jun 24, 2024 12:21 PM', the user 'administrator Portal Administrator', and a 'Configuration' link. The main content area is titled 'Add Account' and features a left-hand navigation menu with options: 'Account Type', 'Account Info', 'General Settings', and 'Summary'. The 'Summary' step is active, displaying the following account details:

Summary	
Review the configured settings, and click Finish to exit the wizard.	
Account	
Type:	Veeam Backup account
Name:	donna_ortiz
Description:	Created by administrator at 6/24/2024 12:19 PM
Role:	Portal Operator
Password:	*****

At the bottom of the wizard, there are three buttons: 'Previous', 'Finish', and 'Cancel'.

Editing User Account Settings

For each user account added to the Veeam Backup for AWS configuration database, you can modify the settings of the account:

1. Switch to the **Configuration** page.
2. Navigate to **Accounts > Portal Users**.
3. Select the user account and click **Edit**.
4. Complete the **Edit Account** wizard.
 - a. To provide a new name and description for the user account, follow the instructions provided in section [Adding User Accounts](#) (step 3).
 - b. To choose a new role for the user account, follow the instructions provided in section [Adding User Accounts](#) (step 4).
 - c. At the **Summary** step of the wizard, review summary information and click **Finish** to confirm the changes.

The screenshot displays the Veeam Backup for AWS configuration interface. The top navigation bar includes the Veeam logo, the text 'Veeam Backup for AWS', the server time 'Jun 24, 2024 12:26 PM', and the user 'administrator Portal Administrator'. The left sidebar shows navigation options: Exit Configuration, Getting Started, Administration, Accounts (selected), Repositories, Workers, Server settings, General, Configuration Backup, Licensing, and Support Information. The main content area is titled 'Portal Users' and contains a search bar, a filter dropdown, and a table of user accounts. The table has columns for Account, Role, Type, Description, and MFA Enabled. The 'donna_ortiz' account is highlighted in green.

Account ↑	Role	Type	Description	MFA Enabled
administrator	Portal Administrator	Veeam Backup for AWS	Created by bd at 5/14/202...	No
backup-operator	Portal Operator	Veeam Backup for AWS	Created by bd at 5/6/2024 ...	No
bd	Portal Administrator	Veeam Backup for AWS	"Veeam Backup for AWS m...	No
bd1	Portal Administrator	Veeam Backup for AWS		No
donna_ortiz	Portal Operator	Veeam Backup for AWS	Created by administrator a...	No
restore	Restore Operator	Veeam Backup for AWS	Created by bd at 5/6/2024 ...	No

Changing User Passwords

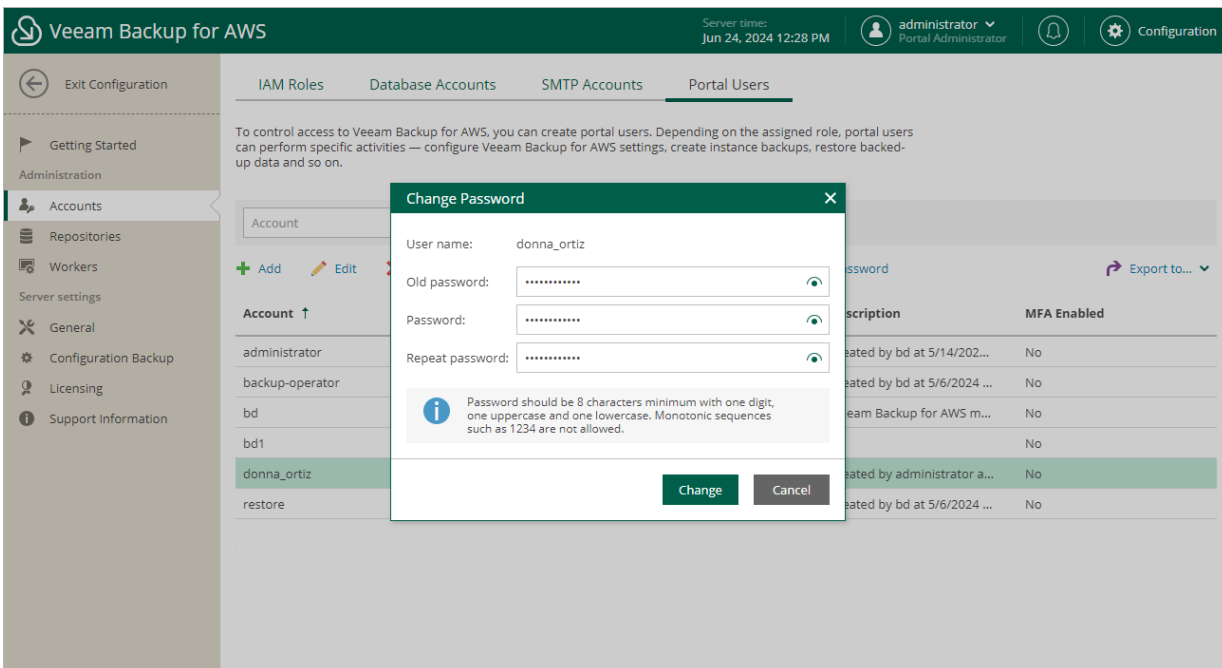
For Veeam Backup for AWS user accounts, you can change the password specified while creating the account:

IMPORTANT

Consider the following:

- You cannot change the password for a user account whose user identity was obtained from an identity provider.
- If your backup appliance is managed by a Veeam Backup & Replication server and you change the password of a user whose credentials Veeam Backup & Replication uses to connect to the backup appliance, you must also change this user password in the Veeam Backup & Replication console as described in the Veeam Backup & Replication User Guide, section [Editing and Deleting Credentials Records](#). Otherwise, the connection will not be established.

1. Switch to the **Configuration** page.
2. Navigate to **Accounts > Portal Users**.
3. Select the user account and click **Change Password**.
4. In the **Change Password** window, enter the currently used password, enter and confirm a new password, and then click **Change**.



Enabling Multi-Factor Authentication

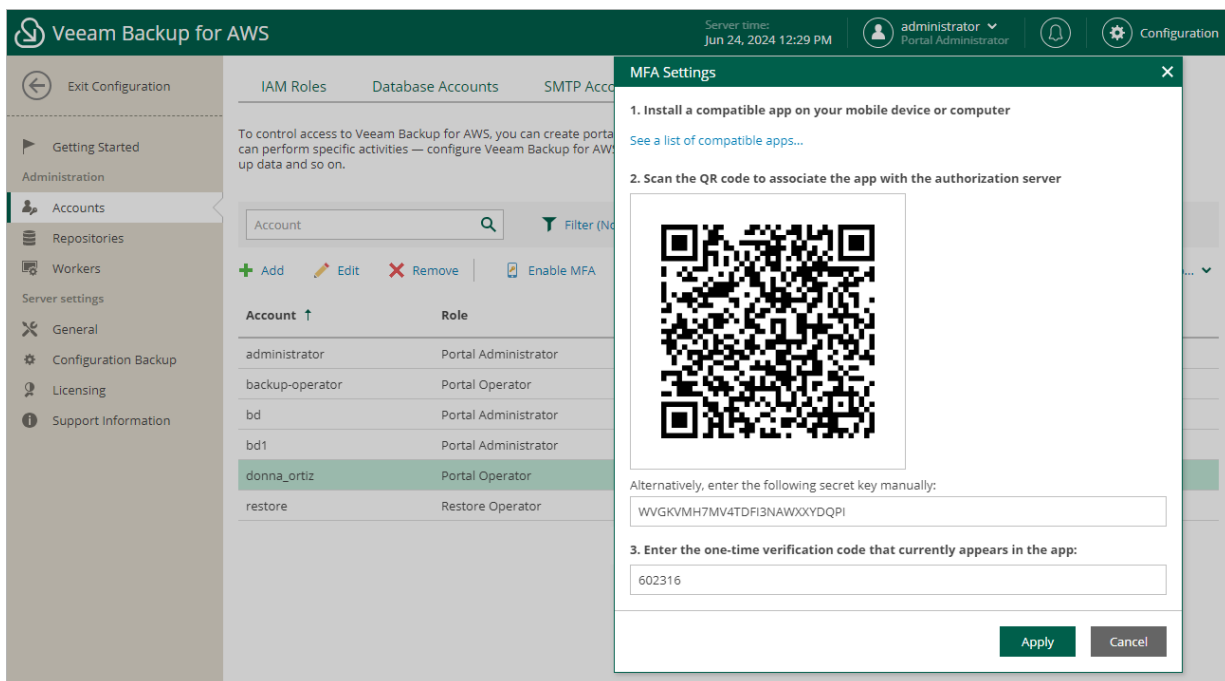
Multi-factor authentication (MFA) in Veeam Backup for AWS is based on the Time-based One-Time Password (TOTP) method that requires users to verify their identity by providing a temporary six-digit code sent by an authentication application to a trusted device.

To enable MFA for a user account, do the following:

1. Switch to the **Configuration** page.
2. Navigate to **Accounts > Portal Users**.
2. Select the user account and click **Enable MFA**.
3. Follow the instructions provided in the **MFA Settings** window:
 - a. Install an authentication application on a trusted device.
You can use any application that supports the TOTP protocol.
 - b. To associate the authentication application with the authorization server, scan the displayed QR code using the camera of the trusted device.
 - c. Enter a verification code generated by the authentication application.
 - d. Click **Apply**.

IMPORTANT

You cannot enable MFA for a user account whose user identity was obtained from an identity provider.



The screenshot shows the Veeam Backup for AWS Configuration page. The 'Accounts' section is active, displaying a table of user accounts. The 'donna_ortiz' account is selected, and the 'Enable MFA' button is visible. An 'MFA Settings' dialog box is open, showing instructions for enabling MFA. The dialog box includes a QR code for scanning and a field for entering a secret key manually. The verification code '602316' is entered in the 'Enter the one-time verification code' field.

Account	Role
administrator	Portal Administrator
backup-operator	Portal Operator
bd	Portal Administrator
bd1	Portal Administrator
donna_ortiz	Portal Operator
restore	Restore Operator

MFA Settings

1. Install a compatible app on your mobile device or computer
[See a list of compatible apps...](#)
2. Scan the QR code to associate the app with the authorization server

Alternatively, enter the following secret key manually:

3. Enter the one-time verification code that currently appears in the app:

Apply **Cancel**

Managing Database Accounts

To allow Veeam Backup for AWS to authenticate against PostgreSQL DB instances protected by backup policies, you must specify credentials of database accounts that will be used to access the databases when performing [image-level backup](#) and [restore operations](#).

NOTE

After you upgrade Veeam Backup for AWS to version 8.0, the list of database accounts will be populated with policy credentials that were previously used to access protected databases.

Adding Database Accounts


To add a new database account, do the following:

1. [Launch the Add Account wizard.](#)
2. [Specify an account name and description.](#)
3. [Specify general settings.](#)
4. [Finish working with the wizard.](#)

Step 1. Launch Add Account Wizard

To launch the **Add Account** wizard, do the following:

1. Switch to the **Configuration** page.
2. Navigate to **Accounts > Database Accounts**.
3. Click **Add**.



Step 2. Specify Account Name and Description

At the **Account Info** step of the wizard, use the Name and Description fields to enter a name for the database account and to provide a description for future reference. The specified name and description will be displayed on the **Database Accounts** tab.

The name must be unique in Veeam Backup for AWS; the maximum length of the name is 127 characters; the maximum length of the description is 1024 characters.

The screenshot shows the 'Add Account' wizard in Veeam Backup for AWS. The interface includes a top navigation bar with the Veeam logo, 'Veeam Backup for AWS', server time 'Jun 13, 2024 8:05 AM', user 'administrator Portal Administrator', and a 'Configuration' link. The main content area is titled 'Add Account' and features a left-hand navigation pane with 'Account Info' (selected), 'General Settings', and 'Summary'. The 'Specify account name and description' step is active, with the instruction 'Enter a name and description for the database account'. The 'Name' field contains 'db-instance-01' and the 'Description' field contains 'account to access db-instance-01 databases'. At the bottom right, there are 'Next' and 'Cancel' buttons.

Step 3. Specify General Settings

At the **General Settings** step of the wizard, specify credentials that the account will use to access databases protected by backup policies.

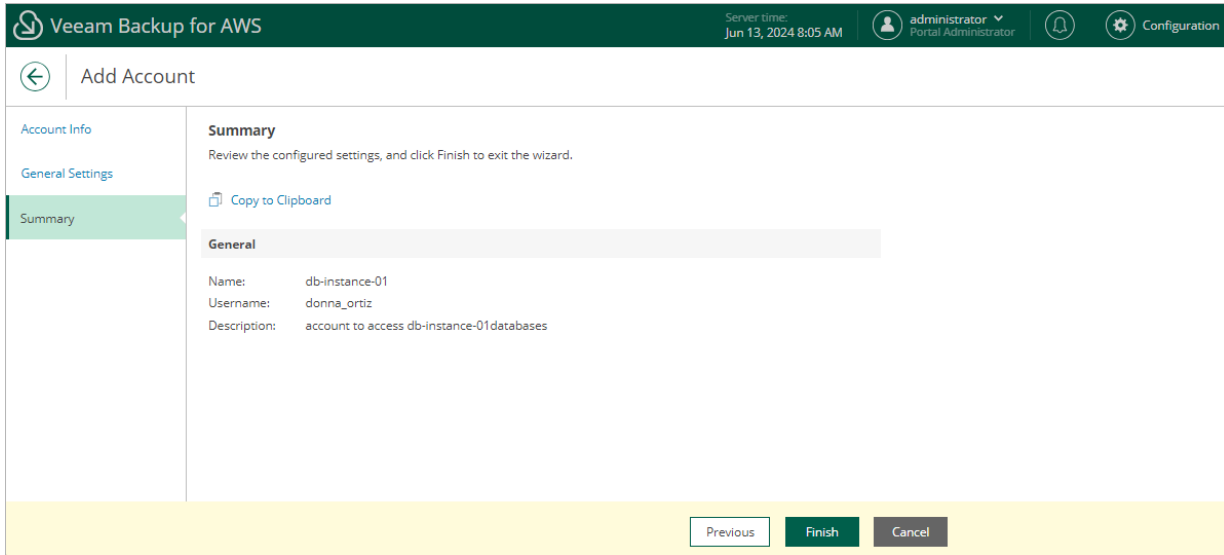
The screenshot shows the 'Add Account' wizard in Veeam Backup for AWS. The interface includes a top navigation bar with the Veeam logo, server time (Jun 13, 2024 8:05 AM), and user information (administrator, Portal Administrator). The main content area is titled 'Add Account' and features a sidebar with three tabs: 'Account Info', 'General Settings' (which is selected and highlighted in green), and 'Summary'. The 'General Settings' section is titled 'Specify account username and password' and contains two input fields: 'Username' with the value 'donna_ortiz' and 'Password' with masked characters '*****'. A 'Previous' button is disabled, while 'Next' and 'Cancel' buttons are active. The bottom of the wizard has a yellow background.

Step 4. Finish Working with Wizard

At the **Summary** step of the wizard, review summary information and click **Finish**.

TIP


After you add the database account, you will be able to specify this account while creating backup policies to allow Veeam Backup for AWS to access databases. For more information, see [Performing RDS Backup](#).



Editing Database Accounts

For each database account added to the Veeam Backup for AWS configuration database, you can modify the settings of the account:

1. Switch to the **Configuration** page.
2. Navigate to **Accounts > Database Accounts**.
3. Select the account and click **Edit**.
4. Complete the **Edit Account** wizard.
 - a. To specify a new name and description for the account, follow the instructions provided in section [Adding Database Accounts](#) (step 2).
 - b. To modify the credentials that are used to access databases added to backup policies, follow the instructions provided in section [Adding Database Accounts](#) (step 3).
 - c. At the **Summary** step of the wizard, review summary information and click **Finish** to confirm the changes.



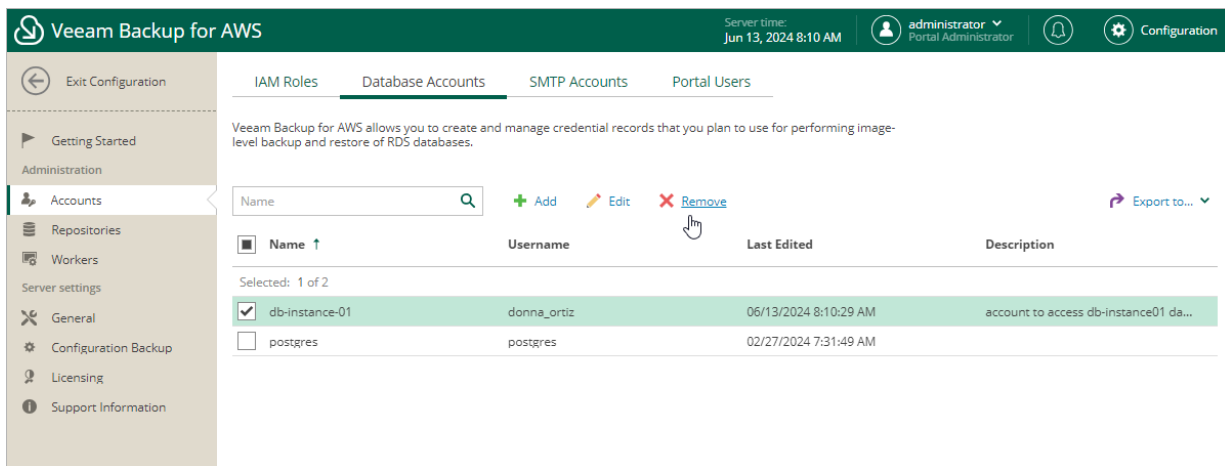
Removing Database Accounts

Veeam Backup for AWS allows you to permanently remove a database account from the configuration database if you no longer need it:

1. Switch to the **Configuration** page.
2. Navigate to **Accounts > Database Accounts**.
3. Select the account and click **Remove**.

IMPORTANT

You cannot remove a database account that is associated with any backup policy. Delete all of the affected policies or [edit their settings](#) — and then try removing the account again.



The screenshot shows the Veeam Backup for AWS Configuration page. The top navigation bar includes the Veeam logo, server time (Jun 13, 2024 8:10 AM), and user information (administrator, Portal Administrator). The left sidebar contains navigation options: Exit Configuration, Getting Started, Administration, Accounts (selected), Repositories, Workers, Server settings, General, Configuration Backup, Licensing, and Support Information. The main content area is titled 'Database Accounts' and includes a search bar, '+ Add', 'Edit', and 'Remove' buttons. Below the buttons is a table with columns: Name, Username, Last Edited, and Description. The table shows two accounts: 'db-instance-01' (selected) and 'postgres'. The 'db-instance-01' account is highlighted in green, and its 'Remove' button is being clicked by a mouse cursor.

Name	Username	Last Edited	Description
<input checked="" type="checkbox"/> db-instance-01	donna_ortiz	06/13/2024 8:10:29 AM	account to access db-instance01 da...
<input type="checkbox"/> postgres	postgres	02/27/2024 7:31:49 AM	

Managing Worker Instances

To perform most data protection and disaster recovery operations (such as creating and removing EC2 and RDS image-level backups, restoring backed-up data, EFS indexing), Veeam Backup for AWS uses worker instances. Worker instances are temporary Linux-based EC2 instances that are responsible for the interaction between the backup appliance and other Veeam Backup for AWS components. Worker instances process backup workload and distribute backup traffic when transferring data to backup repositories.

Each worker instance is deployed in a specific AWS Region for the duration of the backup, restore and retention process. AWS Regions in which Veeam Backup for AWS deploys worker instances to perform operations are predefined and described in section [Worker Instance Locations](#). However you can choose whether you want Veeam Backup for AWS to deploy worker instances in the backup account or in production AWS accounts, specify network settings and instance types that will be used to deploy worker instances. For more information on AWS accounts in which Veeam Backup for AWS deploys worker instances, see [Worker Deployment Options](#).

NOTE

You can tell worker instances from other EC2 instances running in your environment by their names – all worker instances deployed by Veeam Backup for AWS to perform backup and restore operations have the same name – *VBA_Worker*, all worker instances deployed by Veeam Backup for AWS to perform EFS indexing have the same name – *EFS_Worker*.

Managing Worker Configurations

A configuration is a group of network settings that Veeam Backup for AWS uses to deploy worker instances in a specific AWS Region to perform data protection, disaster recovery, backup retention and EFS indexing operations. Veeam Backup for AWS deploys one worker instance per each AWS resource added to a backup policy, restore, indexing or retention task.

Adding Configurations for Backup Account

By default, Veeam Backup for AWS deploys worker instances used for retention, backup and restore operations in the backup account. You can [choose an IAM role](#) and [specify network settings](#) that will be used to deploy these worker instances.

Specifying IAM Role

Out of the box, Veeam Backup for AWS uses the permissions of the *Default Backup Restore* role to deploy worker instances – the role is preconfigured and has all the required permissions. Therefore, the default backup account is an AWS account to which the backup appliance belongs. However, you can specify another IAM role to change the backup account.

To specify an IAM role for worker instances, do the following:


1. Switch to the **Configuration** page.
2. Navigate to **Workers > Network**.
3. At the **Backup Accounts** tab, click the link in the **Service IAM role** field.
4. In the **Choose IAM Role** window, select the necessary IAM role, and then click **Apply**.

For an IAM role to be displayed in the list of available IAM roles, it must be added to Veeam Backup for AWS as described in section [Adding IAM Roles](#).

IMPORTANT

After you choose an IAM role, it is not recommended to change it. Otherwise, all the created worker configurations will be removed automatically as soon as you choose another IAM role.

After you specify the IAM role, it is recommended that you check whether permissions of the specified IAM role are sufficient to deploy worker instances. For information on how to check IAM role permissions, see [Checking IAM Role Permissions](#). To learn what permissions must have the IAM role used to deploy worker instances, see [Worker Deployment Role Permissions in Backup Account](#).



Adding Worker Configurations

To deploy worker instances in the *Backup* account, Veeam Backup for AWS uses the default network settings of AWS Regions (if any). However, to optimize infrastructure costs and to ensure better performance of backup, retention and restore processes, you can add worker configurations to specify network settings for each Availability Zone in which worker instances will be deployed.

To add a worker configuration:

1. In the **Worker configurations** section, click **Add**.
2. Complete the **Add Worker Configuration** wizard.
 - a. At the **General** step of the wizard, select an AWS Region and Availability Zone for which you want to configure network settings.

If you create the worker configuration that will be used to perform EC2 backup operations, you can select any Availability Zone in the specified AWS Region. Veeam Backup for AWS will still be able to perform the operations even if the selected zone will differ from the Availability Zone where the processed EC2 instances reside.
 - b. At the **Network** step of the wizard, select an Amazon VPC and a subnet to which you want to connect worker instances, and specify a security group that must be associated with the instances. For an Amazon VPC, a subnet and a security group to be displayed in the lists of available network specifications, they must be created in AWS as described in [AWS Documentation](#).

Veeam Backup for AWS will apply the specified network settings to all worker instances that will be deployed in the AWS Region and Availability Zone selected at the **General** step of the wizard.

IMPORTANT

When selecting a subnet and security group, consider the following:

- Security rules configured in the selected security group must allow direct network traffic required to communicate with [AWS services](#). To learn how to add rules to security groups, see [AWS Documentation](#).

Proxy redirect and setting a proxy in the Veeam Backup for AWS configuration are not supported.

- If you select an Outpost subnet, backup and restore operations in the AWS Region to which the AWS Outpost is connected may fail to complete successfully. The issue occurs if the default worker instance type is not supported for the AWS Outpost. To work around the issue, change the default worker profiles as described in section [Managing Worker Profiles](#).

By default, Veeam Backup for AWS uses public access to communicate with worker instances. That is why the public IPv4 addressing attribute must be enabled for the selected subnet, the selected VPC must have an internet gateway attached, and the VPC and subnet route tables must have routes that direct internet-bound traffic to this internet gateway. If you want worker instances to operate in a private network, do either of the following:

- Enable public IPv4 addressing for the subnet as described in [AWS Documentation](#).
- Enable the private network deployment functionality, and configure specific VPC endpoints for the subnet to let Veeam Backup for AWS use private IPv4 addresses as described in section [Configuring Private Network Deployment](#).

For the list of specific endpoints required to perform backup and restore operations, see [Configuring Private Networks](#).

- Configure VPC endpoints as described in section [Appendix C. Configuring Endpoints in AWS](#).

c. At the **Summary** step of the wizard, review summary information and click **Finish**.

The screenshot shows the 'Add Worker Configuration' wizard in Veeam Backup for AWS, specifically the 'Summary' step. The interface includes a top navigation bar with the Veeam logo, server time (Oct 12, 2023 2:40 PM), user information (administrator, Portal Administrator), and a 'Configuration' button. The main content area is divided into a left sidebar with 'General', 'Network', and 'Summary' tabs, and a main panel titled 'Review configured settings'. The 'Summary' tab is active, showing a summary of the configured settings. The 'General' section lists 'Region: Europe (Paris)' and 'Availability zone: eu-west-3c'. The 'Network' section lists 'VPC: dept-01-amroz-srv07 VPC (172.28.0.0/16)', 'Subnet: 172.28.0.0/20 (eu-west-3c)', and 'Security group: veeamsecuritygroup'. The 'Verification checks' section shows 'Auto-assign public IPv4' with a green checkmark and 'OK'. At the bottom, there are three buttons: 'Previous', 'Finish', and 'Cancel'.

Section	Property	Value
General	Region	Europe (Paris)
	Availability zone	eu-west-3c
Network	VPC	dept-01-amroz-srv07 VPC (172.28.0.0/16)
	Subnet	172.28.0.0/20 (eu-west-3c)
	Security group	veeamsecuritygroup
Verification checks	Auto-assign public IPv4	✓ OK

Testing Configurations for FLR

When performing file-level recovery for an EC2 instance, Veeam Backup for AWS deploys a worker instance, attaches and mounts EBS volumes of the EC2 instance to the worker instance and launches file-level recovery browser to allow users to browse, download and restore files and folders. To make sure whether worker network settings are configured properly, and the file-level recovery browser is accessible from the your local machine, it is recommended that you run a file-level recovery test before you start file-level recovery operations in an AWS Region.

To run the file-level recovery test for a specific region, do the following:

1. Switch to the **Configuration** page.
2. Navigate to **Workers > Network**.
3. In the **Worker configurations** section, select the necessary configuration, and then click **Test FLR**.
4. Wait until the status of the file-level recovery test in the **FLR Status** column changes to *Running*, and then click the status.

Veeam Backup for AWS will display the **FLR Test Log** window where you can track the progress and view the results of the test.

5. If network settings are configured properly for the AWS Region, Veeam Backup for AWS will deploy the worker instance and display the link to the file-level recovery browser in the **FLR Test Log** window.
 - a. To check that you can access the file-level recovery browser, click the displayed link.

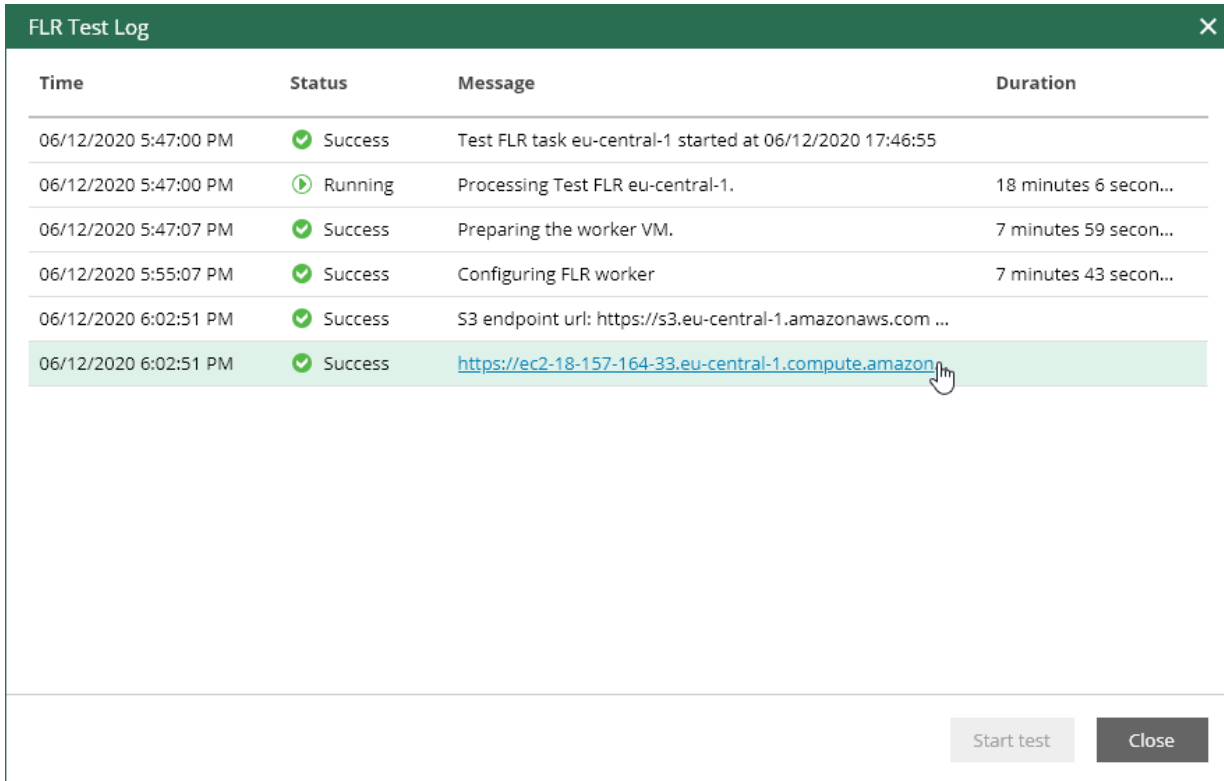
Note that the security group associated with worker instances must allow inbound internet access from the machine from which you plan to open the file-level recovery browser.

- b. To finish the file-level recovery test, click **End Test** in the file-level recovery browser.

If you do not click **End Test** within 30 minutes after Veeam Backup for AWS displays the link to the file-level recovery browser, the file-level recovery test will finish automatically with the *Warning* status.

TIP

If the file-level recovery test finishes with the *Warning* or *Error* status, you can run the test again after fixing issues with the network settings. To do that, select the necessary configuration in the **Worker configurations** section, and then click **Test FLR**.



Time	Status	Message	Duration
06/12/2020 5:47:00 PM	Success	Test FLR task eu-central-1 started at 06/12/2020 17:46:55	
06/12/2020 5:47:00 PM	Running	Processing Test FLR eu-central-1.	18 minutes 6 secon...
06/12/2020 5:47:07 PM	Success	Preparing the worker VM.	7 minutes 59 secon...
06/12/2020 5:55:07 PM	Success	Configuring FLR worker	7 minutes 43 secon...
06/12/2020 6:02:51 PM	Success	S3 endpoint url: https://s3.eu-central-1.amazonaws.com ...	
06/12/2020 6:02:51 PM	Success	https://ec2-18-157-164-33.eu-central-1.compute.amazon...	

Start test Close

Adding Configurations for Production Accounts

To perform EFS indexing operations, as well as RDS backup and restore operations, worker instances are deployed in production accounts by default. However, if you also want Veeam Backup for AWS to deploy worker instances in production accounts for backup and restore operations performed for EC2 instances (for example, to restore instances from cloud-native snapshots encrypted using default AWS managed keys), you must configure the backup policy and restore settings.

Specifying IAM Roles

To deploy worker instances in production accounts, Veeam Backup for AWS employs the following IAM roles:

- An IAM role that is used to retrieve network settings of AWS Regions in a production account when adding new or editing existing working configurations. The role must be assigned permissions listed in section [Worker Configuration IAM Role Permissions](#).

You must specify this IAM role in the **Add Worker Configuration** wizard as described in section [Adding Worker Configurations](#).

- An IAM role that is used to perform a backup or restore operation. Veeam Backup for AWS also uses this role to deploy worker instances in a production account. That is why the role must be assigned additional permissions listed in section [EFS Backup IAM Role Permissions](#), [EC2 Backup IAM Role Permissions](#), [EC2 Restore IAM Permissions](#) or [RDS Backup IAM Role Permissions](#).

You must specify this IAM role in the backup policy or restore settings as described in section [Creating EFS Backup Policies](#), [Creating EC2 Backup Policies](#), [Performing RDS Backup](#), [Performing Entire EC2 Instance Restore](#), [Performing Volume-Level Restore](#) or [Performing RDS Database Restore](#).

- An IAM role that is attached to the deployed worker instances and further used by Veeam Backup for AWS to communicate with the instances. The role must be assigned permissions listed in section [Worker Deployment Role Permissions in Production Accounts](#) or [FLR Worker IAM Role Permissions](#).

You must specify this IAM role when enabling worker deployment in production accounts in the backup policy or restore settings as described in section [Creating EFS Backup Policies](#), [Creating EC2 Backup Policies](#), [Creating RDS Backup Policies](#), [Performing Entire EC2 Instance Restore](#), [Performing Volume-Level Restore](#), [Performing File-Level Recovery](#) or [Performing RDS Database Restore](#).

NOTE

Since you do not specify an IAM role for file-level recovery operations, the role that you specify when enabling worker deployment in production accounts in the restore settings is also used by Veeam Backup for AWS to deploy worker instances.

Adding Worker Configurations

To deploy worker instances in production accounts, Veeam Backup for AWS automatically chooses the most appropriate network settings of AWS Regions (for example, specifies a VPC as a mount target for the processed file system) when performing EFS indexing operations, and uses the default network settings of AWS Regions (if any) when performing EC2 backup and restore operations. However, you can add worker configurations to specify network settings for each region in which worker instances will be deployed. You can add multiple worker configurations with different network settings per AWS Region.

To add a worker configuration:

1. Switch to the **Configuration** page.
2. Navigate to **Workers > Network**.
3. Switch to the **Production Accounts** tab.
4. In the **Worker configurations** section, click **Add**.
5. Complete the **Add Worker Configuration** wizard.
 - a. At the **General** step of the wizard, do the following:
 - i. In the **Account** section, select an AWS account to which resources that you plan to process belong and specify an IAM role that will be used to access and list region network settings in the selected AWS account. The role must be granted the permissions listed in section [Worker Configuration IAM Role Permissions](#).

For an IAM role to be displayed in the IAM role list, it must be added to Veeam Backup for AWS as described in section [Adding IAM Roles](#). If you have not added the necessary IAM role to Veeam Backup for AWS beforehand, you can do it without closing the **Add Worker Configuration** wizard. To do that, click **Add** and complete the **Add IAM Role** wizard.

NOTE

Consider the following:

- After you specify the IAM role, it is recommended that you check whether permissions of the specified IAM role are sufficient to access and list region network settings in the selected AWS account. For information on how to check IAM role permissions, see [Checking IAM Role Permissions](#).
- The selected IAM role will be used only to populate network settings for the **Add Worker Configuration** wizard. IAM roles whose permissions Veeam Backup for AWS will use to configure the specified settings when deploying worker instances will be specified in the backup policy and restore settings.

- ii. In the **Region** section, select an AWS Region and Availability Zone in which AWS resources that you plan to process reside.

Note that for restore operations, the configuration must be created in the same Availability Zone where the restored EC2 instance will operate.

- b. At the **Network** step of the wizard, select an Amazon VPC and a subnet to which you want to connect worker instances created based on the new worker configuration, and specify a security group that will be associated with the instances. For an Amazon VPC, a subnet and a security group to be displayed in the lists of available network specifications, they must be created in AWS as described in [AWS Documentation](#).

Veeam Backup for AWS will apply the specified network settings to all worker instances that will be deployed in the specified location. For EFS indexing, Veeam Backup for AWS will also apply these settings to worker instances deployed to process file systems that have mount targets in the selected VPC.

IMPORTANT

When adding a worker configuration, consider the following:

- [Applies only to worker instances used for EFS indexing] The selected security group must allow outbound access on ports **2049** and **443**. These ports are used by worker instances to mount file systems and to communicate with [AWS services](#). Proxy redirect and setting a proxy in the Veeam Backup for AWS configuration are not supported.
- [Applies only to worker instances used for EFS indexing] The **DNS resolution** option must be enabled for the selected VPC. For more information, see [AWS Documentation](#).
- [Applies only to worker instances used for EC2 backup and restore] The selected security group must allow outbound access on port **443** required to communicate with [AWS services](#). Proxy redirect and setting a proxy in the Veeam Backup for AWS configuration are not supported.

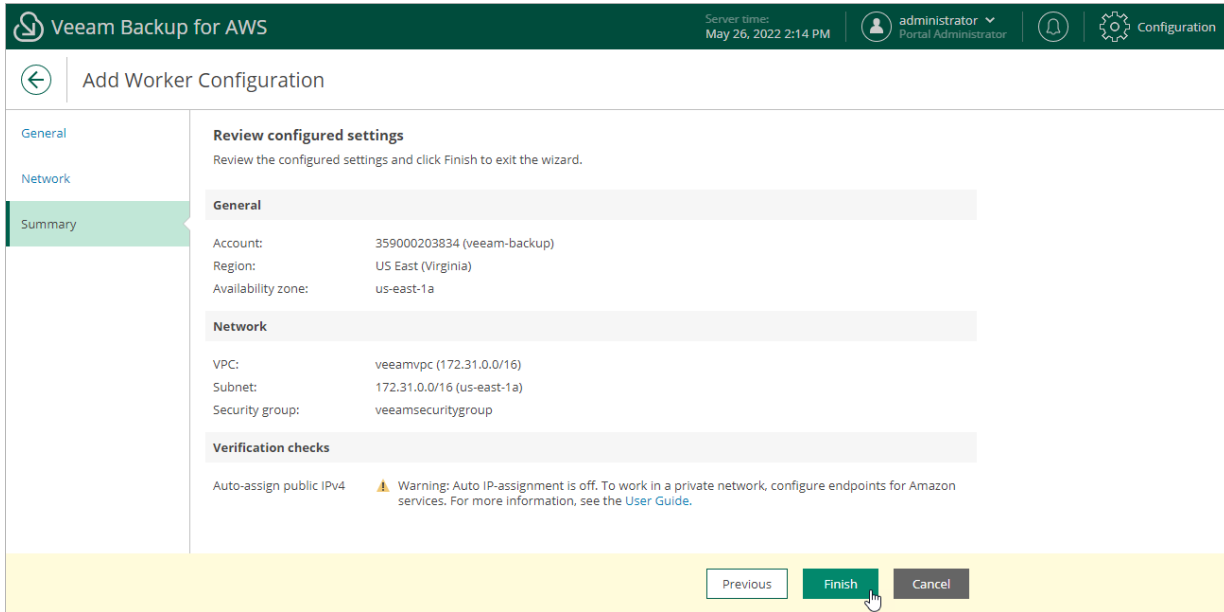
By default, Veeam Backup for AWS uses public access to communicate with worker instances. That is why the [public IPv4 addressing](#) attribute must be enabled for the selected subnet, the selected VPC must have an [internet gateway attached](#), and the VPC and subnet route tables must have routes that direct internet-bound traffic to this internet gateway. If you want worker instances to operate in a private network, do either of the following:

- Enable the private network deployment functionality, and configure specific VPC endpoints for the subnet to let Veeam Backup for AWS use private IPv4 addresses as described in section [Configuring Private Network Deployment](#).

For the list of specific endpoints required to perform backup and restore operations, see [Configuring Private Networks](#).

- Configure VPC endpoints as described in section [Appendix C. Configuring Endpoints in AWS](#).

c. At the **Summary** step of the wizard, review summary information and click **Finish**.



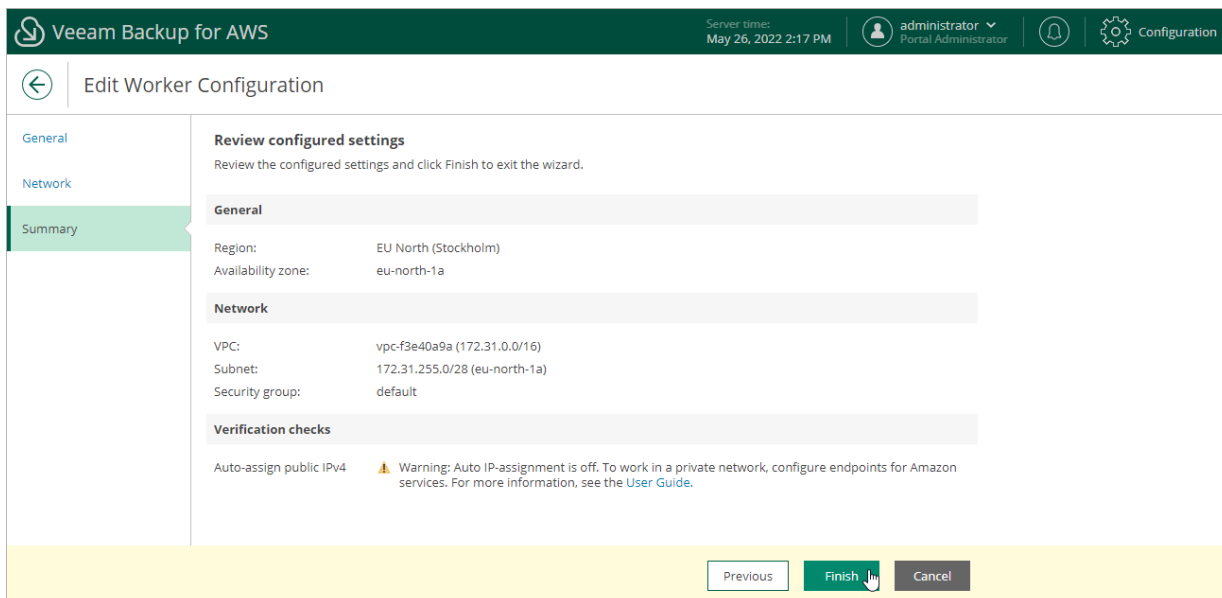
Editing Configurations

You can edit worker configurations added for AWS Regions:

1. Switch to the **Configuration** page.
2. Navigate to **Workers > Network**.
3. Switch to the necessary tab.
4. Select the worker configuration and click **Edit**.
5. Complete the **Edit Worker Configuration** wizard:
 - a. To change the VPC and subnet to which the related worker instances are connected, and the security group associated with the instances, follow the instructions provided in section [Adding Configurations for Backup Account](#) (step 2.b) or in section [Adding Configurations for Production Accounts](#) (step 5.b).
 - b. At the **Summary** step of the wizard, review configuration information and click **Finish** to confirm the changes.

NOTE

If any worker instances are currently deployed in the selected AWS Region, the changes will be applied only when Veeam Backup for AWS removes the instances from infrastructure (that is, when the running backup or restore process completes).



Removing Configurations

Veeam Backup for AWS allows you to permanently remove worker configurations if you no longer need them. When you remove a worker configuration, Veeam Backup for AWS does not remove currently running worker instances that have been created based on this configuration – these instances are removed only when the related operations complete.

To remove a worker configuration from Veeam Backup for AWS, do the following:

1. Switch to the **Configuration** page.
2. Navigate to **Workers > Network**.
3. Switch to the necessary tab.
4. Select the worker configuration and click **Remove**.

NOTE

If there are any worker instances created based on the selected configuration that are currently involved in a backup or restore process, these instances will be removed only when the process completes.

The screenshot shows the Veeam Backup for AWS configuration interface. The top navigation bar includes the Veeam logo, 'Veeam Backup for AWS', server time (Oct 12, 2023 2:42 PM), user (administrator), and a 'Configuration' button. The left sidebar contains navigation options: Getting Started, Administration, Accounts, Repositories, Workers (selected), Server settings, General, Configuration Backup, Licensing, and Support Information. The main content area is divided into tabs: Network, Profile, Tags, and Advanced. Under the 'Network' tab, there are sections for 'Backup Accounts' and 'Production Accounts'. A text block explains that worker instances are temporary Linux-based EC2 instances. Below this is the 'IAM role' section, which is currently empty. The 'Worker configurations' section shows a table with columns: Region, Availability Zone, Virtual Private Cl..., Subnet, Security Group, and FLR Status. Two configurations are listed for the 'Europe (Paris)' region. A dialog box titled 'Remove Worker Configuration' is overlaid on the table, with a message: 'If you remove the worker configuration, the default network settings will be used to launch worker instances in the selected region.' The dialog has 'Remove' and 'No' buttons.

Worker instances are temporary Linux-based EC2 instances responsible for the interaction with backup repositories and performing data copy tasks. They are automatically launched for the duration of a backup or restore operation and removed immediately when the operation completes.

IAM role
Select an IAM role that will be used to launch worker instances.

Service IAM role: Worker deployment

Worker configurations
The default network settings of a worker configuration are used to launch worker instances. To change these settings, add a worker configuration.

Region	Availability Zone	Virtual Private Cl...	Subnet	Security Group	FLR Status
Selected: 1 of 2					
<input type="checkbox"/> Europe (Paris)	eu-west-3b	amroz-srv VPC (17...	172.28.0.0/20 (eu-...	amroz-srv-VcbSecu...	Never Executed
<input checked="" type="checkbox"/> Europe (Paris)	eu-west-3c	dept-01-amroz-srv...	172.28.0.0/20 (eu-...	veeamsecuritygroup	Never Executed

Managing Worker Profiles

Worker profiles are instance types of worker instances that Veeam Backup for AWS deploys in a specific AWS Region to perform backup, restore, archive and health check operations. Veeam Backup for AWS deploys one worker instance per each AWS resource added to a backup policy or restore task. The profile of each deployed worker instance is selected based on the performed operation and the size of EBS volumes attached to the processed instance.

There are 4 types of worker profiles in Veeam Backup for AWS:

- **Small profile** – a profile that is used for EC2 and RDS backup and restore operations if the total size of all EBS volumes of the processed instance is less than 1024 GB.
- **Medium profile** – a profile that is used for EC2 and RDS backup and restore operations if the total size of all EBS volumes of the processed instance is 1024 GB - 16 TB.
- **Large profile** – a profile that is used for EC2 and RDS backup and restore operations if the total size of all EBS volumes of the processed instance is more than 16 TB.
- **Archiving profile** – a profile that is used for creating EC2 and RDS archived.

Out of the box, Veeam Backup for AWS comes with the default set of worker profiles where the small profile is *c5.large*, the medium profile is *c5.2xlarge*, the large profile is *c5.4xlarge*, and the archiving profile is *c5.2xlarge*. However, to boost operational performance, you can add custom sets of worker profiles to specify instance types of worker instances that will be deployed in different regions.

IMPORTANT

You cannot change the default worker profile used to deploy worker instances that perform EC2 file-level recovery, EFS indexing and retention operations – the default instance sizes of these worker instances are described in section [Worker Instance Locations](#). If you want to use a specific instance size for these worker instances, open a [support case](#).

Adding Profiles

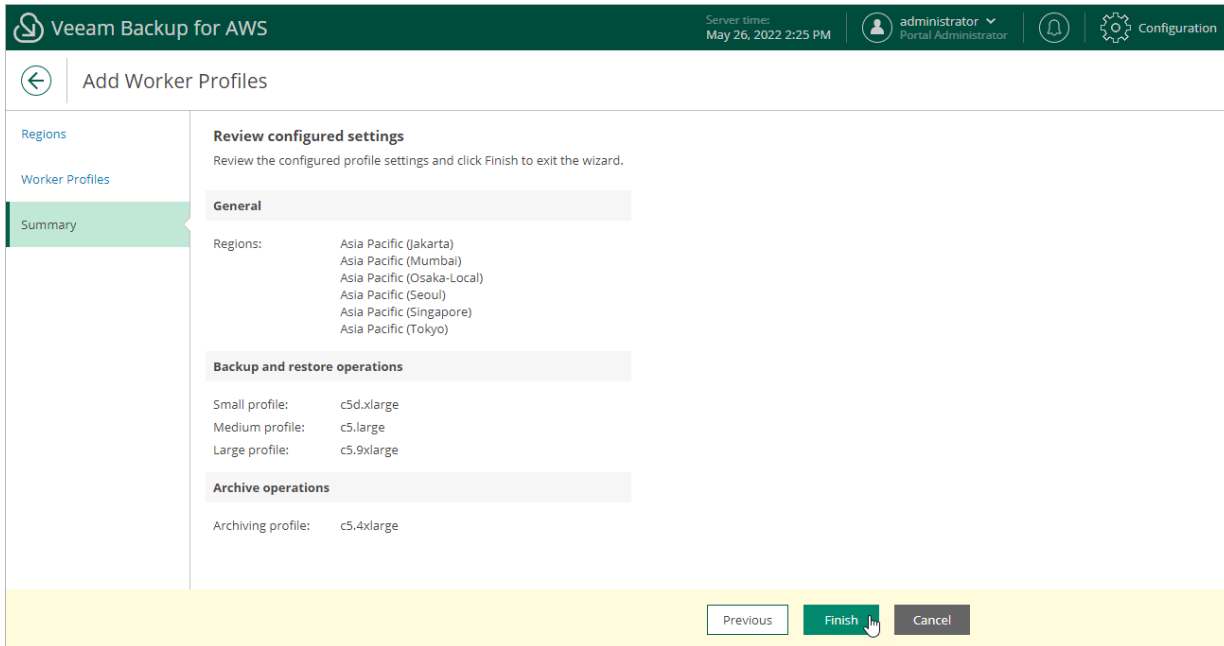
For each AWS Region in which worker instances will be deployed, you can add a custom set of worker profiles:

1. Switch to the **Configuration** page.
2. Navigate to **Workers > Profile** and click **Add**.
3. Complete the **Add Worker Profiles** wizard.
 - a. At the **Regions** step of the wizard, select regions for which you want to specify worker profiles and click **Add**.
 - b. At the **Worker Profiles** step of the wizard, choose profiles that will be used to deploy workers in the selected regions. To help you choose, tables in the **Choose instance type** section will provide information on the number of vCPU cores and the amount of system RAM for each available instance type.

For the full description of instance types that can be used to deploy EC2 instances in AWS, see [AWS Documentation](#).

- c. At the **Summary** step of the wizard, review summary information and click **Finish**.

As soon as you click **Finish**, Veeam Backup for AWS will create a separate set of worker profiles for each of the selected regions.




Editing Profiles

For each set of worker profiles created for an AWS Region, you can modify settings specified while creating the profile set:

1. Switch to the **Configuration** page.
2. Navigate to **Workers > Profiles**.
3. Select the profile set and click **Edit**.
4. Complete the **Edit Worker Profiles** wizard:
 - a. To change profiles that will be used to deploy workers in the selected region, follow the instructions provided in section [Adding Profiles](#) (step 3.b).
 - b. At the **Summary** step of the wizard, review configuration information and click **Finish** to confirm the changes.

NOTE

If there are any worker instances that are currently involved in a backup or archive backup process in the selected region, the changes will be applied only when the process completes.



Removing Profiles

Veeam Backup for AWS allows you to permanently remove sets of worker profiles if you no longer need them. When you remove a profile set, Veeam Backup for AWS does not remove currently running worker instances that have been created based on this set – these instances are removed only when the related operations complete.

To remove a profile set from Veeam Backup for AWS, do the following:

1. Switch to the **Configuration** page.
2. Navigate to **Workers > Profiles**.
3. Select the profile set and click **Remove**.

Veeam Backup for AWS

Server time: Mar 6, 2023 4:10 PM administrator Portal Administrator Configuration

Exit Configuration

Network Profile Advanced

Add worker profiles that will be used to launch worker instances for backup, restore and archive operations in AWS regions.

By default, the following profiles are used:

- Small profile (c5.large) is used if the size of the largest processed EBS volume is less than 1024 GB.
- Medium profile (c5.2xlarge) is used if the size of the largest processed EBS volume is between 1024 GB and 16 TB.
- Large profile (c5.4xlarge) is used if the size of the largest processed EBS volume is greater than 16 TB.

- Archiving profile (c5.4xlarge)

Remove Worker Profile

If you remove the worker profile, the default profile settings will be used to launch worker instances in the selected region.

Region

Region ↑

Selected: 1 of 5

Region	Small Profile	Medium Profile	Large Profile	Archiving Profile
<input type="checkbox"/> Asia Pacific (Tokyo)	c5d.xlarge	c5.large	c5.9xlarge	c5.4xlarge
<input checked="" type="checkbox"/> Asia Pacific (Seoul)	c5d.xlarge	c5.large	c5.9xlarge	c5.4xlarge
<input type="checkbox"/> Asia Pacific (Singapore)	c5d.xlarge	c5.large	c5.9xlarge	c5.4xlarge
<input type="checkbox"/> EU West (London)	c5.large	c5.xlarge	c5.2xlarge	c5.2xlarge
<input type="checkbox"/> EU West (Paris)	t4g.large	t3a.xlarge	t4g.medium	t3.2xlarge

Adding Worker Tags

For all worker instances that are deployed in specific AWS Regions for the duration of backup, restore and retention processes, you can assign custom AWS tags, which may help you differentiate worker instances that have the same or similar names:

1. Switch to the **Configuration** page.
2. Navigate to **Workers > Tags**.
3. Use the **Key** and **Value** fields to specify a key and a value for a new custom AWS tag, and then click **Add**. Note that you cannot add more than 25 custom AWS tags.

Consider the following limitations:

- The maximum length of the tag key is 128 characters.
- The maximum length of the tag value is 256 characters.
- The `aws:` prefix is reserved for AWS use and cannot be added.

For more information on tag limitations, see [AWS Documentation](#).

4. Click **Save**.

TIP

You can use a number of runtime variables as tag values to allow Veeam Backup for AWS to automatically fill in specific information for worker instances deployed during data protection operations. However, for worker instances deployed during restore operations, retention tasks, configuration checks and FLR tests, the values of the `%policyid%` and `%policyName%` variables will be replaced with operation names.

The screenshot shows the Veeam Backup for AWS configuration interface. The top navigation bar includes the Veeam logo, the text "Veeam Backup for AWS", the server time "Oct 26, 2023 10:02 AM", and the user "administrator Portal Administrator". The left sidebar contains navigation options: "Exit Configuration", "Getting Started", "Administration" (Accounts, Repositories, Workers), and "Server settings" (General, Configuration Backup, Licensing, Support Information). The main content area is titled "Tags" and includes a "Save" button and a warning "Your changes are not saved yet.". Below this, there is a message: "You can assign custom tags to workers and use this for billing, security, monitoring and reporting services.". The "Key:" field contains "policyname" and the "Value:" field contains "%policyName%". An "Add" button is next to the value field. Below the fields, there is a list of tags: "applianceid: %applianceid%". A note states "A maximum of 25 custom tags is allowed.". At the bottom, there is an information box with the title "These parameters can be used as a value for tags:" and a table of variables:

<code>%applianceid%</code>	Assigns the unique Veeam Backup appliance ID
<code>%policyid%</code>	Assigns the policy ID for which the worker is deployed
<code>%policyName%</code>	Assigns the policy name for which the worker is deployed

Configuring General Settings

Veeam Backup for AWS allows you to configure general settings that are applied to all performed operations and deployed infrastructure components.

- [Configure private network deployment mode to ensure secure communication between infrastructure components.](#)
- [Define for how long obsolete snapshots and session records must be retained.](#)
- [Configure notification settings for automated delivery of reports.](#)
- [Provide certificates to secure connections between Veeam Backup for AWS components.](#)
- [Change the time zone set on the backup appliance.](#)
- [Configure single sign-on settings to retrieve user identities from an identity provider.](#)

Configuring Private Network Deployment

If you want [worker instances](#) to operate in a private network, you can enable the private network deployment functionality and instruct Veeam Backup for AWS to deploy worker instances without public IPv4 addresses. In this case, worker instances will communicate with the Amazon S3 service through a private S3 endpoint specified in repository settings for data protection and recovery tasks.

To configure private network deployment, do the following:

1. Switch to the **Configuration** page, navigate to **General > Deployment Mode** and set the **Private network deployment** toggle to *On*.
2. To allow worker instances to access AWS services, create VPC interface endpoints for all subnets to which the worker instances will be connected, as described in section [Configuring Private Networks](#) (step 1).
3. To allow worker instances to communicate with the Amazon S3 service, do the following:
 - a. For all VPCs in the AWS Regions where backup repositories are located, create an S3 interface endpoint for all subnets to which worker instances will be connected, as described in section [Configuring Private Networks](#) (step 1).
 - b. For the backup appliance and worker instances, ensure connectivity between them and the Amazon S3 service, as described in section [Configuring Private Networks](#) (steps 2-3).
4. To allow worker instances to access Amazon S3 buckets, configure repository settings to use the created S3 interface endpoint for backup operations:
 - a. Click **Save** to enable the private network deployment functionality.
 - b. Click the **Configure repositories** link.
 - c. In the **Configuration Issues** window, click the link in the **Settings** column.

For a backup repository to be displayed in the list of available repositories, it must be added to Veeam Backup for AWS as described in section [Managing Backup Repositories](#).
 - d. In the **Edit Repository** wizard, navigate to the **Settings** step. Then, from the **Interface VPC endpoint** drop-down list, select the S3 interface endpoint that will be used to communicate with the Amazon S3 service.

For an S3 interface endpoint to be displayed in the **Interface VPC endpoint** list, it must be created in the Amazon VPC console for all subnets to which the worker instances will be connected, as described in section [Configuring Private Networks](#) (step 1).

To check whether you have configured all the necessary settings correctly, run your backup policies as described in section [Performing Backup](#).

The screenshot shows the Veeam Backup for AWS configuration interface. The top navigation bar includes the Veeam logo, the text "Veeam Backup for AWS", the server time "Sep 29, 2023 2:15 PM", the user "administrator Portal Administrator", and a "Configuration" button. The left sidebar contains a navigation menu with items: "Exit Configuration", "Getting Started", "Administration", "Accounts", "Repositories", "Workers", "Server settings", "General" (highlighted), "Configuration Backup", "Licensing", and "Support Information". The main content area is titled "Deployment Mode" and contains the following text: "This setting allows a backup appliance to operate in infrastructures with restricted public access. For more information, see the [User Guide](#)." Below this is a "Save" button and a yellow warning message: "Your changes are not saved yet." A toggle switch for "Private network deployment" is set to "On". Below the toggle, it says "To configure settings for private environment, it is required to follow the steps below." A numbered list of prerequisites follows: 1. **Workers Prerequisites**: Configure worker region requirements as described in the [User Guide](#). 2. **Repository Prerequisites**: Configure S3 repository region requirements as described in the [User Guide](#). 3. **Repository Configuration Verification**: [Configure repositories](#) to utilize an interface VPC endpoint for backup operations. 4. **Test Backup Policies**: [Run the backup policies](#) to verify everything is correctly configured.

Configuring Global Retention Settings

You can configure global retention settings to specify for how long the following data must be retained in the configuration database:

- [Obsolete snapshots and replicas](#)
- [Session records](#)

Configuring Retention Settings for Obsolete Snapshots and Replicas

If an instance is no longer processed by a backup policy (for example, it was removed from the backup policy or the backup policy no longer exists), its cloud-native snapshots and snapshot replicas become obsolete. Retention policy settings configured when creating backup policies do not apply to obsolete snapshots – these snapshots are removed from the configuration database according to their own retention settings.

NOTE

Global retention settings apply to all EC2 and RDS cloud-native snapshots, as well as to snapshot replicas created by the Veeam backup service. If an instance is still processed by a backup policy, but some of its cloud-native snapshots and snapshot replicas are older than the number of days (or months) specified in the global retention settings, these cloud-native snapshots and snapshot replicas will not be removed from Veeam Backup for AWS.

To configure retention settings for obsolete snapshots and replicas, do the following:

1. Switch to the **Configuration** page.
2. Navigate to **General > Retention**.
3. In the **Obsolete snapshots retention** section, select one of the following options:
 - Select the **Never** option if you do not want Veeam Backup for AWS to remove obsolete snapshots and replicas.
 - Select the **After** option to specify the number of days (or months) during which Veeam Backup for AWS must keep obsolete snapshots in the configuration database. For days, the number must be between 15 and 36135. For months, the number must be between 1 and 1188.

If you select this option, Veeam Backup for AWS will remove obsolete snapshots of an instance as soon as the specified period is over.

4. Click **Save**.

NOTE

When Veeam Backup for AWS removes an obsolete snapshot from the configuration database, it also removes the snapshot from AWS.

Configuring Retention Settings for Session Records

Veeam Backup for AWS stores records for all sessions of performed data protection and disaster recovery operations in the configuration database on the additional data disk attached to the backup appliance. These session records are removed from the configuration database according to their own retention settings. By default, session logs are stored for 3 months.

To configure retention settings for session records, do the following:

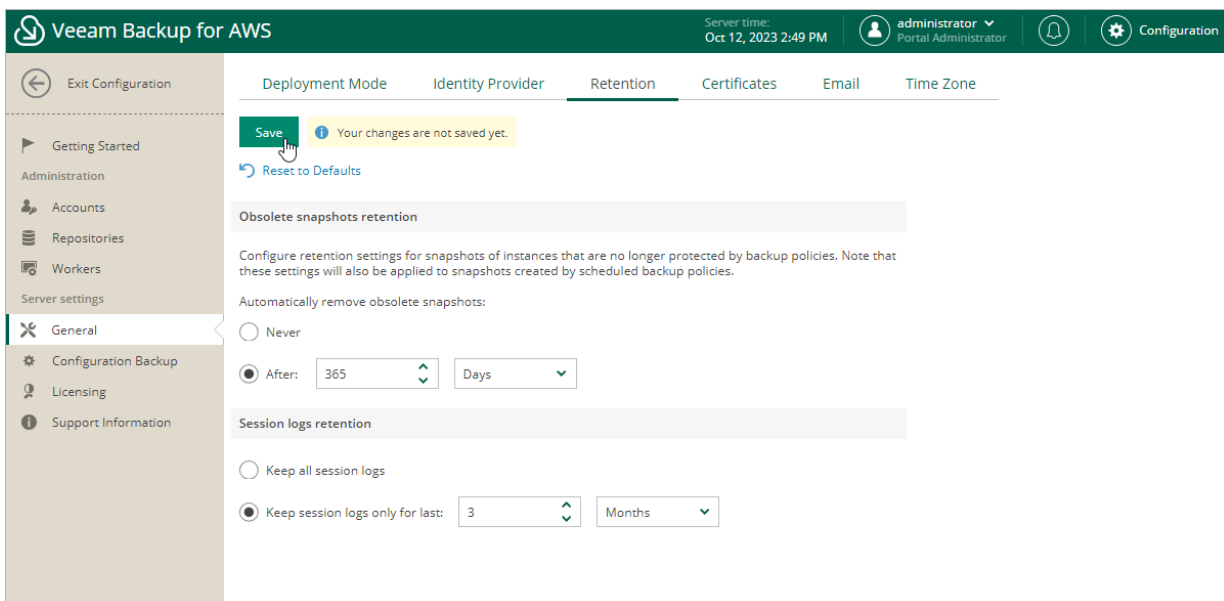
1. In the **Session logs retention** section, select one of the following options:
 - Select the **Keep all session logs** option if you do not want Veeam Backup for AWS to remove session records.
 - Select the **Keep session logs only for last** option if you want to specify the number of days (or months) during which Veeam Backup for AWS must keep session records in the configuration database.

If you select this option, Veeam Backup for AWS will remove all session records that are older than the specified time limit.

2. Click **Save**.

IMPORTANT

Retaining all session records in the configuration database may overload the data EBS volume. By default, the volume comes with 20 GB of storage capacity. If you choose not to remove sessions records at all, consider increasing the volume capacity to avoid runtime problems.



Configuring Global Notification Settings

You can specify email notification settings for automated delivery of backup policy results and daily reports. Every daily report contains cumulative statistics on all backup and restore sessions, as well as retention sessions performed within the past 24-hour period.

IMPORTANT

Veeam Backup for AWS does not support sending e-mails through TLS Wrapper.

To connect an email service that will be used for sending email notifications:

1. Switch to the **Configuration** page.
2. Navigate to **General > Email**.
3. Select the **Enable email notifications** check box.
4. Click the link in the **Email server** field and configure [email server settings](#).
5. In the **From** field, enter an email address of the notification sender. This email address will be displayed in the **From** field of notifications.
6. In the **To** field, enter an email address of a recipient.

For each particular policy, you can configure specific notification settings. For more information on backup policies, see [Performing Backup](#).

NOTE

If you specify the same email recipient in both backup policy notification and global notification settings, Veeam Backup for AWS will override the configured global notification settings and will send each notification to this recipient only once to avoid notification duplicates.

7. In the **Subject** field, specify a subject for notifications. You can use the following runtime variables:
 - *%JobName%* – a backup policy name.
 - *%JobResult%* – a backup policy result.
 - *%ObjectCount%* – the number of instances in a backup policy.
 - *%Issues%* – the number of instances in a backup policy that encountered any issues (errors and warnings) while being processed.

The default subject for email notifications is: *[%JobResult%] %JobName% (%ObjectCount% instances) %Issues%*.

8. In the **Notify me immediately on policy** section, choose whether you want to receive email notifications in case backup policies complete successfully, complete with warnings or complete with errors.
9. To receive daily reports, select the **Send daily report at** check box and specify the exact time when the reports will be sent.
10. Click **Save**.

TIP

Veeam Backup for AWS allows you to send a test message to check whether you have configured all settings correctly. To do that, click **Send Test Email**. A test message will be sent to the specified email address.

Configuring Email Server Settings

To configure email server settings, choose whether you want to employ [Basic \(SMTP\)](#) or [Modern \(OAuth 2.0\)](#) authentication for your email service.

Using Basic Authentication

To employ the Basic authentication to connect to your email server, in the **Email Server Settings** window:

1. From the **Authentication** drop-down list, select *Basic*.
2. In the **Mail server name or address** field, enter a DNS name or an IP address of the SMTP server. All email notifications (including test messages) will be sent by this SMTP server.
3. In the **Port** field, specify a communication port for SMTP traffic. The default SMTP port is 25.
4. In the **Timeout** field, specify a connection timeout for responses from the SMTP server.
5. For an SMTP server with SSL/TLS support, select the **Connect using SSL** check box to enable SSL data encryption.
6. If your SMTP server requires authentication, select the **This server requires authentication** check box and choose an account that will be used when authenticating against the SMTP server from the **Connect as** drop-down list.

For an account to be displayed in the list of available accounts, it must be added to Veeam Backup for AWS as described in section [Adding SMTP Accounts](#). If you have not added an account beforehand, click **Add** and complete the **Add Account** wizard.

7. Click **OK**.

Using Modern Authentication

To employ the Modern authentication to connect to your email service:


1. From the **Authentication** drop-down list, select *Modern*.
2. In **Email Server Settings** window, copy the URL from the **Redirect URL** field.
If you plan to send notifications using the Google email service, make sure that the Veeam Backup for AWS Web UI is open using the public IPv4 DNS.
3. For Veeam Backup for AWS to be able to use OAuth 2.0 to access Google Cloud or Microsoft Azure APIs, register a new client application either in the [Google Cloud Console](#) or in the [Microsoft Azure portal](#).

When registering the application, make sure that the redirect URL specified for the application matches the URL copied from the Veeam Backup for AWS Web UI.

IMPORTANT

Due to DNS name differences, the Google email service does not support the redirect URL of the backup appliance that was deployed in the US East (N.Virginia) region.

4. Back to the Veeam Backup for AWS Web UI, do the following in the **Email Server Settings** window:
 - a. Use the **Mail service** drop-down list to choose whether the service that you want to use to send email notifications is a *Google* or *Microsoft* email service.
 - b. In the **Application client ID** and **Client secret** fields, provide the Client ID and Client secret created for the application as described in [Google Cloud documentation](#) or [Microsoft Docs](#).
 - c. [Applies only if you have selected the **Microsoft** option] In the **Tenant ID** field, provide the ID of an Azure AD tenant in which the application has been registered.
 - d. Click **Authorize**. You will be redirected to the authorization page. Sign in using a Google or Microsoft Azure account to validate the configured settings.



Adding SMTP Accounts


To add an account that will be used to connect to an SMTP server, do the following:

1. Switch to the **Configuration** page.
2. Navigate to **Accounts > SMTP Accounts**.
3. Click **Add**.

Complete the **Add Account** wizard.

- a. At the **Account Name** step of the wizard, specify a name and description for the SMTP account. The name must be unique in Veeam Backup for AWS and the length of the name must not exceed 255 characters. The description length must not exceed 255 characters.
- b. At the **Account** step of the wizard, specify credentials of a user account that has permissions to access the SMTP server. Veeam Backup for AWS will use the specified credentials to authenticate against the SMTP server.

c. At the **Summary** step of the wizard, review summary information and click **Finish**.




Editing SMTP Accounts

For each SMTP account, you can modify the settings configured while adding the account:

1. Switch to the **Configuration** page.
2. Navigate to **Accounts > SMTP Accounts**.
3. Select the check box next to the necessary SMTP account and click **Edit**.

Complete the **Edit Account** wizard.

- a. To provide a new name and description for the account, follow the instructions provided in section [Adding SMTP Accounts](#) (step 3a).
- b. To specify credentials of another user account to be used to authenticate against the SMTP server, follow the instructions provided in section [Adding SMTP Accounts](#) (step 3b).



Replacing Security Certificates

To establish secure data communications between the backup appliance and web browsers running on user workstations, Veeam Backup for AWS uses Transport Layer Security (TLS) certificates.

IMPORTANT

Starting from Veeam Backup for AWS version 5.0, only the TLS v1.3 certificates are supported. Therefore, Veeam Backup for AWS will automatically recreate the previously generated self-signed certificate when updating the backup appliance.

When you install Veeam Backup for AWS, it automatically generates a default self-signed certificate. You can replace this default certificate with your own self-signed certificate or with a certificate obtained from a Certificate Authority (CA). To replace the currently used TLS certificate, do the following:

1. Switch to the **Configuration** page.
2. Navigate to **General > Certificates**.
3. Click **Replace Web Certificate**.

Complete the **New Certificate Wizard**.

- a. At the **Certificate Source** step of the wizard, do the following:
 - Select the **Recreate a self-sign certificate** option if you want to replace the existing certificate with a new self-signed certificate automatically generated by Veeam Backup for AWS.
 - Select the **Upload certificate** option if you want to upload a certificate that you obtained from a CA or generated using a 3rd party tool.
- b. [This step applies only if you have selected the **Upload certificate(s)** option] At the **Upload certificate(s)** step of the wizard, browse to the certificate that you want to install, and provide a password for the certificate file if required.

NOTE

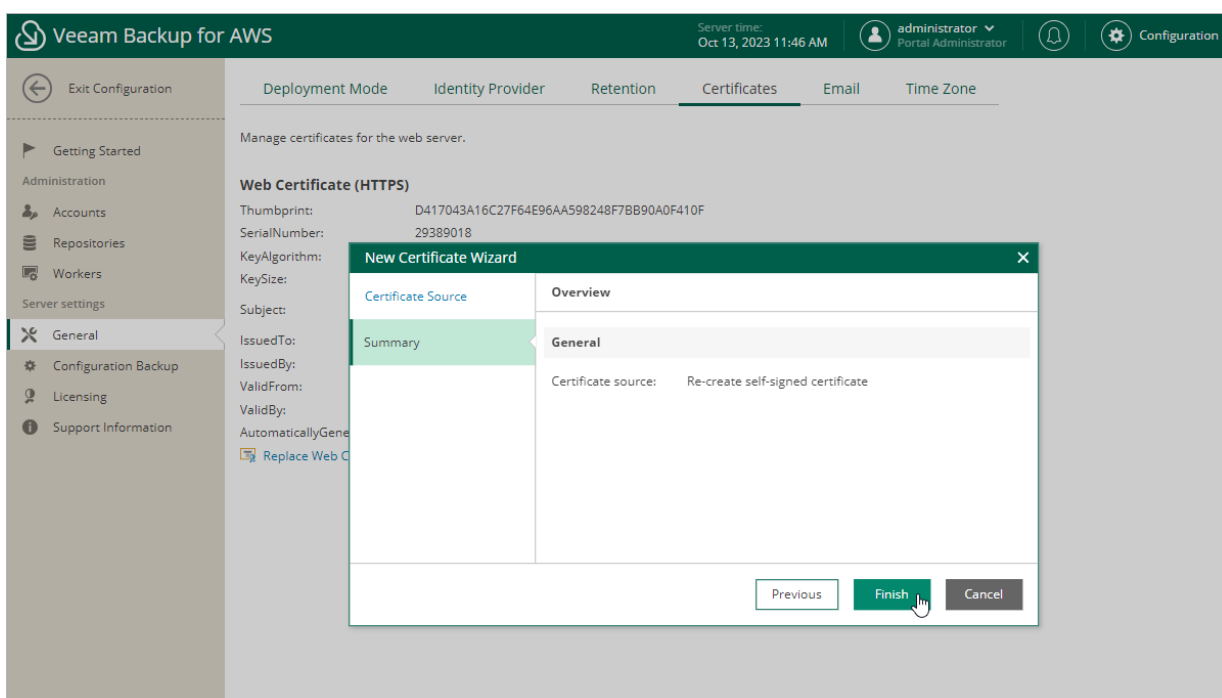
Veeam Backup for AWS supports certificates only in the .PFX and .P12 formats.

- c. At the **Summary** step of the wizard, review summary information and click **Finish**.

NOTE

Consider the following:

- If you have recreated the self-signed certificate, the browser from which you will try to access Veeam Backup for AWS next time will display a warning notifying that the connection is untrusted (although it is secured with SSL). To eliminate the warning, import the self-signed certificate to user workstations.
- If you have recreated the certificate for a backup appliance managed by the Veeam Backup & Replication server, Veeam Backup & Replication will not be able to license resources and to collect resource data. To work around the issue, accept the newly created certificate using the Veeam Backup & Replication console as described in section [Editing Appliance Settings](#).



Changing Time Zone

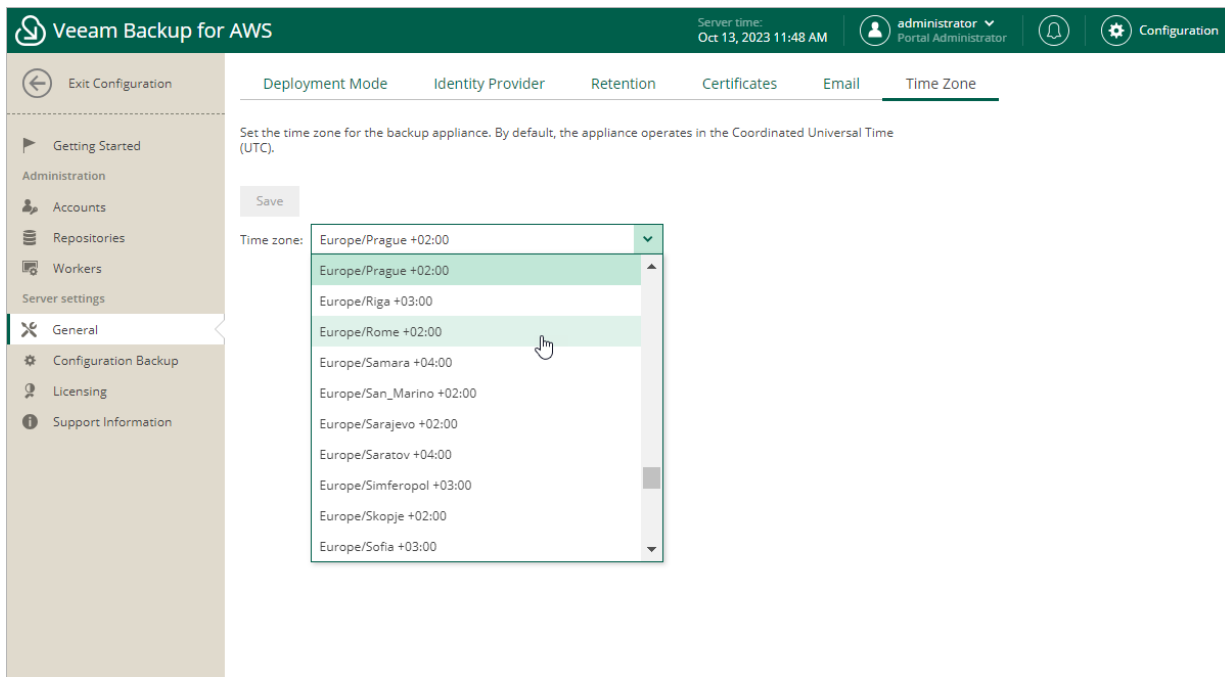
Veeam Backup for AWS runs daily reports and performs all data protection and disaster recovery operations according to the time zone set on the backup appliance. Since the backup appliance is deployed on an EC2 instance in Amazon EC2, the time zone is set to Coordinated Universal Time (UTC) by default. However, you can change the time zone if required. For example, you may want the time on the backup appliance to match the time on the workstation from which you access Veeam Backup for AWS.

To change the time zone set on the backup appliance:

1. Switch to the **Configuration** page.
2. Navigate to **General > Time Zone**.
3. Select the necessary time zone from the **Time zone** drop-down list.
4. Click **Save**.

NOTE

It is not recommended to change the time zone if any data protection or disaster recovery session is currently running. Wait for all the running sessions to complete or stop them manually – and then change the time zone. To learn how to track real-time statistics of all running and completed operations, see [Viewing Session Statistics](#).



Configuring SSO Settings

Veeam Backup for AWS supports single sign-on (SSO) authentication based on the SAML 2.0 protocol. SSO authentication scheme allows a user to log in to different software systems with the same credentials using the identity provider service.

To configure SSO settings for Veeam Backup for AWS, do the following:

1. Switch to the **Configuration** page.
2. Navigate to **General > Identity Provider**.
3. In the **Identity Provider Configuration** section, import identity provider settings from a file obtained from your identity provider:
 - a. Click **Upload Metadata**.
 - b. In the **Upload Identity Provider Configuration** window, click **Browse** to locate the file with the identity provider settings.
 - c. Click **Upload**.
4. Forward the service provider authentication settings to the identity provider – to obtain the settings, in the **Veeam Backup for AWS Configuration** section, click **Download**. Veeam Backup for AWS will download a metadata file with the service provider authentication settings to your local machine.

Alternatively, you can copy the service provider settings manually:

- a. Click **Copy Link** in the **SP Entity ID / Issuer** field.
 - b. Click **Copy Link** in the **Assertion Consumer URL** field.
5. [Optional] If you want to sign and encrypt authentication requests sent from Veeam Backup for AWS to the identity provider, select a certificate with a private key that will be used to sign and encrypt the requests:
 - a. In the **Veeam Backup for AWS Configuration** section, click **Select** in the **Certificate** field.
 - b. In the **Upload Veeam Backup certificate** window, click **Browse** to locate the certificate file. In the **Password** field, specify a password used to open the file.
 - c. Click **Upload**.

NOTE

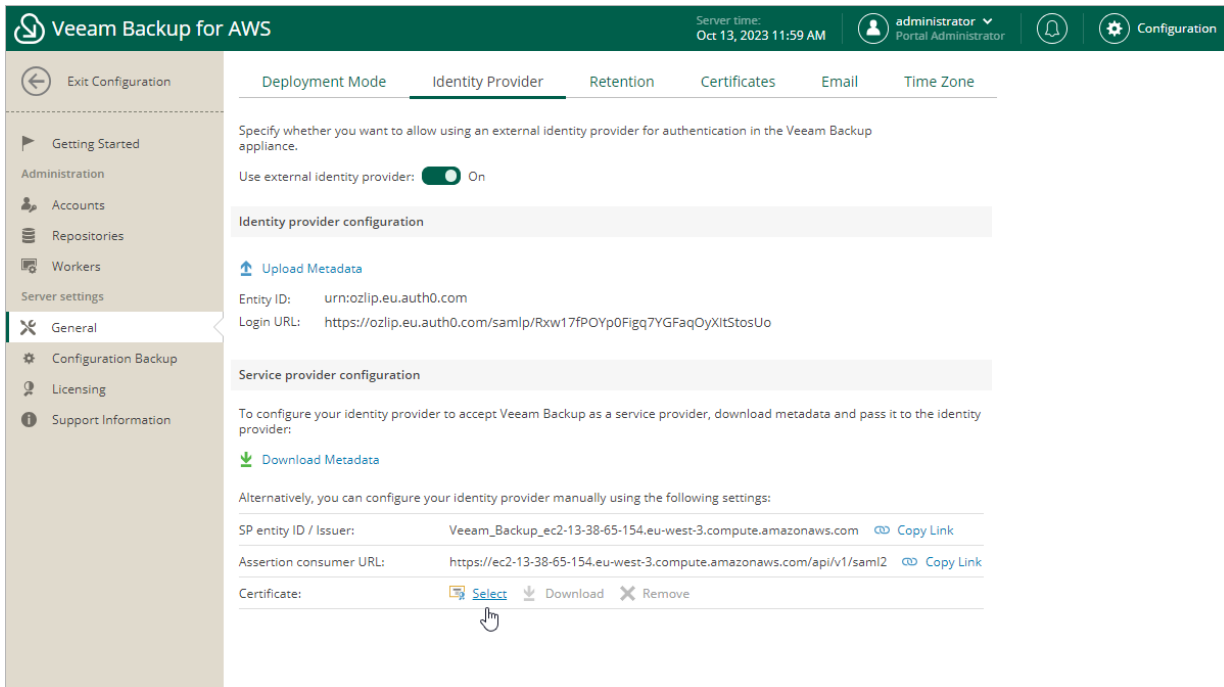
Veeam Backup for AWS supports certificates only in the .PFX and .P12 formats.

After you configure SSO settings, you can add user accounts that will be able to log in to Veeam Backup for AWS using single sign-on. For more information, see [Adding User Accounts](#).

IMPORTANT

To authenticate a user whose identity has been received from the identity provider, Veeam Backup for AWS redirects the user to the identity provider portal. After the user logs in to the portal, the identity provider sends a SAML authentication response to Veeam Backup for AWS. The SAML response must contain the `UserName` attribute to allow Veeam Backup for AWS to identify the user. The attribute value must match the user name that you specify [when creating the user account](#).

If your identity provider does not send the `UserName` attribute by default, you must create a claim rule on the identity provider side to send this attribute in the SAML authentication response to the Veeam Backup for AWS request.



The screenshot displays the Veeam Backup for AWS configuration interface. The top navigation bar includes the Veeam logo, the text "Veeam Backup for AWS", the server time "Oct 13, 2023 11:59 AM", and the user "administrator Portal Administrator". The left sidebar contains navigation options: "Exit Configuration", "Getting Started", "Administration", "Accounts", "Repositories", "Workers", "Server settings", "General", "Configuration Backup", "Licensing", and "Support Information". The main content area is titled "Identity Provider" and includes the following sections:

- Deployment Mode:** "Specify whether you want to allow using an external identity provider for authentication in the Veeam Backup appliance." The "Use external identity provider" toggle is set to "On".
- Identity provider configuration:** Includes an "Upload Metadata" button, "Entity ID: urn:ozlip.eu.auth0.com", and "Login URL: https://ozlip.eu.auth0.com/samlp/Rxw17fPOYp0Figq7YGFaqOyXItStosUo".
- Service provider configuration:** Includes a "Download Metadata" button and instructions: "To configure your identity provider to accept Veeam Backup as a service provider, download metadata and pass it to the identity provider:".
- Manual configuration settings:** Lists "SP entity ID / Issuer: Veeam_Backup_ec2-13-38-65-154.eu-west-3.compute.amazonaws.com" and "Assertion consumer URL: https://ec2-13-38-65-154.eu-west-3.compute.amazonaws.com/api/v1/saml2", both with "Copy Link" buttons.
- Certificate:** Includes a "Select" button, a "Download" button, and a "Remove" button.

Performing Configuration Backup and Restore

You can back up and restore the configuration database that stores data collected from Veeam Backup for AWS for the existing backup policies, protected AWS resources, created worker instance configurations and profiles, added IAM roles and users, logged session records and so on. If the backup appliance goes down for some reason, you can reinstall it and quickly restore its configuration from a backup. You can also use a configuration backup to migrate the configuration of one backup appliance to another backup appliance in AWS.

It is recommended that you regularly perform configuration backup for every backup appliance present in AWS. Periodic configuration backups reduce the risk of data loss and minimize the administrative overhead costs in case any problems with the backup appliances occur.

You can run configuration backup manually on demand, or instruct Veeam Backup for AWS to do it automatically on a regular basis.

Performing Configuration Backup

During configuration backup, data from configuration database of an appliance is exported and saved to a backup file in a repository. The configuration database contains the following information: existing backup policies, protected AWS resources, created worker instance configurations and profiles, added IAM roles and users, logged session records and so on.

IMPORTANT

If your backup appliance is managed by a Veeam Backup & Replication server, you will neither be able to perform manual or scheduled configuration backup of Veeam Backup for AWS from the Web UI, nor to export the configuration data from the Web UI. In this case, you can perform configuration backup using the Veeam Backup & Replication console as described in section [Performing Configuration Backup Using Console](#).

Performing Configuration Backup Using Console

While performing configuration backup, Veeam Backup & Replication backs up the configuration of the backup server and also configurations of all backup appliances added to the backup infrastructure. The results of every configuration backup session are displayed in the **History** view under the **System** node.

You can perform configuration backup manually or instruct Veeam Backup & Replication to do it automatically on a regular basis:

- To perform configuration backup manually, follow the instructions provided in the Veeam Backup & Replication User Guide, section [Running Configuration Backups Manually](#).
- To instruct Veeam Backup & Replication to perform configuration backup automatically, follow the instructions provided in the Veeam Backup & Replication User Guide, section [Scheduling Configuration Backups](#).

Before You Begin

If you plan to back up the configuration of a managed backup appliance, keep in mind the following limitations and considerations:

- You must enable backup file encryption in the configuration backup settings. Otherwise, Veeam Backup & Replication will back up only the backup server configuration.

To learn how to create encrypted configuration backup, see the Veeam Backup & Replication User Guide, section [Creating Encrypted Configuration Backups](#).

- You cannot store configuration backups in scale-out backup repositories and external repositories.
- For Veeam Backup & Replication to be able to back up the appliance configuration, the backup appliance must be available and must run a Veeam Backup for AWS version that is compatible with the Veeam Backup & Replication version.

For the list of compatible versions, see [System Requirements](#).

- During configuration backup, Veeam Backup & Replication processes only 3 appliances at a time – the appliances exceeding this limit are queued.

- To enable data loss protection in case you lose or forget the password used for data encryption, you can use Veeam Backup Enterprise Manager to decrypt backup files.

To learn how to let Veeam Backup & Replication encrypt and decrypt data with Enterprise Manager, see the Veeam Backup Enterprise Manager Guide, section [Managing Encryption Keys](#).

Configuration Backup Location

Veeam Backup & Replication stores configuration backups of backup appliances in a repository specified in configuration backup settings. Backups are saved in the `\\VeeamConfigBackup\AWS` folder.

NOTE

Consider the following:

- It is not recommended to store configuration backups on the backup server. Otherwise, you will not be able to restore configuration of managed backup appliances in case the backup server goes down.
- If the name of an appliance contains unsupported characters, these characters are replaced with the '_' underscore symbol in the name format for a subfolder and a backup files.

Performing Configuration Backup Using Web UI

While performing configuration backup, Veeam Backup for AWS exports data from the configuration database and saves it to a backup file in a backup repository. You can back up the configuration database of a backup appliance either manually or automatically.

IMPORTANT

If your backup appliance is managed by a Veeam Backup & Replication server, you will neither be able to perform manual or scheduled configuration backup of Veeam Backup for AWS from the Web UI, nor to export the configuration data from the Web UI. In this case, you can perform configuration backup using the Veeam Backup & Replication console as described in section [Performing Configuration Backup Using Console](#).

Performing Configuration Backup Manually

To back up the configuration database manually, do the following:

1. Switch to the **Configuration** page.
2. Navigate to **Configuration Backup**.
3. In the **Overview** section, click **Take Backup Now**.
4. In the **Create Manual Backup** window, select a repository where the configuration backup will be stored, and click **Create**.

For a backup repository to be displayed in the **Repository** list, it must be added to Veeam Backup for AWS as described in section [Adding Backup Repositories Using Web UI](#). The **Repository** list shows only backup repositories of the *S3 Standard* storage class that have encryption enabled and immutability disabled.

As soon as you click **Create**, Veeam Backup for AWS will start creating a new backup file in the selected repository. To track the progress, click **Go to Sessions** in the **Session Info** window to proceed to the [Session Logs](#) tab.

TIP

Once Veeam Backup for AWS creates a successful configuration backup, you can click **Export Last Backup** to download the backup file to a local machine and then use it to [restore configuration data](#).

Performing Configuration Backup Automatically

To instruct Veeam Backup for AWS to back up the configuration database automatically by schedule, do the following:

1. Switch to the **Configuration** page.
2. Navigate to **Configuration Backup**.
3. In the **Backup Schedule** section, set the **Enable scheduling** toggle to *On*.
4. Click the link in **Repository** field, and select a repository where configuration backups will be stored in the **Choose Repository** window.

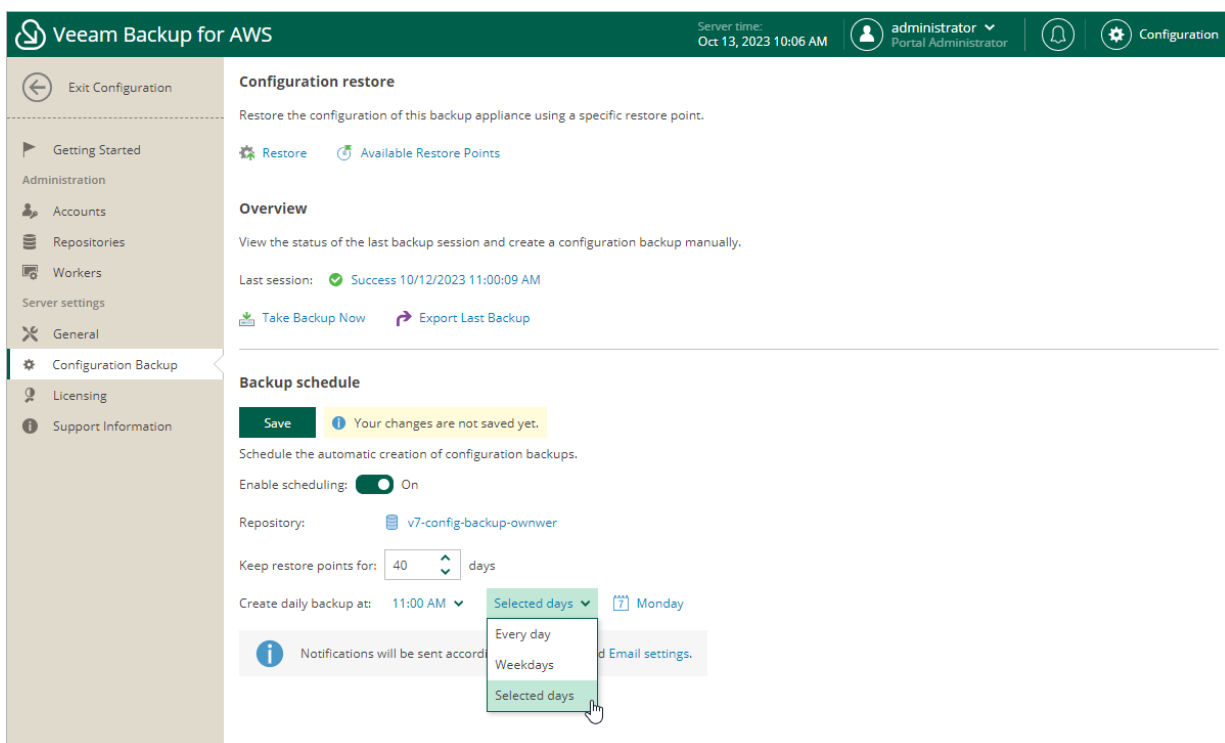
For a backup repository to be displayed in the list of available repositories, it must be added to Veeam Backup for AWS as described in section [Adding Backup Repositories Using Web UI](#). The list shows only backup repositories of the *S3 Standard* storage class that have encryption enabled and immutability disabled.

5. In the **Keep restore points for** field, specify the number of days for which you want to keep restore points in a backup chain in the selected backup repository.

If a restore point is older than the specified time limit, Veeam Backup for AWS removes the restore point from the backup chain.

6. In the **Create daily backup at** field, choose whether configuration backups will be created every day, on weekdays (Monday through Friday), or on specific days.

7. Click **Save**.



The screenshot shows the Veeam Backup for AWS web interface. The top navigation bar includes the Veeam logo, the text 'Veeam Backup for AWS', the server time 'Oct 13, 2023 10:06 AM', and the user 'administrator Portal Administrator'. The left sidebar contains navigation options: Exit Configuration, Getting Started, Administration (Accounts, Repositories, Workers), Server settings (General), Configuration Backup (selected), Licensing, and Support Information. The main content area is titled 'Configuration restore' and includes a 'Restore' button and a link to 'Available Restore Points'. Below this is an 'Overview' section showing the last session status as 'Success 10/12/2023 11:00:09 AM' and buttons for 'Take Backup Now' and 'Export Last Backup'. The 'Backup schedule' section is active, showing a 'Save' button and a warning 'Your changes are not saved yet.'. It includes a toggle for 'Enable scheduling' set to 'On', a 'Repository' dropdown set to 'v7-config-backup-ownwer', a 'Keep restore points for' field set to '40 days', and a 'Create daily backup at' field set to '11:00 AM' with a 'Selected days' dropdown menu open, showing options for 'Every day', 'Weekdays', and 'Selected days'. A 'Monday' calendar icon is also visible. A notification banner at the bottom states 'Notifications will be sent according to your Email settings.' with a link to 'Email settings.'

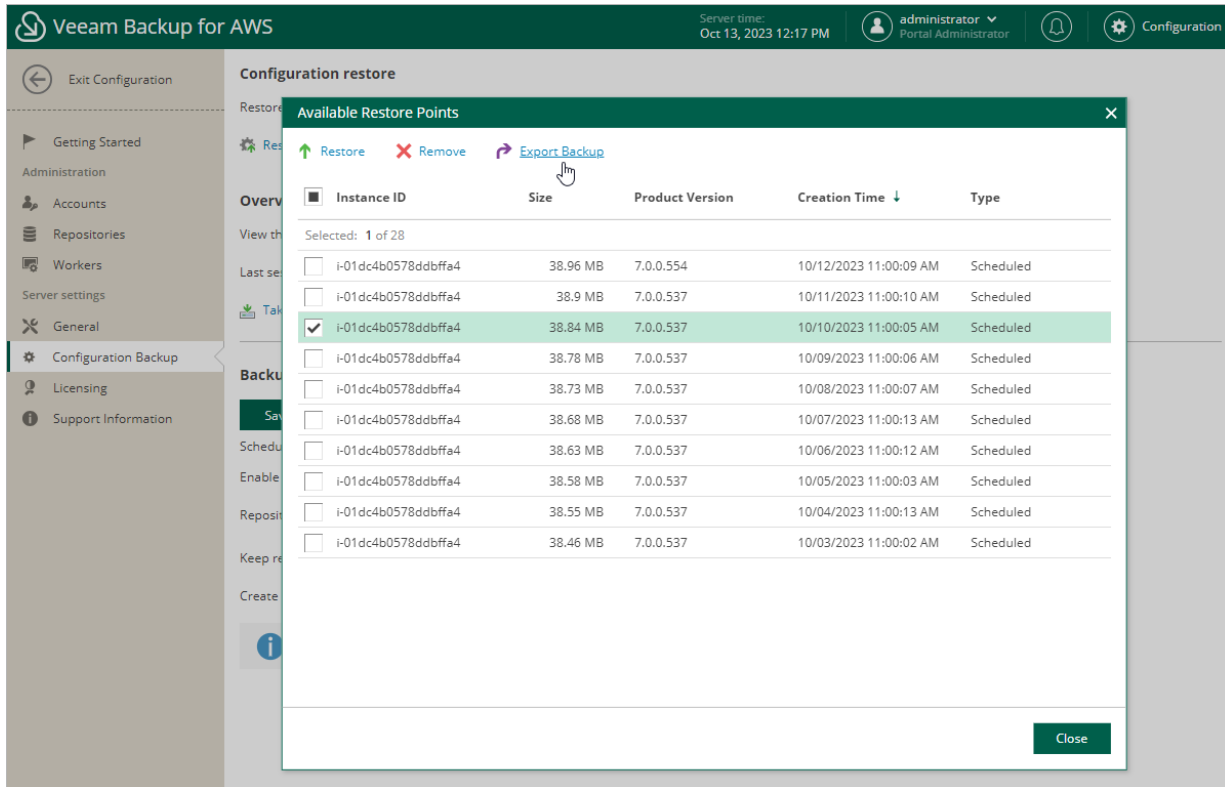
Exporting Configuration Backup Data

Once Veeam Backup for AWS creates a successful configuration backup, you can export the configuration backup file and use it to [restore configuration data](#) on another backup appliance.

To export the configuration backup file to a local machine, do the following:

1. Switch to the **Configuration** page.
2. Navigate to **Configuration Backup**.
3. Use one of the following options:
 - o To export the last successful configuration backup:
 - i. In the **Overview** section, click **Export Last Backup**.
 - ii. In the **Export Last Backup** window, specify a password that will be used to encrypt the exported file, provide a hint for the specified password, and click **Export**.
 - o To export a specific configuration backup file:
 - i. In the **Configuration restore** section, click **Available Restore Points**.
 - ii. In the **Available Restore Points** window, select the necessary backup and click **Export Backup**.
 - iii. In the **Export Backup** window, specify a password that will be used to encrypt the exported file, provide a hint for the specified password, and click **Export**.

As soon as you click **Export**, Veeam Backup for AWS will save the exported backup file to the default download directory on the local machine.



Performing Configuration Restore

Veeam Backup for AWS offers restore of the configuration database that can be helpful in the following situations:

- The configuration database got corrupted, and you want to recover data from a configuration backup.
- You want to roll back the configuration database to a specific point in time.
- A backup appliance got corrupted, and you want to recover its configuration from a configuration backup.
- A backup appliance went down, and you want to apply its configuration to a new backup appliance.

IMPORTANT

If your backup appliance is managed by a Veeam Backup & Replication server, you will not be able to restore the configuration of Veeam Backup for AWS from the Web UI. In this case, you can perform configuration restore using the Veeam Backup & Replication console as described in section [Restoring Configuration Data Using Console](#).

Restoring Configuration Data Using Console

To restore the configuration database of a backup appliance using the Veeam Backup & Replication console, do the following:

IMPORTANT

Before you start the restore process, stop all policies that are currently running.

1. [Check prerequisites and limitations](#).
2. [Launch the Configuration Restore wizard](#).
3. [Choose a backup file](#).
4. [Review the backup file info](#).
5. [Specify a decryption password](#).
6. [Choose restore options](#).
7. [Specify a user whose credentials will be used to connect to the appliance](#).
8. [Wait for the restore process to complete](#).
9. [Finish working with the wizard](#).

Limitations and Considerations

Before you restore the configuration database of a backup appliance, consider the following:

- Make sure there are no sessions currently running on the backup appliance. Also, make sure there are no backup policies scheduled to run during restore. Otherwise, backups created by these policies may be corrupted.

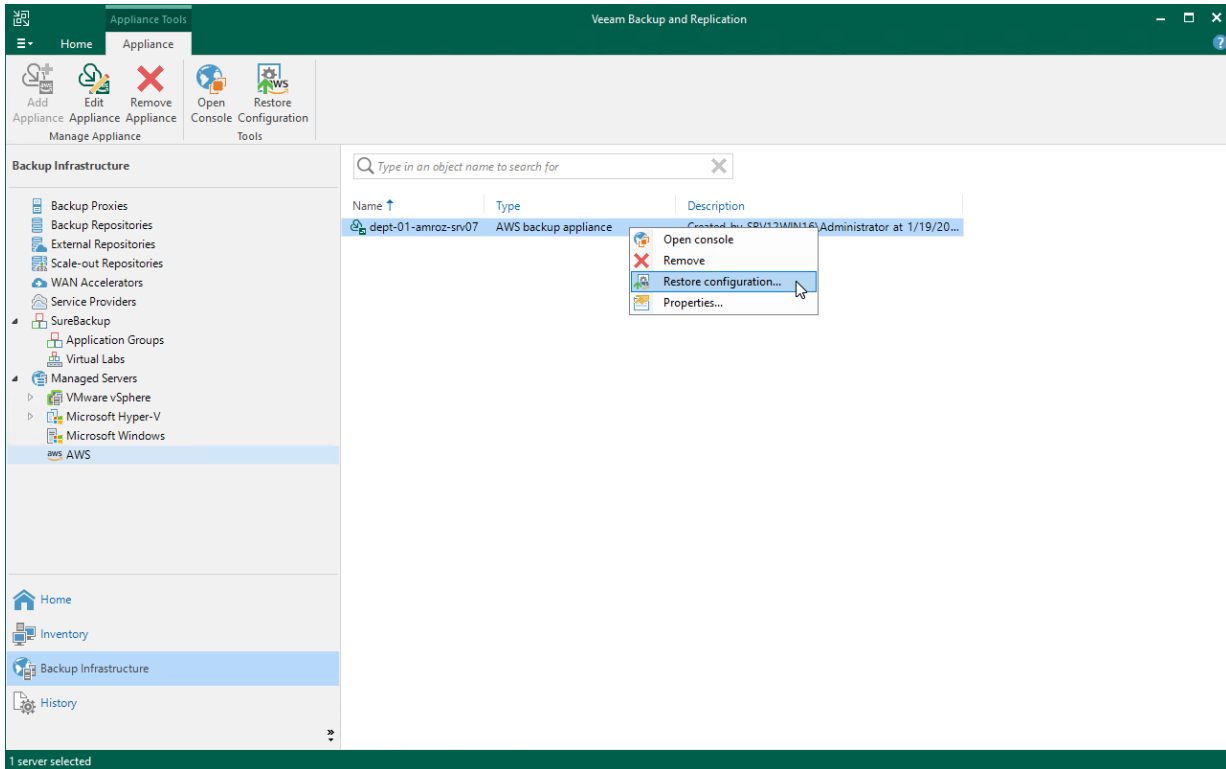
- If the backup appliance requires an upgrade, perform it before you start configuration restore. Otherwise, Veeam Backup & Replication will not be able to perform the restore operation. To learn how to upgrade appliances, see [Updating Appliances Using Console](#).
- If you remove the backup appliance from the backup infrastructure, you will not be able to restore its configuration. However, you will be able to restore the configuration to another backup appliance currently added to the backup infrastructure.
- If you want to restore the configuration to another backup appliance, you must remove the initial appliance from the backup infrastructure beforehand.
- Make sure that repositories added to the restored backup appliance are not managed by any other backup appliances. Otherwise, retention sessions running on different appliances may corrupt backup files stored in the repositories, which may result in unpredictable data loss.
- The appliance to which you restore the configuration preserves its TLS certificate.
- [Applies only if you restore the configuration to another backup appliance] During restore, Veeam Backup & Replication removes the initial appliance and its repositories from the backup infrastructure. If the restore operation fails, re-add the appliance and its repositories to the backup infrastructure.

Step 1. Launch Configuration Restore Wizard

To launch the **Configuration Restore** wizard, do the following:

1. In the Veeam Backup & Replication console, open the **Backup Infrastructure** view.
2. Navigate to **Managed Servers > AWS**.
3. Select a backup appliance for which you want to perform the restore operation, and click **Restore Configuration** on the ribbon.

Alternatively, you can right-click the necessary appliance and select **Restore configuration**.



Step 2. Choose Backup File

At the **Configuration Backup** step of the wizard, do the following:

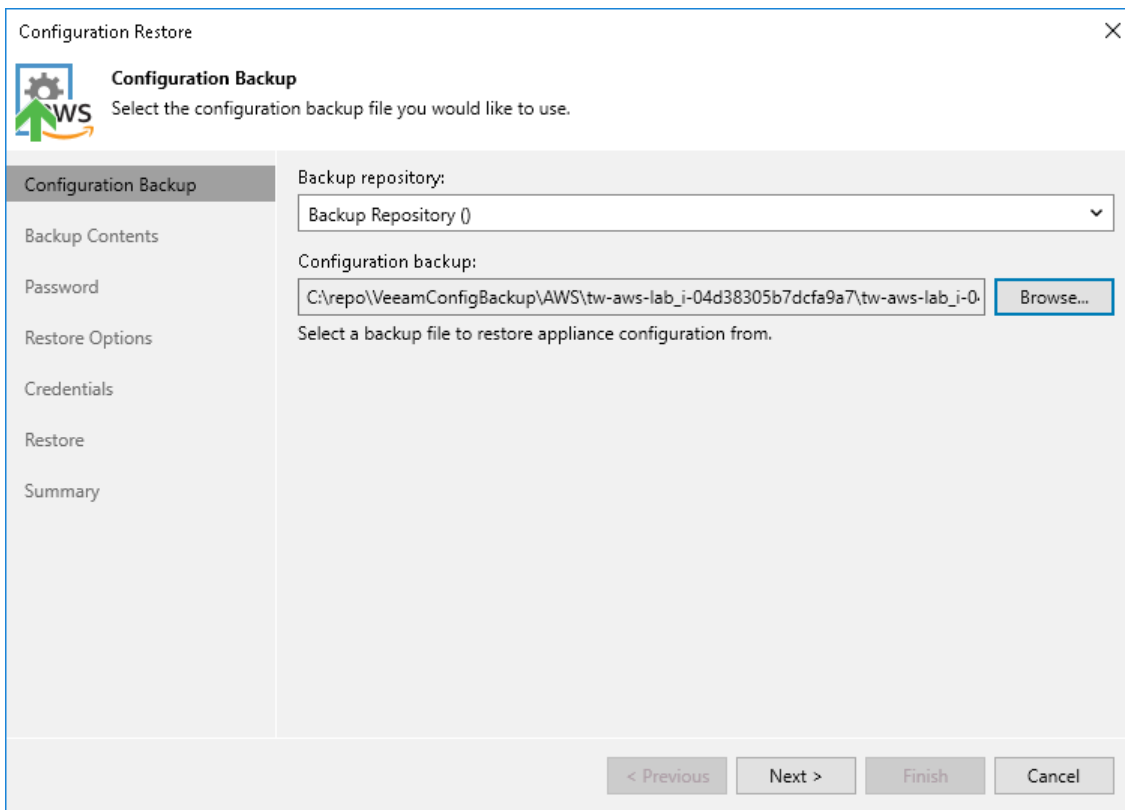
1. From the **Backup repository** list, select a repository where the configuration backup file is stored.

For a repository to be displayed in the **Backup repository** list, it must be added to the backup infrastructure as described Veeam Backup & Replication User Guide, section [Adding Backup Repositories](#).

2. Click **Browse** and select the necessary file.

NOTE

If the selected configuration backup file is not stored on the backup server, Veeam Backup & Replication will copy the file to a temporary folder on the server and automatically delete it from the folder as soon as the restore process completes.



The screenshot shows the 'Configuration Restore' wizard window, specifically the 'Configuration Backup' step. The window title is 'Configuration Restore' with a close button (X) in the top right corner. On the left side, there is a navigation pane with the following items: 'Configuration Backup' (selected), 'Backup Contents', 'Password', 'Restore Options', 'Credentials', 'Restore', and 'Summary'. The main area of the wizard is titled 'Configuration Backup' and contains the following elements:


- A sub-header 'Configuration Backup' with a gear icon and the text 'Select the configuration backup file you would like to use.'
- A 'Backup repository:' dropdown menu currently showing 'Backup Repository ()'.
- A 'Configuration backup:' text box containing the path 'C:\repo\VeeamConfigBackup\AWS\tw-aws-lab_j-04d38305b7dcfa9a7\tw-aws-lab_j-0-'. To the right of this text box is a 'Browse...' button.
- Below the text box is the instruction: 'Select a backup file to restore appliance configuration from.'
- At the bottom of the wizard, there are four buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'.

Step 3. Review Backup File Info

At the **Backup Contents** step of the wizard, Veeam Backup & Replication will analyze the content of the selected backup and display the following information:

- Backup file – the date and time when the backup file was created, the size of the file, the file location and so on.
- [Applies If the configuration backup file selected at step 2 is not stored on the backup server] Downloaded backup file – the temporary location of the configuration backup file on the backup server.
- Product – the name of the product and its version that was installed on the initial appliance.
- Catalogs – configuration data saved in the file (such as the number of configured backup policies, added user accounts, created repositories, logged session records and so on).

At the **Backup Contents** step of the wizard, review the provided information and click **Next** to confirm that you want to use the selected file to restore the configuration data.



Step 4. Specify Password

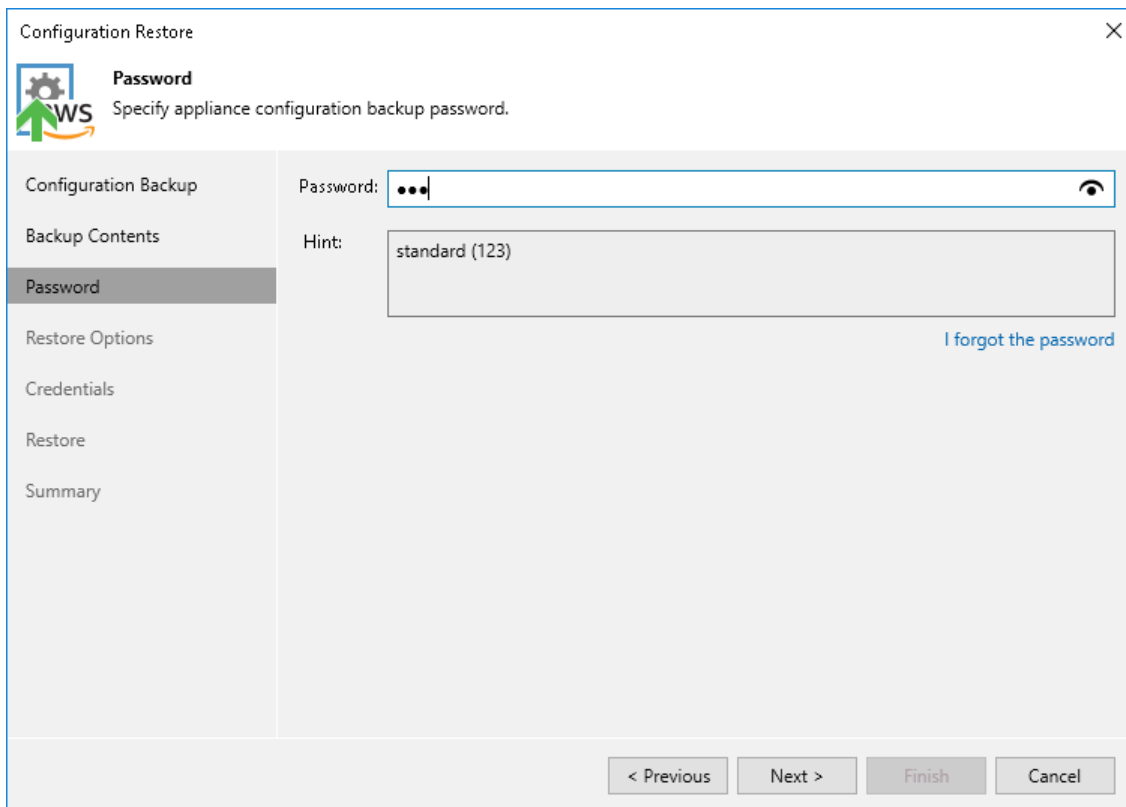
At the **Password** step of the wizard, specify a password used to encrypt the configuration backup.

If you do not remember the password, you can restore configuration backup data without providing it. To do that, click the **I forgot the password** link and follow the instructions provided in the Veeam Backup & Replication User Guide, section [Decrypting Data Without Password](#).

NOTE

To restore configuration data without a password, the following requirements must be met:

- You must have either the Veeam Universal License or a legacy socket-based license (Enterprise edition or higher) installed on the backup server.
- The backup server must be connected to Veeam Backup Enterprise Manager, and password loss protection must be enabled on the Veeam Backup Enterprise Manager side for the duration of both the backup and restore operations. For more information, see the [Veeam Backup Enterprise Manager Guide](#).



The screenshot shows the 'Configuration Restore' wizard window. The title bar reads 'Configuration Restore' with a close button (X) on the right. Below the title bar is a navigation pane on the left with the following items: 'Configuration Backup', 'Backup Contents', 'Password' (highlighted), 'Restore Options', 'Credentials', 'Restore', and 'Summary'. The main area of the window is titled 'Password' and contains the instruction 'Specify appliance configuration backup password.' Below this instruction are two input fields: 'Password:' with a masked password field (three dots) and a toggle icon (eye), and 'Hint:' with a text box containing 'standard (123)'. A blue link 'I forgot the password' is located to the right of the hint field. At the bottom of the window are four buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'.

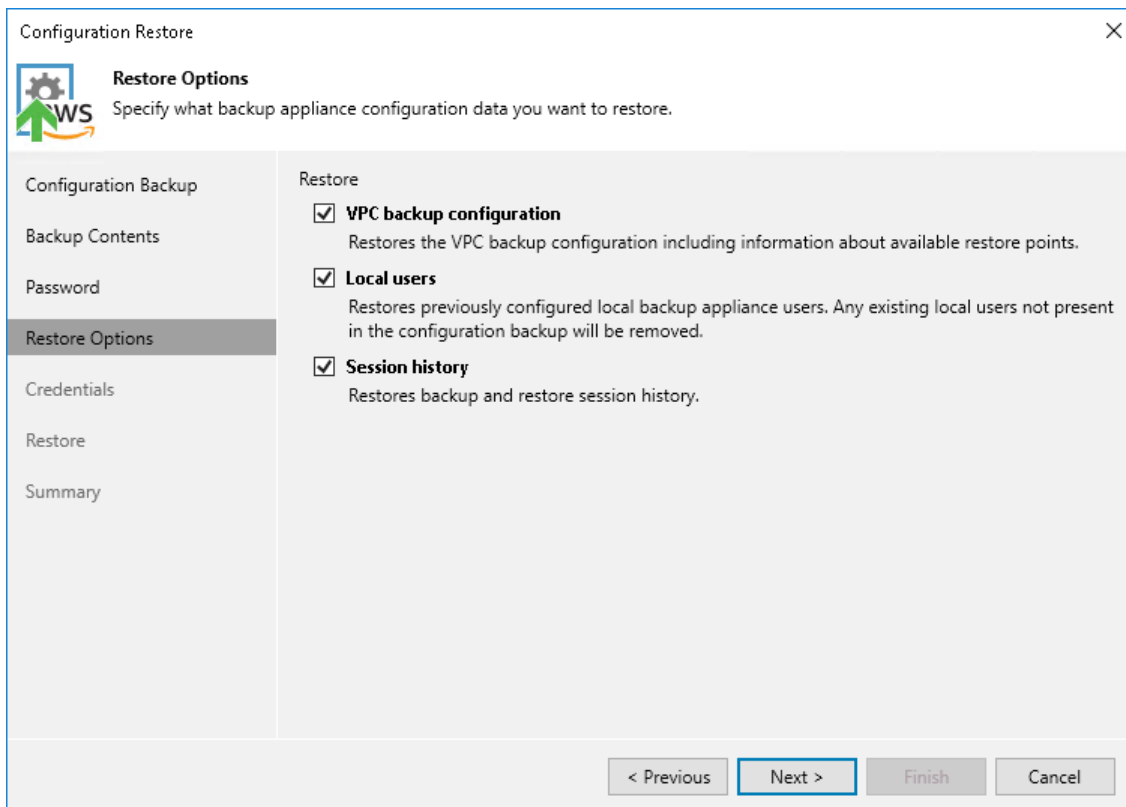
Step 5. Choose Restore Options

By default, Veeam Backup & Replication restores configuration data for the existing infrastructure components, created backup policies, configured global settings. At the **Restore Options** step of the wizard, you can choose whether you want to restore VPC configuration backups, portal users of the source backup appliance and session logs as well.

If you select the **VPC backup configuration** check box, Veeam Backup & Replication will restore VPC configurations of AWS Regions added to a backup policy running on the initial backup appliance and information on available restore points. If you select the **Local users** check box, Veeam Backup & Replication will restore all Portal Administrators, Portal Operators and Restore Operators saved to the configuration backup file – and overwrite the currently added portal users. If you select the **Session history** option, Veeam Backup & Replication will restore backup sessions, restore sessions, rescan sessions and service sessions – in this case, the restore process may take more time to complete.

IMPORTANT

After you click **Next**, the restore process will start. You will not be able to halt the process or edit the restore settings.



Step 6. Specify User Credentials

[This step applies only if you have selected the **Local users** option at the **Restore Options** step of the wizard]

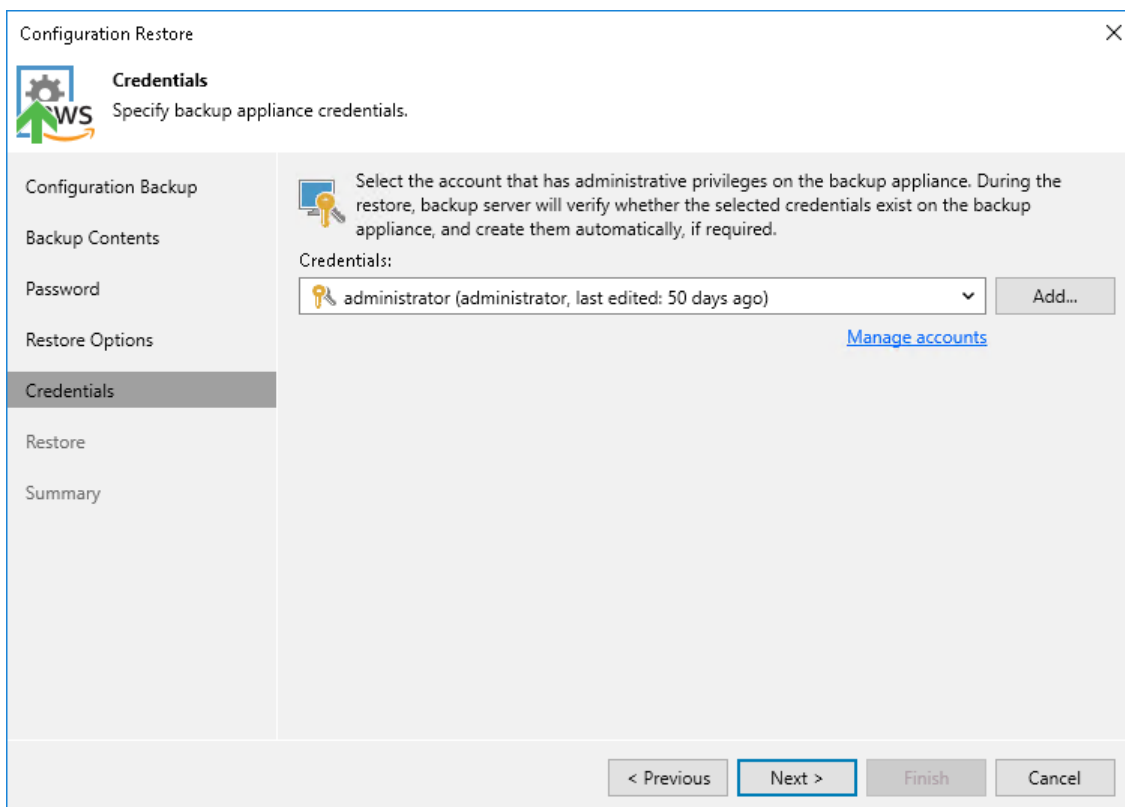
After the configuration restore process completes, Veeam Backup & Replication will try to connect to the backup appliance using credentials of the user specified [when adding the appliance](#) to the backup infrastructure. However, since you have chosen to restore all users saved to the configuration backup file, this user may be overwritten and Veeam Backup & Replication will fail to connect to the appliance.

That is why at the **Credentials** step of the wizard, you will be prompted to specify a user whose credentials Veeam Backup & Replication will use to connect to the backup appliance. You can specify a new or an existing user. If you specify an existing user, the user must have been assigned the *Portal Administrator* role on the initial appliance and the credentials of the user must match the credentials saved in the configuration backup file.

For a user to be displayed in the **Credentials** list, it must be added to the Credentials Manager as described in the Veeam Backup & Replication User Guide, section [Standard Accounts](#). If you have not added a user to the Credentials Manager beforehand, you can do it without closing the **Configuration Restore** wizard. To do that, click either the **Manage accounts** link or the **Add** button and specify a user name, password and description in the **Credentials** window.

IMPORTANT

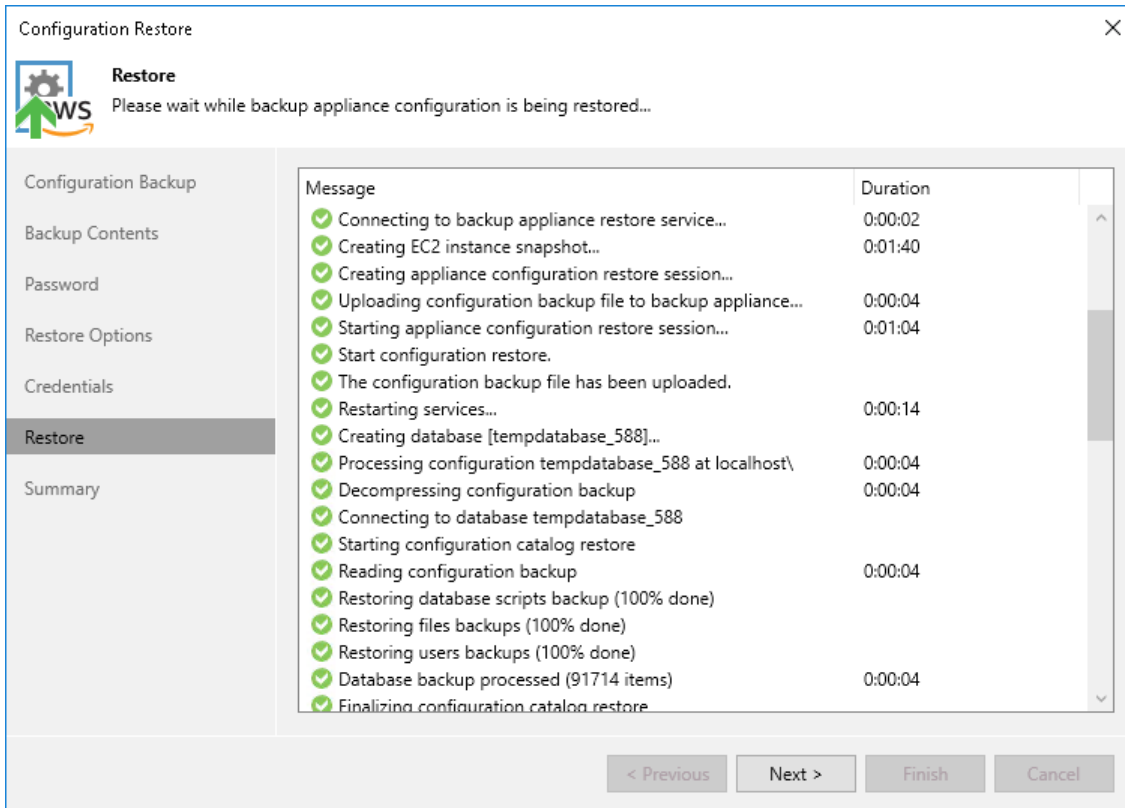
After you click **Next**, the restore process will start. You will not be able to halt the process or edit the restore settings.



The screenshot shows the 'Configuration Restore' wizard window, specifically the 'Credentials' step. The window title is 'Configuration Restore' with a close button (X) in the top right corner. On the left, there is a navigation pane with the following items: 'Configuration Backup', 'Backup Contents', 'Password', 'Restore Options', 'Credentials' (which is highlighted), 'Restore', and 'Summary'. The main area of the wizard is titled 'Credentials' and contains the instruction: 'Specify backup appliance credentials.' Below this, there is a key icon and the text: 'Select the account that has administrative privileges on the backup appliance. During the restore, backup server will verify whether the selected credentials exist on the backup appliance, and create them automatically, if required.' Underneath, there is a 'Credentials:' label followed by a dropdown menu showing 'administrator (administrator, last edited: 50 days ago)' and an 'Add...' button. A blue link labeled 'Manage accounts' is positioned below the dropdown. At the bottom of the wizard, there are four buttons: '< Previous', 'Next >' (which is highlighted with a blue border), 'Finish', and 'Cancel'.

Step 7. Track Progress

Veeam Backup & Replication will display the results of every step performed while executing the configuration restore. At the **Restore** step of the wizard, wait for the restore process to complete and click **Next**.



The screenshot shows the 'Configuration Restore' wizard window. The 'Restore' step is selected in the left-hand navigation pane. The main area displays a list of messages and their durations, indicating the progress of the restore process. The messages are as follows:


Message	Duration
✓ Connecting to backup appliance restore service...	0:00:02
✓ Creating EC2 instance snapshot...	0:01:40
✓ Creating appliance configuration restore session...	
✓ Uploading configuration backup file to backup appliance...	0:00:04
✓ Starting appliance configuration restore session...	0:01:04
✓ Start configuration restore.	
✓ The configuration backup file has been uploaded.	
✓ Restarting services...	0:00:14
✓ Creating database [tempdatabase_588]...	
✓ Processing configuration tempdatabase_588 at localhost\	0:00:04
✓ Decompressing configuration backup	0:00:04
✓ Connecting to database tempdatabase_588	
✓ Starting configuration catalog restore	
✓ Reading configuration backup	0:00:04
✓ Restoring database scripts backup (100% done)	
✓ Restoring files backups (100% done)	
✓ Restoring users backups (100% done)	
✓ Database backup processed (91714 items)	0:00:04
✓ Finalizing configuration catalog restore	

At the bottom of the window, there are four buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'. The 'Next >' button is highlighted, indicating it is the next step in the process.

Step 8. Finish Working with Wizard

At the **Summary** step of the wizard, click **Finish** to finalize the process of configuration data restore.

If Veeam Backup & Replication encounters an issue while performing configuration restore, the wizard will display the **Open backup appliance console and validate the restored configuration manually** link. This link redirects you to the Veeam Backup for AWS Web UI where you can view the details on the occurred issues. To learn how to resolve issues, see [Restoring Configuration Data Using Web UI](#).



Restoring Configuration Data Using Web UI

To restore the configuration database of a backup appliance using the Veeam Backup for AWS Web UI, do the following:

IMPORTANT

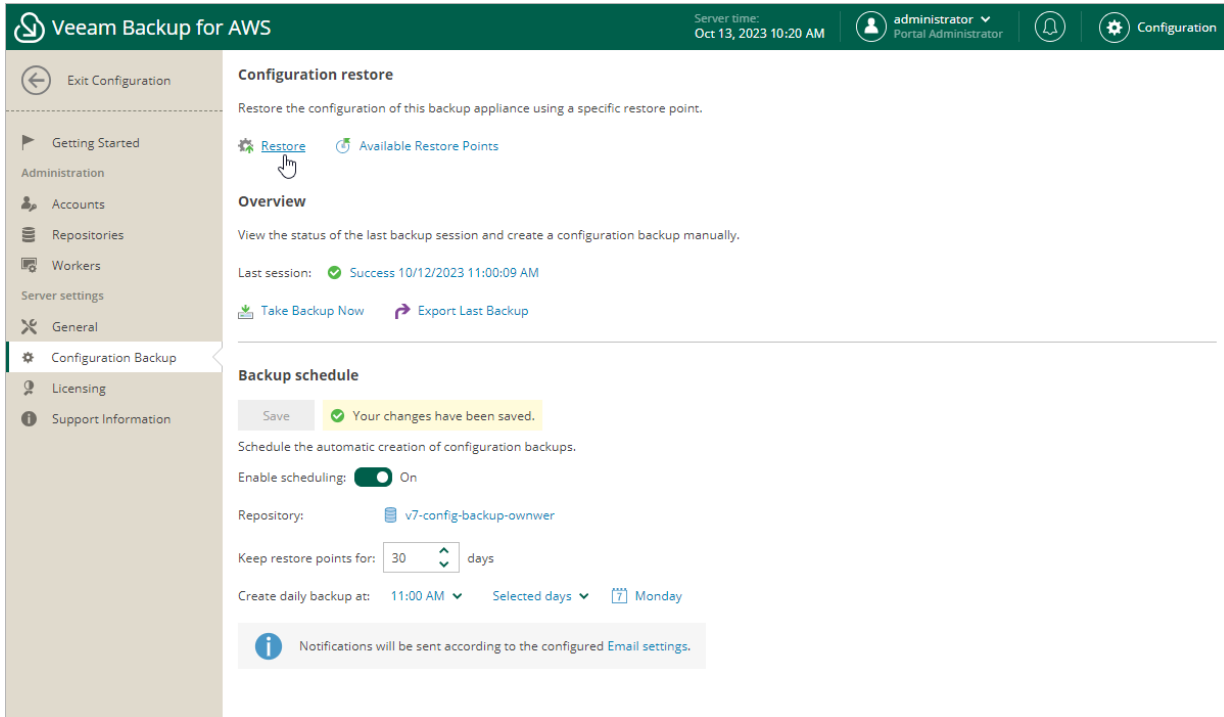
Before you start the restore process, stop all policies that are currently running.

1. [Launch the Configuration Restore wizard.](#)
2. [Choose a backup file.](#)
3. [Review the backup file info.](#)
4. [Choose restore options.](#)
5. [Track the restore progress.](#)
6. [View the results of verification steps.](#)
7. [Finish working with the wizard.](#)

Step 1. Launch Configuration Restore Wizard

To launch the **Configuration Restore** wizard, do the following:

1. Switch to the **Configuration** page.
2. Navigate to **Configuration Backup**.
3. In the **Configuration restore** section, click **Restore**.

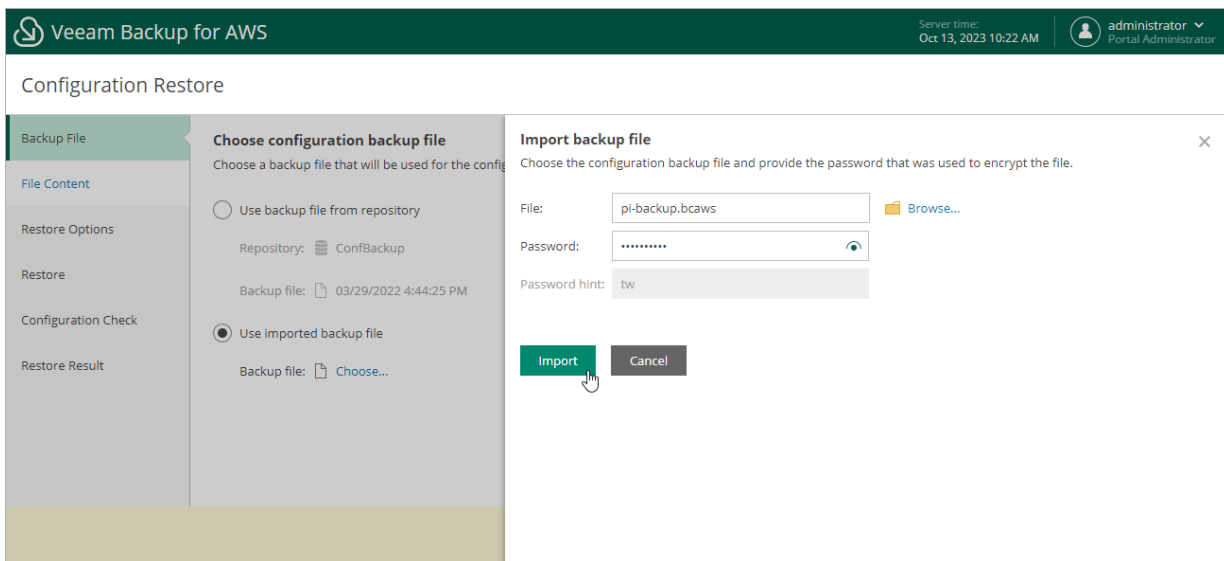


Step 2. Choose Backup File

At the **Backup File** step of the wizard, choose whether you want to use an exported backup file or a backup file stored in a backup repository.

- If you want to use a file stored in a backup repository, select the **Use backup file from repository** option and do the following:
 - a. Click the link in the **Repository** field, and use the list of available repositories in the **Choose repository** window to select the repository where the configuration backup file is stored.

For a backup repository to be displayed in the list of available repositories, it must be added to Veeam Backup for AWS as described in section [Adding Backup Repositories Using Web UI](#). The repository list shows only backup repositories that store configuration backup files.
 - b. Click the link in the **Backup file** field, select the necessary file in the **Choose backup file** window and click **Apply**.
- If you want to use a file that was exported from this or another backup appliance, select the **Use imported backup file** option, and do the following:
 - a. Click the link in the **Backup file** field.
 - b. In the **Import backup file** window, browse to the necessary backup file, provide the password that was used to encrypt the file, and click **Import**.



Step 3. Review Backup File Info

Veeam Backup for AWS will analyze the content of the selected backup file and display the following information:

- File information – the date and time when the backup file was created.
- Product information – the version of Veeam Backup for AWS that was installed on the initial backup appliance and the version of the File-Level Recovery service that was running on the appliance.

NOTE

Consider that if the current version of Veeam Backup for AWS installed on the backup appliance is later than the version saved in the configuration backup file, the configuration restore operation will not downgrade the backup appliance version.

- Product configuration – configuration data saved in the file (such as number of existing backup policies, added IAM roles and repositories, logged session records and so on).

At the **File Content** step of the wizard, review the provided information and click **Next** to confirm that you want to use the selected file to restore the configuration data.

The screenshot shows the 'Configuration Restore' wizard in Veeam Backup for AWS. The 'File Content' step is active, displaying the following information:

Review file content
Review the content of the selected configuration backup file.

File information

Restore point:	03/29/2022 4:44:25 PM
----------------	-----------------------

Product information

Product name:	Veeam Backup for AWS
Product version:	5.0.0.xxx
File-level recovery service version:	5.0.0.xxx

Product configuration

Standard repositories:	2
Archive repositories:	1
IAM roles:	3
EC2 backup policies:	1
VPC backup policy:	1
EFS backup policies:	1
Sessions:	3040


At the bottom of the wizard, there are three buttons: 'Previous', 'Next', and 'Cancel'. The 'Next' button is highlighted, indicating it is the current step.

Step 4. Choose Restore Options

By default, Veeam Backup for AWS restores only configuration data for the existing infrastructure components, created backup policies and configured global settings. At the **Restore Options** step of the wizard, you can choose whether you want to restore session logs, user accounts of the initial backup appliance and VPC configuration backups as well.

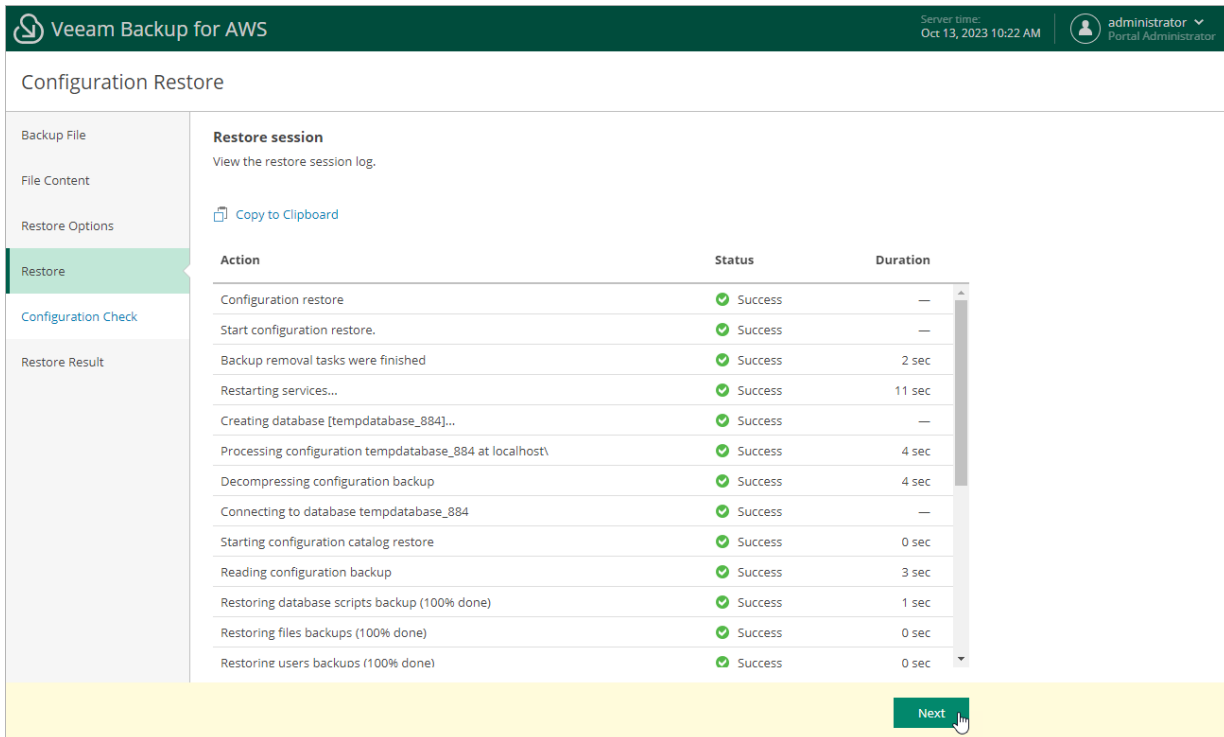
IMPORTANT

After you click **Restore**, the restore process will start. You will not be able to halt the process or edit the restore settings.



Step 5. Track Restore Progress

Veeam Backup for AWS will display the results of every step performed while executing the configuration restore. At the **Restore** step of the wizard, wait for the restore process to complete and click **Next**.



The screenshot shows the Veeam Backup for AWS interface during a configuration restore. The top navigation bar includes the Veeam logo, the product name, the server time (Oct 13, 2023 10:22 AM), and the user role (administrator Portal Administrator). The main content area is titled "Configuration Restore" and features a left-hand navigation menu with options: Backup File, File Content, Restore Options, Restore (highlighted), Configuration Check, and Restore Result. The main panel displays a "Restore session" summary with a "Copy to Clipboard" button. Below this is a table showing the progress of various actions, all of which are marked as "Success". A "Next" button is visible at the bottom right of the main panel.

Action	Status	Duration
Configuration restore	Success	—
Start configuration restore.	Success	—
Backup removal tasks were finished	Success	2 sec
Restarting services...	Success	11 sec
Creating database [tempdatabase_884]...	Success	—
Processing configuration tempdatabase_884 at localhost\	Success	4 sec
Decompressing configuration backup	Success	4 sec
Connecting to database tempdatabase_884	Success	—
Starting configuration catalog restore	Success	0 sec
Reading configuration backup	Success	3 sec
Restoring database scripts backup (100% done)	Success	1 sec
Restoring files backups (100% done)	Success	0 sec
Restoring users backups (100% done)	Success	0 sec

Step 6. View Configuration Check Results

After the restore process is over, Veeam Backup for AWS will run a number of verification checks to confirm that the configuration data has been restored successfully. At the **Configuration Check** step of the wizard, wait for the verification checks to complete and check whether Veeam Backup for AWS encountered any configuration issues.

If Veeam Backup for AWS encounters an issue while performing a verification check, the **Result** column will display a description of the issue, and the **Action** column will provide instructions on how to resolve it. For example, to resolve the issue with IAM role permissions, do the following:

1. In the **Action** column, click **View** in the **Role permissions** field.
2. In the **IAM role permissions** window, review IAM roles that are missing permissions required to perform operations, and choose one of the following options:
 - If you do not plan to use an IAM role to perform Veeam Backup for AWS operations, skip the notification and, after the configuration restore operation completes, specify a new role in the repository, policy and worker settings shown in the **Used As** column.
 - If you want to grant the missing permissions to an IAM role in the AWS Management Console, select the necessary role and click **Export Missing Permissions** to download the full list of missing permissions as a single JSON policy document.
 - If you want to instruct Veeam Backup for AWS to assign the missing permissions to an IAM role, select the necessary role and click **Grant**.

In the **Grant permissions** window, provide one-time access keys of an IAM user that is authorized to update permissions of IAM roles, and then click **Grant**.

The IAM user must have the following permissions:

```
"iam:AttachRolePolicy",
"iam:CreatePolicy",
"iam:CreatePolicyVersion",
"iam:CreateRole",
"iam:GetAccountSummary",
"iam:GetPolicy",
"iam:GetPolicyVersion",
"iam:GetRole",
"iam:ListAttachedRolePolicies",
"iam:ListPolicyVersions",
"iam:SimulatePrincipalPolicy",
"iam:UpdateAssumeRolePolicy"
```

NOTE

Veeam Backup for AWS does not store one-time access keys in the configuration database.

After you resolve all issues, click **Recheck** to ensure the backup appliance is now fully functional, and click **Next**.

IMPORTANT

Restored repositories must not be managed by multiple backup appliances simultaneously – retention sessions running on different backup appliances may corrupt backup files stored in the repositories, which may result in unpredictable data loss. That is why Veeam Backup for AWS verifies whether the restored backup repositories are managed by any backup appliances – but only for those repositories that were added to Veeam Backup for AWS version 7.0 or later. If the backup repositories are already managed by any backup appliances, Veeam Backup for AWS encounters an issue while performing a verification check. To resolve the issue, you must change the owner of these repositories to complete the restore session. To do that, in the **Action** column, click **View** in the **Repositories ownership** field. Then, click **Take Ownership** in the **Repository ownership** window.

The screenshot shows the Veeam Backup for AWS Configuration Restore interface. The main window displays a table of verification steps. The 'Role permissions' step is marked as 'Failed'. A modal dialog titled 'Grant Permissions' is open, showing a table of IAM roles with missing permissions. The dialog also includes a 'Provide temporary credentials' section with fields for 'Access key' and 'Secret key'.

Configuration Restore

Backup File
File Content
Restore Options
Restore
Configuration Check
Restore Result

Verification steps
The check will confirm that the configuration has been restored successfully.

Recheck Export

IAM role permissions
The following IAM roles are missing permissions. To allow Veeam Backup for AWS to perform backup and restore operations, grant the missing permissions to the IAM roles.

Grant Export Missing Permissions Recheck

<input type="checkbox"/>	IAM Role ↑	AWS Account	Used As	Result
Selected: 1 of 2				
<input checked="" type="checkbox"/>	Default Backup R...	537095525393 (vee...	Repository role, P...	Missing permissi...
<input type="checkbox"/>	Service role	537095525393 (vee...	Repository role, S...	Missing permissi...

Grant Permissions

Provide temporary credentials

You can grant permissions automatically using the form below or manually in the AWS Management Console. These keys will not be saved or stored. To learn how to assign missing permissions to IAM roles, see the [User Guide](#).

Access key: AKIAX2DLKQIWIUFFYT4

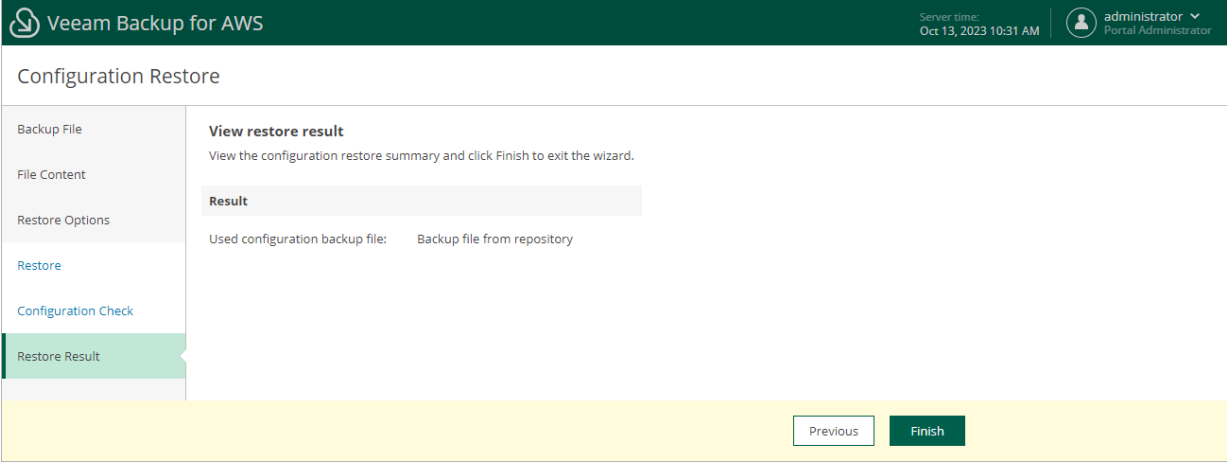
Secret key:

Grant Cancel

Close

Step 7. Finish Working with Wizard

At the Summary step of the wizard, click **Finish** to finalize the process of configuration data restore.



Viewing Available Resources

After you create a backup policy to protect a specific type of AWS resources (EC2 instances, RDS resources, DynamoDB tables, Redshift clusters, EFS or FSx file systems), Veeam Backup for AWS rescans AWS Regions specified in the policy settings and populates the resource list on the **Resources** page with all resources of that type residing in these regions. If an AWS Region is no longer specified in any configured backup policy, Veeam Backup for AWS removes all resources residing in the region from the list of available resources.

The **Resources** page displays AWS resources that can be protected by Veeam Backup for AWS. Each resource is represented with a set of properties, such as:

- **Instance, Cluster or Name** – the name of the resource.
- **Instance ID, Table ID, Cluster ID or File System ID** – the unique identification number of the resource.
- **Instance Size, Source Size, Cluster Size or Table Size** – the size of the resource storage.

NOTE

Veeam Backup for AWS does not show sizes of Aurora DB clusters due to AWS REST API limitations.

- **AWS Account** – the AWS account to which the resource belongs.
- **Region** – the AWS Region where the resource resides.
- **Last Backup** – the date and time of the latest restore point created for the resource (if any).
- **Backup Policy** – the name of the backup policy that protects the resource (if any).
- **Restore Points** – the number of restore points created for the resource (if any).
- **Destination** – types of restore points created for the EC2 or RDS resource (if any).

On the **Resources** page you can also perform the following actions:

- Manually create cloud-native snapshots of RDS and EC2 instances, as well as backups of Redshift clusters, DynamoDB tables, EFS file systems and FSx file systems. For more information, see sections [Creating EC2 Snapshots Manually](#), [Creating RDS Snapshots Manually](#), [Creating Redshift Backups Manually](#), [Creating DynamoDB Backups Manually](#), [Creating EFS Backups Manually](#) and [Creating FSx Backups Manually](#).
- Add resources to existing backup policies. For more information, see [Adding Resources to Policy](#).
- Restore entire EC2 instances, EBS volumes attached to EC2 instances, as well as individual files and folders of EC2 instances.

To do that, select an EC2 instance, click the link in the **Restore Points** column, select the necessary restore point and click **Restore** > **Instance Restore**, **Volume Restore** or **File-level Recovery** in the **Available Restore Points** window. Then, complete the wizard as described in section [Performing Entire EC2 Instance Restore](#), [Performing Volume-Level Restore](#) or [Performing File-Level Recovery](#).

- Restore entire DB instances, specific DB instance databases and Aurora DB clusters.

To do that, select the RDS resource, click the link in the **Restore Points** column, select the necessary restore point and click **Restore** > **Instance Restore** or **Database Restore** in the **Available Restore Points** window. Then, complete the wizard as described in section [Performing RDS Instance Restore](#) or [Performing Database Restore](#).

- Restore DynamoDB tables.

To do that, select the DynamoDB table, click the link in the **Restore Points** column, select the necessary restore point and click **Restore** in the **Available Restore Points** window. Then, complete the wizard as described in section [DynamoDB Restore Using Web UI](#).

- Restore Redshift clusters.

To do that, select the Redshift cluster, click the link in the **Restore Points** column, select the necessary restore point and click **Restore** in the **Available Restore Points** window. Then, complete the wizard as described in section [Redshift Restore Using Web UI](#).

- Restore entire EFS file systems, as well as individual files and folders stored in file systems.

To do that, select the EFS file system, click the link in the **Restore Points** column, select the necessary restore point and click **Restore > Entire EFS** or **File-level Recovery** in the **Available Restore Points** window. Then, complete the wizard as described in section [Performing Entire File System Restore](#) or [Performing File-Level Recovery](#).

- Restore FSx file systems.

To do that, select the FSx file system, click the link in the **Restore Points** column, select the necessary restore point and click **Restore** in the **Available Restore Points** window. Then, complete the wizard as described in section [FSx Restore Using Web UI](#).

- Remove all cloud-native snapshots created for EC2 instances, DB instances or Aurora DB clusters manually, as well as remove all backups created for DynamoDB tables, Redshift clusters, EFS file systems and FSx file systems manually.

To do that, select the necessary resource, click the link in the **Restore Points** column. Then, select the necessary manual snapshot or backup you want to remove in the **Available Restore Points** window, and click **Remove Manual Snapshot** or **Remove Manual Backup**.

- Retrieve archived data from EC2 backups that are stored in repositories of the S3 Glacier Flexible Retrieval or S3 Glacier Deep Archive storage class.

To do that, select the resource, click the link in the **Restore Points** column, select a restore point that contains the archived data you want to retrieve and click **Retrieve Backup** in the **Available Restore Points** window. Then, complete the wizard as described in section [Retrieving EC2 Data From Archive](#).

To extend time for which you want to keep the retrieved data available for restore operations, select the restore point that contains the retrieved data in the **Available Restore Points** window, and click **Extend Availability**. In the **Extend Data Availability Period** window, specify the number of days for which you want to keep the data available for restore operations, and click **Extend**.

Adding Resources to Policy

If you want to protect additional resources by configured backup policies, you can either [edit the backup policy settings](#), or quickly add the resources to the backup policies on the **Resources** tab.

To add a resource to a backup policy, do the following:

1. Navigate to **Resources**.
2. Switch to the necessary tab and select the resource that you want to protect by a backup policy.

For a resource to be displayed in the list of available resources, an AWS Region where the resource resides must be specified in any of configured backup policies that protects this kind of resources, and the IAM role specified in the backup policy settings must have permissions to access the resource.


3. Click **Add to Policy**.

4. In the **Add to Policy** window:

- a. Choose the backup policy that must protect the selected resource and click **Add**.

For a backup policy to be displayed in the list of available policies, an AWS Region where the selected resource resides must be specified in the policy settings, and the IAM role used by Veeam Backup for AWS for this backup policy must have permissions to access the selected resource.

- b. Review the configured settings and click **OK**.



Performing Backup

With Veeam Backup for AWS, you can protect data in the following ways:

- **Create cloud-native snapshots of EC2 instances and RDS resources**

A cloud-native snapshot of a EC2 instance includes point-in-time snapshots of EBS volumes attached to the processed instance. Snapshots of EBS volumes (also referred to as EBS snapshots) are taken using native [AWS capabilities](#).

A cloud-native snapshot of a DB instance includes a storage volume snapshot of the instance. Snapshots of DB instances (also referred to as DB snapshots) are taken using native [AWS capabilities](#).

A cloud-native snapshot of an Aurora DB cluster includes a storage volume snapshot of the cluster that backs up the entire cluster and not just individual databases. Snapshots of Aurora DB clusters (also referred to as DB cluster snapshots) are taken using native [AWS capabilities](#).

- **Replicate cloud-native snapshots to a remote site**

By default, cloud-native snapshots are stored only in the AWS Region where the processed instance resides. For enhanced data safety, you can instruct Veeam Backup for AWS to create copies of cloud-native snapshots and store them in any other AWS Region within any AWS account. You can also combine the snapshot replication functionality with various [data recovery options](#) to migrate instance data between AWS Regions and AWS accounts.

- **Create image-level backups of EC2 instances and RDS resources**

In addition to cloud-native snapshots, you can protect your EC2 and DB instances with image-level backups.

An image-level backup of an EC2 instance captures the whole image of the processed instance (including instance configuration, OS data, application data and so on) at a specific point in time. The backup is saved to a backup repository in the native Veeam format.

An image-level backup of a DB instance captures the PostgreSQL databases of the processed instance. The backup is saved to a backup repository in the native Veeam format.

- **Create backups of your Redshift clusters**

An Amazon Redshift backup captures the whole image of the Redshift cluster at a specific point of time. Redshift backups are taken using native [AWS capabilities](#).

- **Create backups and backup copies of your DynamoDB tables**

An Amazon DynamoDB backup captures the whole image of the DynamoDB table at a specific point of time. DynamoDB backups are taken using native [AWS capabilities](#).

By default, DynamoDB backups are stored only in the AWS Region where the processed resources reside. For enhanced data safety, you can instruct Veeam Backup for AWS to create copies of these backups and store them in any other AWS Region within the same AWS account.

- **Create backups and backup copies of your EFS and FSx file systems**

An Amazon EFS and FSx file system backup captures the whole image of the EFS and FSx file system (including file system configuration, files, directories and so on) at a specific point of time. EFS and FSx backups are taken using native [AWS capabilities](#).

By default, EFS and FSx backups are stored only in the AWS Region where the processed resources reside. For enhanced data safety, you can instruct Veeam Backup for AWS to create copies of these backups and store them in other AWS Regions within the same AWS account. For EFS file system, you can also combine the backup copy functionality with various [data recovery options](#) to migrate file system data between AWS Regions.

- **Create backups of your VPC configuration**

An Amazon VPC configuration backup captures the whole image of a VPC configuration of an AWS account (including multiple VPC configuration settings and components) at a specific point in time. By default, the VPC configuration backup is stored in the Veeam Backup for AWS database. For enhanced data safety, you can instruct Veeam Backup for AWS to create copies of Amazon VPC configuration backups and store them in a backup repository.

IMPORTANT

Veeam Backup for AWS does not support backup of the following VPC configuration components: VPC Traffic Mirroring, AWS Network Firewall, Route 53 Resolver DNS Firewall, AWS Verified Access, VPC Flow Logs, carrier gateways, customer IP pools, transit gateway policy tables, and core networks in route tables.

To schedule data protection tasks to run automatically, create backup policies. You will be able to [run the backup policies on demand](#) and manually perform backup of EC2 instances, RDS resources, DynamoDB tables, Redshift clusters, EFS and FSx file systems. To learn how to perform backup manually, see sections [Creating EC2 Snapshots Manually](#), [Creating RDS Snapshots Manually](#), [Creating DynamoDB Backups Manually](#), [Creating Redshift Backups Manually](#), [Creating EFS Backups Manually](#), [Creating FSx Backups Manually](#).

TIP

You can perform advanced data protection operations with image-level backups from the Veeam Backup & Replication console. For more information, see the Veeam Backup & Replication User Guide, section [External Repository](#).

Performing Backup Using Console

Veeam Backup for AWS runs backup policies for every data protection operation. A backup policy is a collection of settings that define the way backup operations are performed: what data to back up, where backups must be stored, when the backup process must start, and so on.

You can create multiple backup policies for AWS resources. One backup policy can be used to process multiple resources within different regions, but you can back up each resource with one backup policy at a time. For example, if an instance is added to more than one backup policy, it will be processed only by a backup policy that has the highest priority. Other backup policies will skip this instance from processing. For information on how to set a priority for a backup policy, see [Settings Policy Priority](#).

After you install AWS Plug-in for Veeam Backup & Replication and add backup appliances to the backup infrastructure, you can manage backup policies directly from the Veeam Backup & Replication console.

Creating Backup Policies

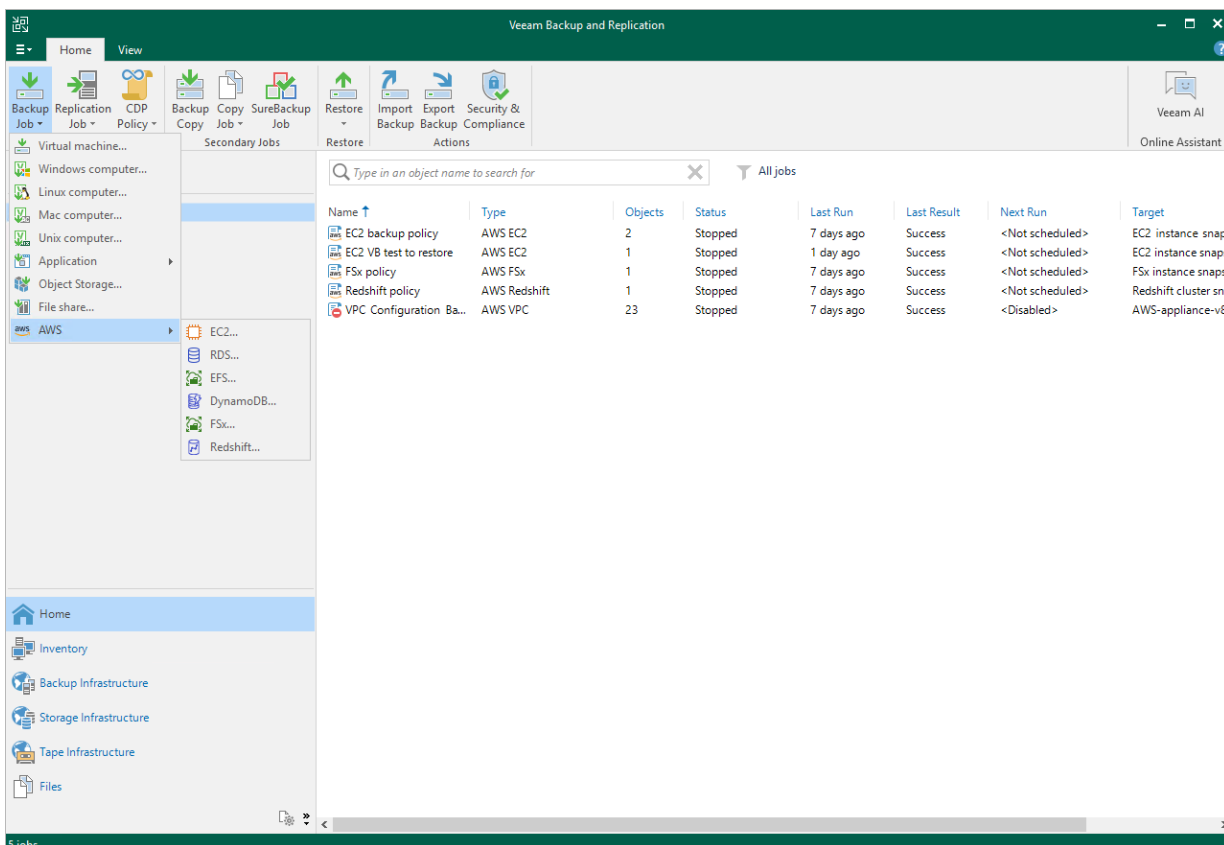
You can create backup policies in the Veeam Backup for AWS Web UI only. However, you can launch the **Add Policy** wizard directly from the Veeam Backup & Replication console – to do that, use either of the following options:

- Switch to the **Home** tab, click **Backup Job** on the ribbon, navigate to **AWS > EC2, RDS, EFS, DynamoDB, Redshift** or **FSx** and select the backup appliance on which you want to create the backup policy.
- Open the **Home** view, right-click **Jobs**, navigate to **Backup > AWS > EC2, RDS, EFS, DynamoDB, Redshift** or **FSx** and select the backup appliance on which you want to create the backup policy.

Veeam Backup & Replication will open the **Add EC2 Policy**, **Add RDS Policy**, **Add EFS Policy**, **Add DynamoDB Policy**, **Add Redshift Policy** or **Add FSx Policy** wizard in a web browser. Complete the wizard as described in section [Creating EC2 Backup Policies](#), [Creating RDS Backup Policies](#), [Creating EFS Backup Policies](#), [Creating DynamoDB Backup Policies](#), [Creating Redshift Backup Policies](#) or [Creating FSx Backup Policies](#).

NOTE

Backup appliance comes with a preconfigured VPC Configuration Backup policy that is disabled by default. To start protecting your Amazon VPC configuration, you must edit the VPC Configuration Backup policy settings and enable the policy. For more information, see [Editing VPC Configuration Backup Policy](#).



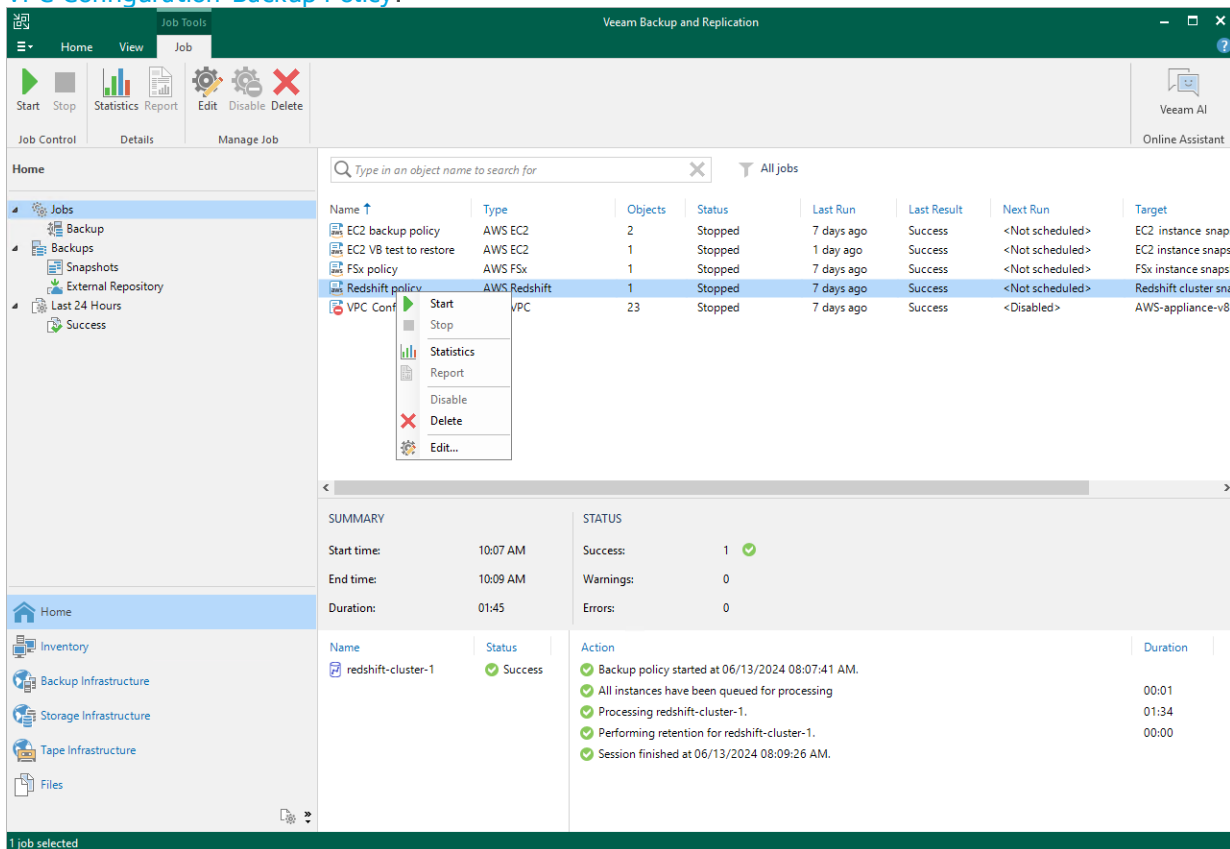
Editing Backup Policy Settings

You can edit backup policies in the Veeam Backup for AWS Web UI only. However, you can launch the edit policy wizard directly from the Veeam Backup & Replication console. To do that, do the following:

1. In the Veeam Backup & Replication console, open the **Home** view.
2. Navigate to **Jobs**.
3. Select the necessary backup policy and click **Edit** on the ribbon.

Alternatively, you can right-click the policy and select **Edit**.

Veeam Backup & Replication will open the **Edit Policy** wizard in a web browser. Complete the wizard as described in section [Creating EC2 Backup Policies](#), [Creating RDS Backup Policies](#), [Creating DynamoDB Backup Policies](#), [Creating Redshift Backup Policies](#), [Creating EFS Backup Policies](#), [Creating FSx Backup Policies](#) or [Editing VPC Configuration Backup Policy](#).



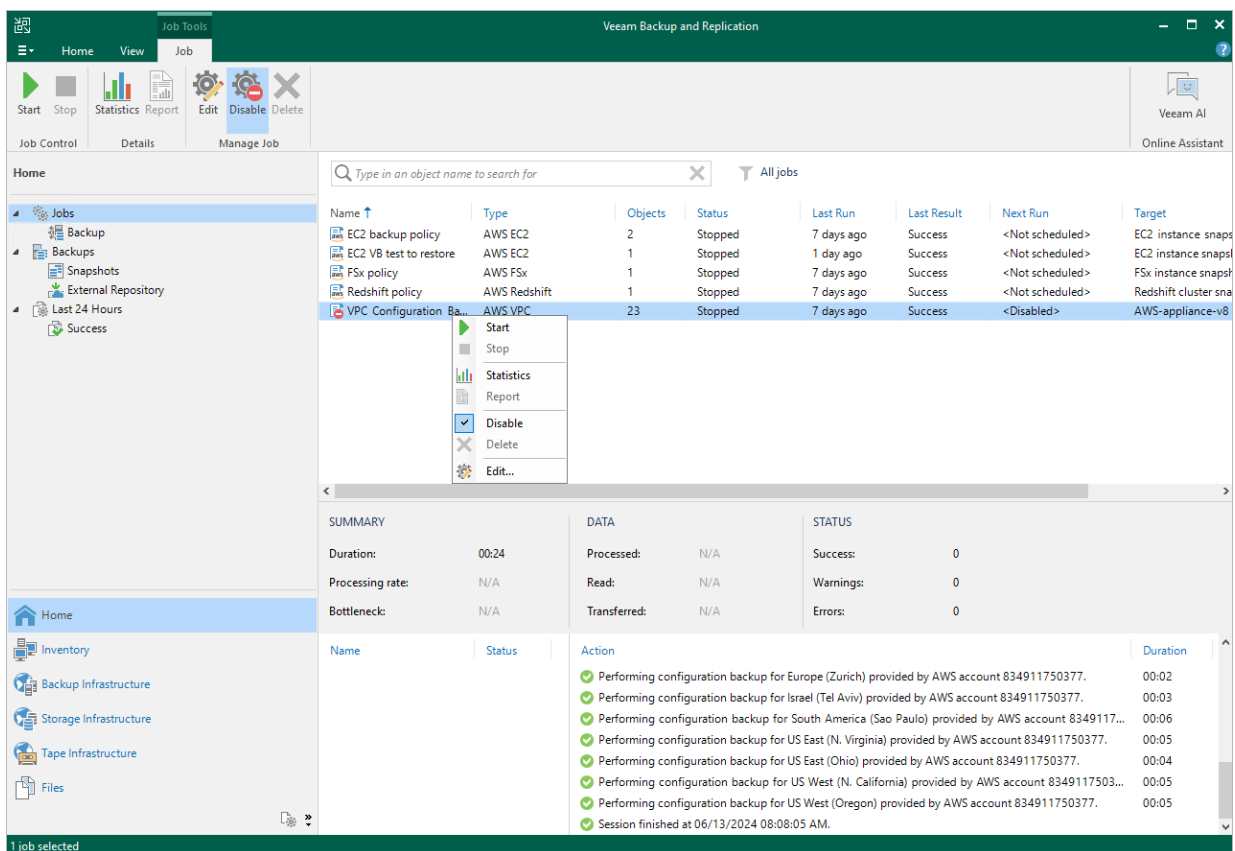
Enabling and Disabling Backup Policies

By default, Veeam Backup for AWS runs all created backup policies according to the specified schedules. However, you can temporarily disable a backup policy so that Veeam Backup for AWS does not run the backup policy automatically. You will still be able to [manually start](#) or enable the disabled backup policy at any time you need.

To disable an enabled backup policy or to enable a disabled backup policy, do the following:

1. In the Veeam Backup & Replication console, open the **Home** view.
2. Navigate to **Jobs**.
3. Select the necessary backup policy and click **Disable** on the ribbon.

Alternatively, you can right-click the necessary backup policy and select **Disable**.



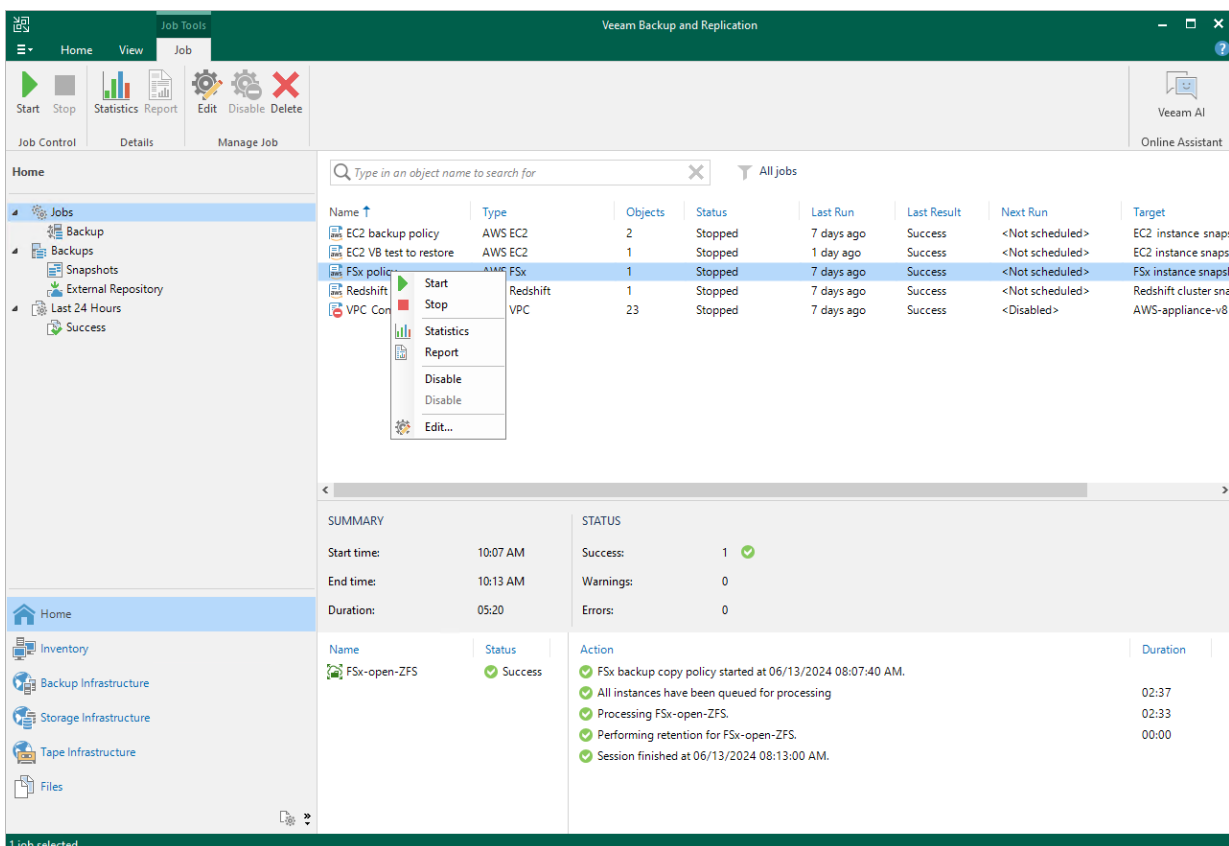
Starting and Stopping Backup Policies

You can start a backup policy manually, for example, if you want to create an additional restore point in the snapshot or backup chain and do not want to modify the configured backup policy schedule. You can also stop a backup policy if processing of an instance is about to take too long, and you do not want the policy to have an impact on the production environment during business hours.

To start or stop a backup policy, do the following:

1. In the Veeam Backup & Replication console, open the **Home** view.
2. Navigate to **Jobs**.
3. Select the necessary backup policy, and click **Start** or **Stop** on the ribbon.

Alternatively, you can right-click the selected policy, and select **Start** or **Stop**.



Deleting Backup Policies

Veeam Backup & Replication allows you to permanently delete backup policies created by Veeam Backup for AWS.

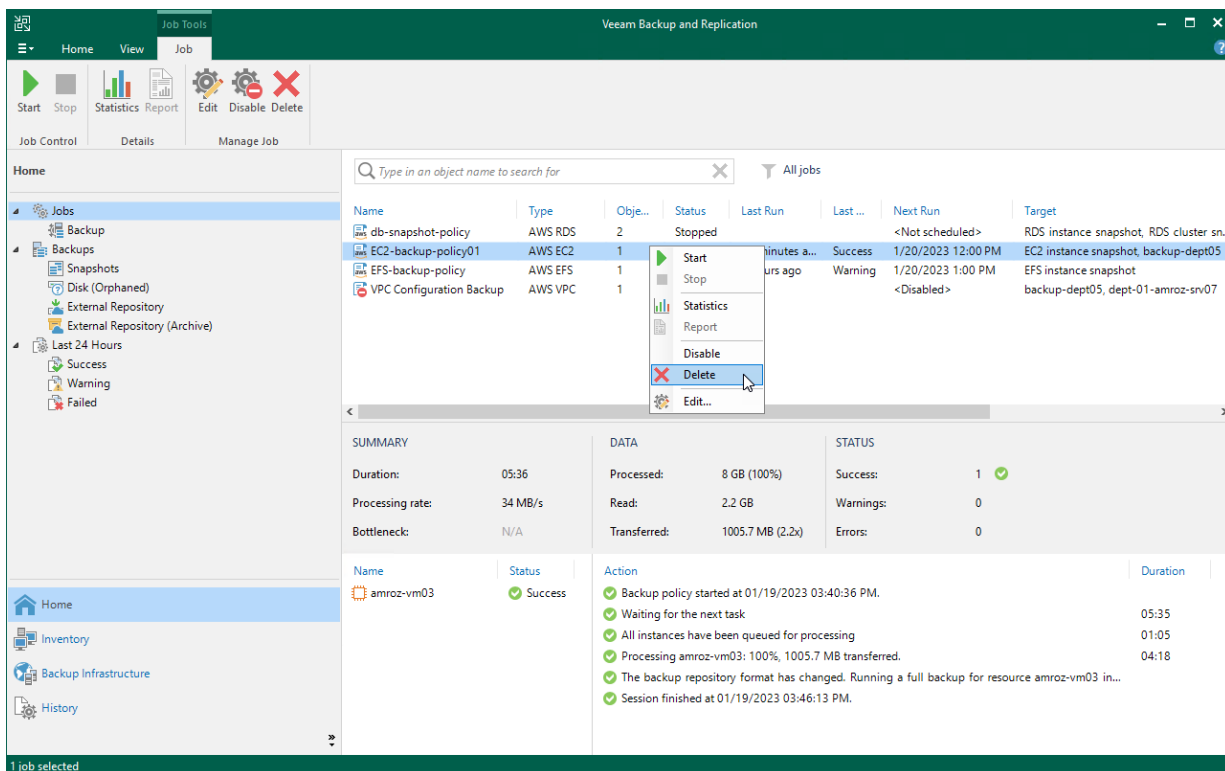
To delete a backup policy, do the following:

1. In the Veeam Backup & Replication console, open the **Home** view.
2. Navigate to **Jobs**.
3. Select the necessary backup policy and click **Delete** on the ribbon.

Alternatively, you can right-click the necessary backup policy and select **Delete**.

IMPORTANT

When you delete a backup policy from Veeam Backup & Replication, the policy is automatically deleted from the backup appliance as well.



Creating Backup Copy Jobs

Backup copy is a technology that helps you copy and store backed-up data of EC2 instances in different locations. Storing data in different locations increases its availability and ensures that data can be recovered in case a disaster strikes.

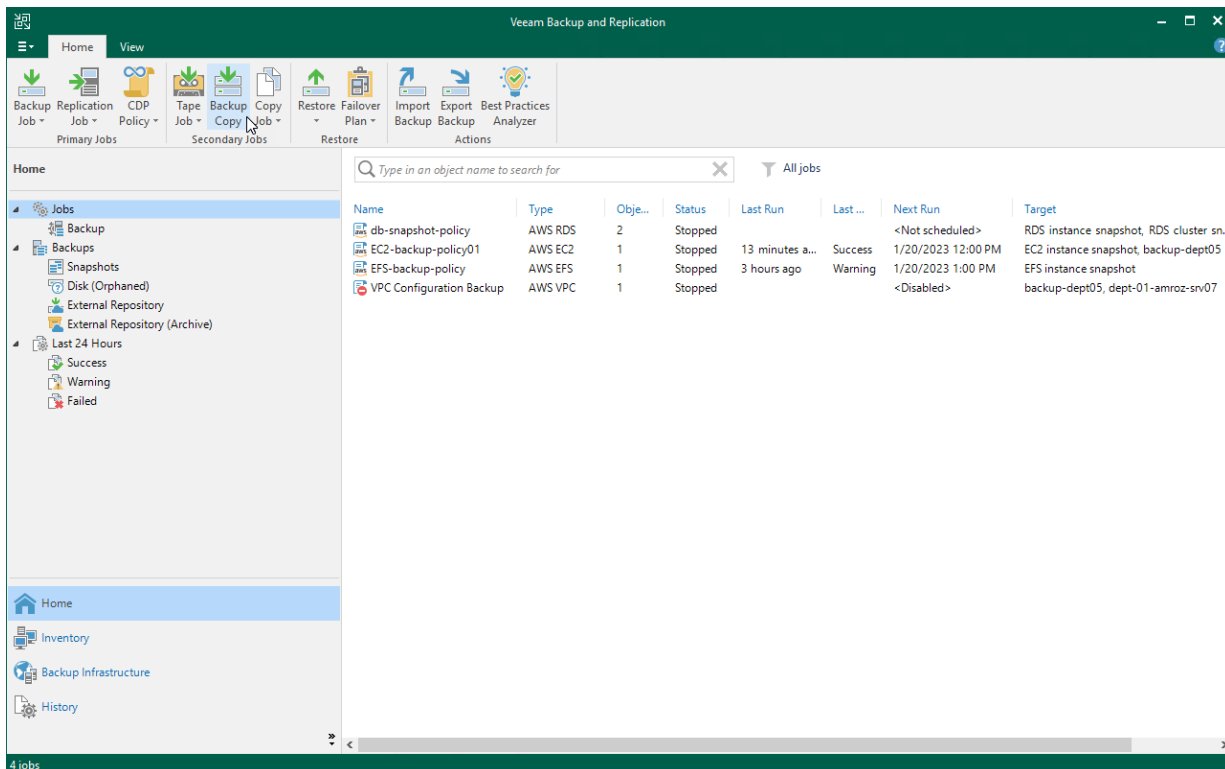
Backup-copy is a job-driven process. Veeam Backup & Replication fully automates the backup copy process and lets you specify retention settings to maintain the desired number of restore points, as well as full backups for archival purposes. For more information on the backup copy functionality, see the Veeam Backup & Replication User Guide, section [Backup Copy](#).

IMPORTANT

Backup copy can be performed only using EC2 backup files stored in standard backup repositories for which you have specified access keys of an IAM user whose permissions are used to access the repositories. To learn how to specify credentials for the repositories, see sections [Creating New Repositories](#) and [Connecting to Existing Appliances](#).

To create a backup copy job, do the following:

1. In the Veeam Backup & Replication console, open the **Home** view.
2. Click **Backup Copy** on the ribbon.
3. Complete the **New Backup Copy Job** wizard as described in the Veeam Backup & Replication User Guide, section [Creating Backup Copy Jobs for VMs and Physical Machines](#).




Copying Backups to Tapes

Veeam Backup & Replication allows you to automate copying of image-level backups of EC2 instances to tape devices and lets you specify scheduling, archiving and media automation options. For more information on supported tape libraries, see the Veeam Backup & Replication User Guide, section [Tape Devices Support](#).

Before you start copying backup to tapes:

- Copy EC2 instance backups to on-premises backup repositories. To learn how to copy backups, see the instructions provided in [Creating Backup Copy Jobs](#).
- Connect tape devices to Veeam Backup & Replication as described in the Veeam Backup & Replication User Guide, section [Tape Devices Deployment](#).
- Configure the tape infrastructure as described in steps 1-3 in the Veeam Backup & Replication User Guide, section [Getting Started with Tapes](#).

To copy EC2 instance backups to tapes, create a backup to tape job as described in the Veeam Backup & Replication User Guide, section [Creating Backup to Tape Jobs](#).



Performing Backup Using Web UI

Veeam Backup for AWS runs backup policies for every data protection operation. A backup policy is a collection of settings that define the way backup operations are performed: what data to back up, where to store backups, when to start the backup process, and so on.

One backup policy can be used to process multiple resources within different regions, but you can back up each resource with one backup policy at a time. For example, if an instance is added to more than one backup policy, it will be processed only by a backup policy that has the highest priority. Other backup policies will skip this instance from processing. For information on how to set a priority for a backup policy, see [Setting Policy Priority](#).

Performing EC2 Backup

One backup policy can be used to process one or more instances within one AWS account. The scope of data that you can protect in an AWS account is limited by permissions of an IAM role that is specified in the backup policy settings.

Before you create an EC2 backup policy, check the following prerequisites:

- If you plan to create image-level backups of EC2 instances, backup infrastructure components that will take part in the backup process must be added to the backup infrastructure and configured properly. These include [backup repositories](#) and [worker instances](#).
- If you plan to receive email notifications on backup policy results, configure [email notification settings](#) first.
- If you plan to create transactionally consistent backups of EC2 instances, check the requirements for [application-aware processing](#) and [guest scripting](#).

To schedule data protection tasks to run automatically, [create backup policies](#). For each protected EC2 instance, you can also [take cloud-native snapshots manually](#) when needed.

Creating EC2 Backup Policies

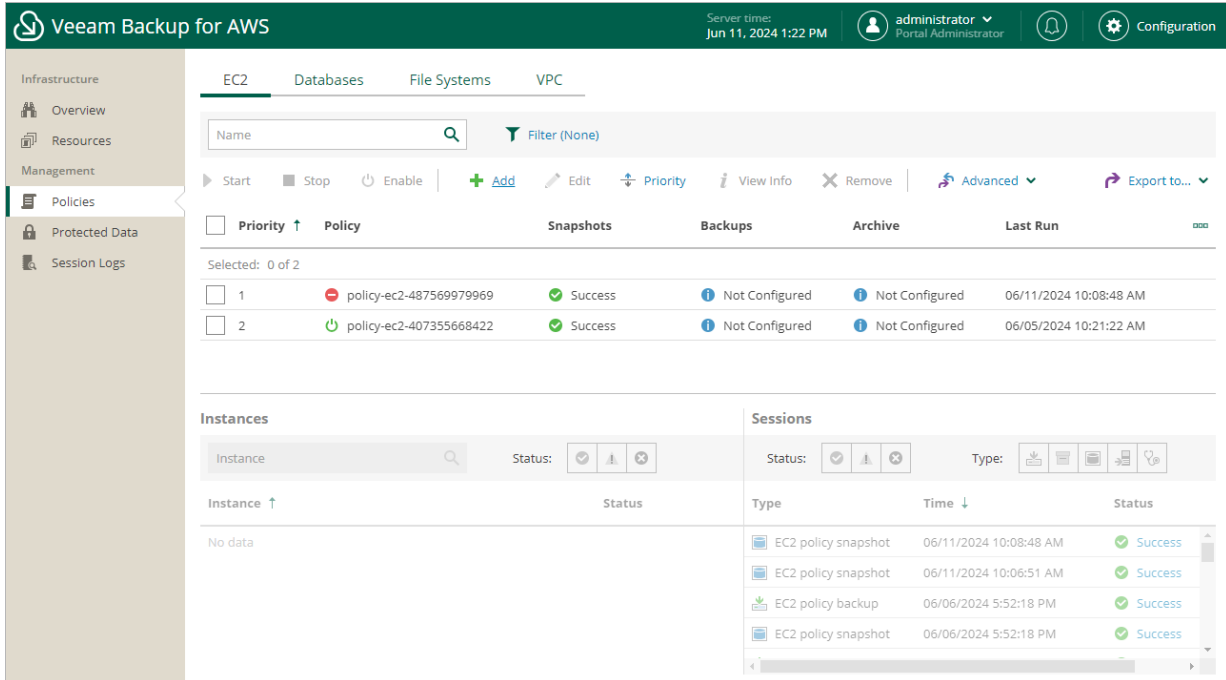
To create a backup policy, do the following:

1. [Launch the Add EC2 Policy wizard](#).
2. [Specify a backup policy name and description](#).
3. [Configure backup source settings](#).
4. [Enable guest processing](#).
5. [Configure backup target settings](#).
6. [Specify a schedule for the backup policy](#).
7. [Enable AWS tags assigning](#).
8. [Specify automatic retry, health check and notification settings for the backup policy](#).
9. [Review estimated cost of the selected EC2 instances](#).
10. [Finish working with the wizard](#).

Step 1. Launch Add EC2 Policy Wizard

To launch the **Add EC2 Policy** wizard, do the following:

1. Navigate to **Policies > EC2**.
2. Click **Add**.



Step 2. Specify Policy Name and Description

At the **Info** step of the wizard, use the **Name** and **Description** fields to specify a name for the new backup policy and to provide a description for future reference. The name must be unique in Veeam Backup for AWS; the maximum length of the name is 127 characters, the maximum length of the description is 255 characters.

The screenshot shows the 'Add EC2 Policy' wizard in Veeam Backup for AWS. The interface is in the 'Info' step, where the user is prompted to 'Specify policy name and description'. The 'Name' field contains 'ec2-frankfurt' and the 'Description' field contains 'EC2 instances protection in Frankfurt'. The 'Cost' is listed as 'N/A'. The 'Next' button is highlighted in green, and the 'Cancel' button is in grey. The left sidebar shows the wizard steps: Info, Sources, Guest Processing, Targets, Schedule, Tags, General Settings, Cost Estimation, and Summary.

Veeam Backup for AWS

Server time: Jun 11, 2024 1:23 PM

administrator Portal Administrator

Configuration

← Add EC2 Policy Cost: N/A ⚠

Info **Specify policy name and description**
Enter a name and description for the policy.

Sources

Guest Processing

Targets

Schedule

Tags

General Settings

Cost Estimation

Summary

Name: ec2-frankfurt

Description: EC2 instances protection in Frankfurt

Next Cancel

Step 3. Configure Backup Source Settings

At the **Sources** step of the wizard, specify backup source settings:

1. [Select an IAM role whose permissions will be used to perform EC2 instance backup.](#)
2. [Select AWS Regions where EC2 instances that you plan to back up reside.](#)
3. [Select EC2 instances to back up.](#)
4. [Select EBS volumes of the selected EC2 instances to exclude from the backup policy.](#)

Step 3.1 Specify IAM Role

In the **IAM role** section of the **Sources** step of the wizard, specify an IAM role whose permissions will be used to access AWS services and resources, and to create cloud-native snapshots of EC2 instances. The specified IAM role must belong to an AWS account in which the EC2 instances that you want to protect reside, and must be assigned the permissions listed in section [EC2 Backup IAM Role Permissions](#).

For an IAM role to be displayed in the **IAM role** list, it must be added to Veeam Backup for AWS with the *Amazon EC2 Backup* operation selected as described in section [Adding IAM Roles](#). If you have not added the necessary IAM role to Veeam Backup for AWS beforehand, you can do it without closing the **Add EC2 Policy** wizard. To do that, click **Add** and complete the [Add IAM Role](#) wizard.

IMPORTANT

It is recommended that you check whether the selected IAM role has all the permissions required to perform the operation. If some permissions of the IAM role are missing, the backup policy will fail to complete successfully. To run the IAM role permission check, click **Check Permissions** and follow the instructions provided in section [Checking IAM Role Permissions](#).

The screenshot shows the 'Add EC2 Policy' wizard in Veeam Backup for AWS. The 'Sources' step is active, and a 'Permission check' dialog is open. The dialog displays a success message: 'Your account meets the required permissions.' Below this, there are buttons for 'Grant', 'Recheck', and 'Export Missing Permissions'. A table lists various permissions, all of which are marked as 'Passed'.

Type	Status	Missing Permissions
IAM permissions	Passed	—
EVENTS permissions	Passed	—
SNS permissions	Passed	—
SQS permissions	Passed	—
EBS permissions	Passed	—
EC2 permissions	Passed	—
SSM permissions	Passed	—
SERVICEQUOTAS permissions	Passed	—
KMS permissions	Passed	—

Step 3.2 Select AWS Regions

In the **Specify region** section of the **Sources** step of the wizard, select AWS Regions where EC2 instances that you plan to back up reside:

1. Click **Choose regions**.
2. In the **Choose regions** window, select the necessary AWS Regions from the **Available Regions** list, and click **Add**.
3. To save changes made to the backup policy settings, click **Apply**.

The screenshot shows the Veeam Backup for AWS interface. The top navigation bar includes the Veeam logo, 'Veeam Backup for AWS', server time 'Jun 11, 2024 1:26 PM', user 'administrator Portal Administrator', and 'Configuration' settings. The main window is titled 'Add EC2 Policy' with a 'Cost: N/A' indicator. The left sidebar shows navigation options: Info, Sources (selected), Guest Processing, Targets, Schedule, Tags, General Settings, Cost Estimation, and Summary. The 'Specify source settings' section is active, showing 'IAM role' as 'admin-407355668422 (Created by nm at 9/28/2023)' and 'Regions' as 'Specify one or more regions.' with a 'Choose regions...' button. The 'Choose regions' dialog box is open, displaying a list of 'Available Regions (24)'. The regions listed are: Asia Pacific (Singapore), Asia Pacific (Sydney), Asia Pacific (Tokyo), Canada (Central), Canada West (Calgary), Europe (Ireland), Europe (London), Europe (Paris), Europe (Spain), Europe (Stockholm) (highlighted), Europe (Zurich), Israel (Tel Aviv), Middle East (UAE), and South America (Sao Paulo). To the right of the list is a 'Selected Regions (1)' box containing 'Europe (Frankfurt)'. There are 'Add' and 'Remove' buttons between the lists, and 'Apply' and 'Cancel' buttons at the bottom of the dialog.

Step 3.3 Select EC2 Instances

In the **Resources** section of the **Sources** step of the wizard, specify the backup scope – select EC2 instances that Veeam Backup for AWS will back up:

1. Click **Choose resources to protect**.
2. In the **Choose resource to protect** window, choose whether you want to back up all EC2 instances from AWS Regions selected at [step 3.2](#) of the wizard, or only specific EC2 instances.

If you select the **All resources** option, Veeam Backup for AWS will regularly check for new EC2 instances launched in the selected regions and automatically update the backup policy settings to include these instances into the backup scope.

If you select the **Protect only following resources** option, you must also specify the resources explicitly:

- a. Use the **Type** drop-down list to choose whether you want to add individual EC2 instances or AWS tags to the backup scope.

If you select the **Tag** option, Veeam Backup for AWS will back up only those EC2 instances that reside in the selected regions under specific AWS tags.

- b. Use the search field of the **Name or ID** drop-down list to find the necessary resource, and then click **Protect** to add the resource to the backup scope.

For a resource to be displayed in the list of available resources, it must reside in an AWS Region that has ever been specified in any backup policy. Otherwise, the only option to discover the available resources is to click **Browse to select specific resources from the global list** and to wait for Veeam Backup for AWS to populate the resource list.

TIP

You can simultaneously add multiple resources to the backup scope. To do that, click **Browse to select specific resources from the global list**, select check boxes next to the necessary EC2 instances or AWS tags in the list of available resources, and then click **Protect**.

If the list does not show the resources that you want to back up, click **Rescan** to launch the data collection process. As soon as the process is over, Veeam Backup for AWS will update the resource list.

If you add an AWS tag to the backup scope, Veeam Backup for AWS will regularly check for new EC2 instances assigned the added AWS tag and automatically update the backup policy settings to include these instances in the scope. However, this applies only to EC2 instances from the regions selected at [step 3.2](#) of the wizard. If you select a tag assigned to EC2 instances from other regions, these instances will not be protected by the backup policy. To work around the issue, either go back to [step 3.2](#) and add the missing regions, or create a new backup policy.

3. To save changes made to the backup policy settings, click **Apply**.

TIP

As an alternative to selecting the **Protect only following resources** option and specifying the resources explicitly, you can select the **All resources** option and exclude a number of resources from the backup scope. To do that, click **Choose resources to exclude** and specify the instances or tags that you do not want to protect – the procedure is the same as described for including resources in the backup scope.

Note that if a resource appears both in the list of included and excluded resources, Veeam Backup for AWS will still not process the resource because the list of excluded resources has a higher priority.

The screenshot shows the 'Add EC2 Policy' configuration window in Veeam Backup for AWS. The 'Choose resources to protect' dialog is open, showing the 'Protect only following resources' option selected. The dialog includes a search bar, a list of protected resources, and buttons for 'Apply' and 'Cancel'.

Item	ID	Value	Region
jf_vb_vpn	I-079d0447ec0d58f9a	—	Europe (Frankfurt)
pi-WIP-backup-frank-2	I-07ab618ed537ac329	—	Europe (Frankfurt)

Step 3.4 Select EBS Volumes

In the **Volumes** section of the **Sources** step of the wizard, you can exclude from processing EBS volumes attached to the selected EC2 instances:

1. Set the **Exclude volumes** toggle to *On*.
2. In the **Choose volumes to exclude** window, choose whether you want to exclude system volumes of the selected EC2 instances from processing.
3. To exclude specific EBS volumes, specify the EBS volumes explicitly:

- a. Use the **Type** list to choose whether you want to exclude individual EBS volumes or AWS tags from the backup scope.

If you select the **Tag** option, Veeam Backup for AWS will exclude from processing only those EBS volumes that reside in the selected regions under specific AWS tags.

- b. Use the search field to the right of the **Type** list to find the necessary resource, and then click **Exclude** to exclude the resource from the backup scope.

For a resource to be displayed in the list of available resources, it must reside in an AWS Region specified at [step 3.2](#) of the wizard. Consider that the list will display resources only if the region has ever been specified in any backup policy. Otherwise, the only option to discover the resources is to click **Browse to select specific resources from the global list** and to wait for Veeam Backup for AWS to rescan the region and to populate the resource list.

TIP

You can simultaneously exclude multiple resources from the backup scope. To do that, click **Browse to select specific resources from the global list**, select check boxes next to the necessary EBS volumes or AWS tags in the list of available resources, and then click **Exclude**.

If the list does not show the resources that you want to exclude, click **Rescan** to launch the data collection process. As soon as the process is over, Veeam Backup for AWS will update the resource list.

If you exclude an AWS tag from the backup scope, Veeam Backup for AWS will regularly check for new EBS volumes assigned the excluded AWS tag and automatically update the backup policy settings to exclude these volumes from the scope.

4. To save changes made to the backup policy settings, click **Apply**.

IMPORTANT

For Windows EC2 instances running VSS-aware applications, it is recommended that you do not exclude specific volumes other than system (root) volumes, since there is a limitation on the AWS System Manager side – only system volumes can be excluded. For more information on creating VSS snapshots, see [AWS Documentation](#).

The screenshot shows the 'Add EC2 Policy' configuration window in Veeam Backup for AWS. The 'Choose volumes to exclude' modal is open, displaying the following settings:

- Exclude system volumes:** Off
- Exclude specific volumes:**
 - Type: Volume
 - Volume ID: vol-086e8d4668a34a0a3
 - Exclude button: Exclude
- Excluded resources (1):**

Item	ID	Value	Region
—	vol-086e8d4668a34a0a3	—	Europe (Frankfurt)

Buttons for 'Apply' and 'Cancel' are visible at the bottom of the modal.

Step 4. Specify Guest Processing Settings

If you back up EC2 instances that are currently running, at the **Guest Processing** step of the wizard, you can configure guest processing settings. These settings allow you to specify what actions Veeam Backup for AWS will perform when communicating with the instance guest OS.

Particularly, you can specify the following guest processing settings:

- **Enable application-aware processing.** For Windows EC2 instances running VSS-aware applications, you can enable application-aware processing to ensure that the applications will be able to recover successfully, without data loss.

Application-aware processing is the Veeam technology based on Microsoft VSS. Microsoft VSS is responsible for quiescing applications on EC2 instances and creating a consistent view of application data. For more information on Microsoft VSS, see [Microsoft Docs](#).

- **Enable guest scripting.** For all processed EC2 instances, you can instruct Veeam Backup for AWS to run custom scripts on the instance before and after the backup operation. For example, for an EC2 instance running applications that do not support Microsoft VSS, Veeam Backup for AWS can execute a pre-snapshot script on the instance to quiesce these applications. This will allow Veeam Backup for AWS to create a transactionally consistent snapshot while no write operations occur on the instance volumes. After the snapshot is created, a post-snapshot script can start the applications again.

Limitations and Requirements

To be able to communicate with instance guest OSES, Veeam Backup for AWS uses the AWS Systems Manager (SSM) service. Thus, if you plan to enable guest processing for EC2 instances protected by the policy, you must consider the following:

- The backup appliance must have outbound internet access to the SSM service.
-
- EC2 instances must have the **443** network port opened for outbound internet access to the SSM service.
- The EC2 instances must be configured to communicate with AWS System Manager. To learn how to configure instance permissions for Systems Manager, see [AWS Documentation](#).
- SSM Agent must be installed on the EC2 instances. To learn how to install SSM Agent, see [AWS Documentation](#).

Note that SSM Agent is already preinstalled on EC2 instances launched from certain AMIs.

For more information on the SSM service, see [AWS Documentation](#).

Enabling Application-Aware Processing

To enable application-aware processing, at the **Guest Processing** step of the wizard, set the **Enable application-aware snapshots** toggle to *On*.

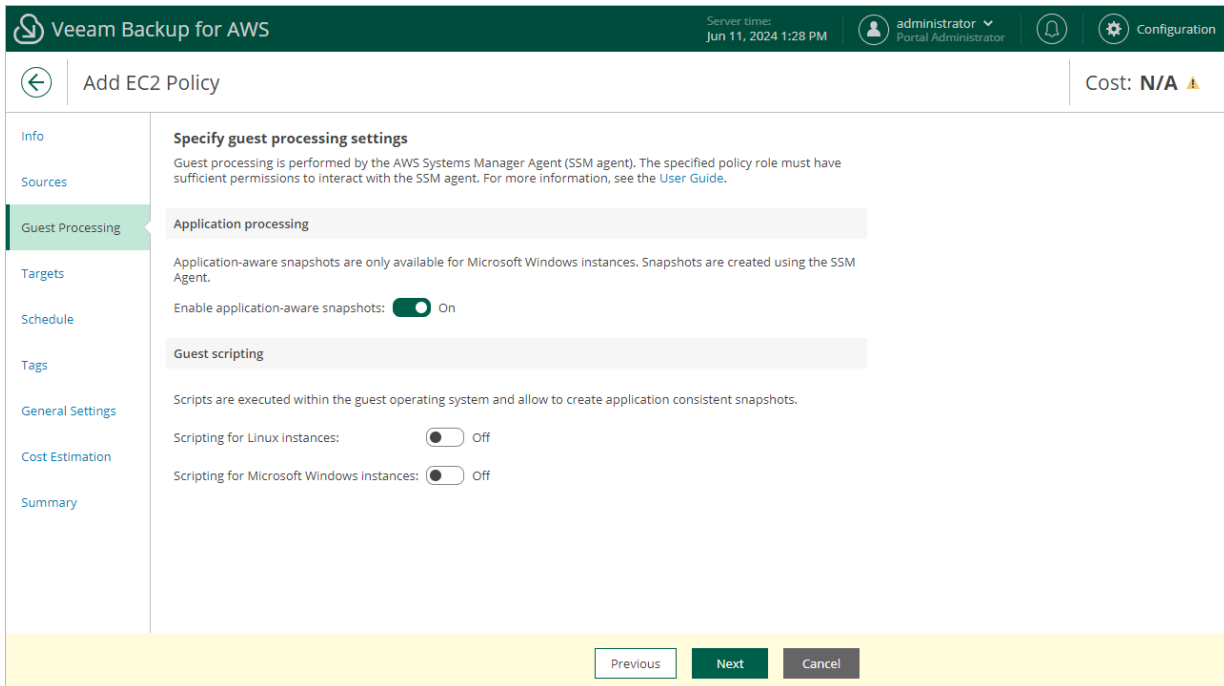
Limitations and Requirements for Application-Aware Processing

If you plan to instruct Veeam Backup for AWS to create transactionally consistent backups using application-aware processing, in addition to the [limitations and requirements for guest processing](#), consider the following:

- Application-aware processing is available only for EC2 instances running Microsoft Windows Server 2008 R2 or later.
- EC2 instances for which you plan to enable application-aware processing must meet the following prerequisites:
 - The EC2 instances must have VSS components installed. To learn how to download and install VSS components, see [AWS Documentation](#).
 - To allow Veeam Backup for AWS to take VSS-enabled snapshots for the EC2 instances, the following permissions must be granted to the IAM roles attached to the instances:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "ec2:CreateTags",
      "Resource": [
        "arn:aws:ec2:*::snapshot/*"
      ]
    },
    {
      "Effect": "Allow",
      "Action": [
        "ec2:DescribeInstances",
        "ec2:CreateSnapshot",
        "ec2:DescribeSnapshots"
      ],
      "Resource": "*"
    }
  ]
}
```

To learn how to create IAM roles for VSS-enabled snapshots and grant permissions to them, see [AWS Documentation](#).



Enabling Guest Scripting

Before you enable guest scripting for processed EC2 instances, check [limitations and requirements](#).

To enable guest scripting, at the **Guest Processing** step of the wizard, do the following

- For EC2 instances running Linux OS, set the **Scripting for Linux instances** toggle to *On*.
The **Specify scripting settings for Linux instances** window will open.
- For EC2 instances running Microsoft Windows OS, set the **Scripting for Microsoft Windows instances** toggle to *On*.
The **Specify scripting settings for Microsoft Windows instances** window will open.

In the opened window, specify pre-snapshot and post-snapshot scripts that must be executed before and after the backup operation:

1. In the **Pre-snapshot script** section, do the following:
 - a. In the **Path in guest** field, specify a path to the pre-snapshot script file on an EC2 instance.
 - b. In the **Arguments** field, specify additional arguments that must be passed to the script when the script is executed.

You can use runtime variables as arguments for the script. To see the list of available variables, click **Parameters**.

NOTE

Veeam Backup for AWS will run the script from the specified directory for all EC2 instances added to the backup policy. If you want to execute different scripts for different EC2 instances, ensure that script files uploaded to these instances are located under the same path and have the same name.

2. Repeat step 1 for post-snapshot scripts in the **Post-snapshot script** section.

3. In the **Additional options** section, choose whether you want to instruct Veeam Backup for AWS to:
 - Run scripts only while taking snapshot that will be used to create an image-level backup.
 - Proceed with snapshot creation even though scripts are missing on some of the processed instances.
 - Ignore exit codes returned while executing the scripts.
4. To save changes made to the backup policy settings, click **Apply**.

Limitations and Requirements for Guest Scripting

If you plan to instruct Veeam Backup for AWS to run custom scripts on the processed EC2 instances, in addition to the [limitations and requirements for guest processing](#), consider the following:

- Scripts must be created beforehand.
- For EC2 instances running Microsoft Windows OS, Veeam Backup for AWS supports scripts in the EXE, BAT, CMD, WSF, JS, VBS and PS1 file formats.
- For EC2 instances running Linux OS, Veeam Backup for AWS supports scripts in the SH file format.
- IAM instance profiles used to grant permissions for SSM to interact with the processed EC2 instances must be created beforehand and attached to these instances. To learn how to create IAM instance profiles for AWS Systems Manager, see [AWS Documentation](#).

The screenshot displays the 'Add EC2 Policy' configuration interface in the Veeam Backup for AWS console. The top navigation bar shows the server time as 'Jun 11, 2024 1:30 PM' and the user as 'administrator Portal Administrator'. The main content area is split into three panels:

- Specify guest processing settings:** This panel includes a description of guest processing and a toggle for 'Enable application-aware snapshots' which is currently turned 'On'.
- Specify scripting settings for Linux instances:** This panel is the active focus. It contains:
 - Pre-snapshot script:** A text input field for 'Path in guest' containing '/scripts/pre-snapshot.cmd' and an 'Arguments' field containing '%policy% %instance%'.
 - Post-snapshot script:** A text input field for 'Path in guest' containing '/scripts/post-snapshot.cmd' and an 'Arguments' field containing '%instance%'.
 - Additional options:** Three toggle switches:
 - 'Run scripts only for snapshots that will be copied to a repository': On
 - 'Ignore missed guest scripts and continue snapshot creation': Off
 - 'Ignore exit codes of the specified scripts': Off
- Buttons:** 'Apply' and 'Cancel' buttons are located at the bottom of the configuration area.

Step 5. Configure Backup Target Settings

By default, backup policies create only cloud-native snapshots of processed instances. At the **Targets** step of the wizard, you can enable the following additional data protection scenarios:

- [Instruct Veeam Backup for AWS to replicate cloud-native snapshots to other AWS accounts or AWS Regions.](#)
- [Instruct Veeam Backup for AWS to create image-level backups.](#)

Configuring Snapshot Replica Settings

If you want to replicate cloud-native snapshots to other AWS accounts or regions, do the following:

1. In the **Replicas** section of the **Targets** step of the wizard, set the **Replicate snapshots** toggle to *On*.
2. In the **Replication settings** window, configure the following mapping settings for each AWS Region where source instances reside:

- a. Select a source AWS Region from the list and click **Edit Region Mapping**.

- b. In the **Edit Region Mapping** window, specify the following settings:

- i. From the **Target account** drop-down list, select an IAM role whose permissions will be used to copy and store cloud-native snapshots in a target AWS Region.

If you select an IAM role created in another AWS account, the cloud-native snapshots will be copied to a target AWS Region in that AWS account.

- ii. From the **Target region** drop-down list, select a target AWS Region to which Veeam Backup for AWS must copy cloud-native snapshots.

- iii. If you want to encrypt the cloud-native snapshots copied to the target AWS Region, select the **Enable encryption** check box and choose the necessary KMS key from the **Encryption key** drop-down list. For a KMS key to be displayed in the list of available encryption keys, it must be stored in the target AWS Region, and the IAM role specified for the copy operation must have permissions to access the key. For more information on KMS keys, see [AWS Documentation](#).

Then, use the **Key usage** drop-down list to choose whether you want to encrypt snapshots for all volumes or only snapshots of the encrypted volumes.

NOTE

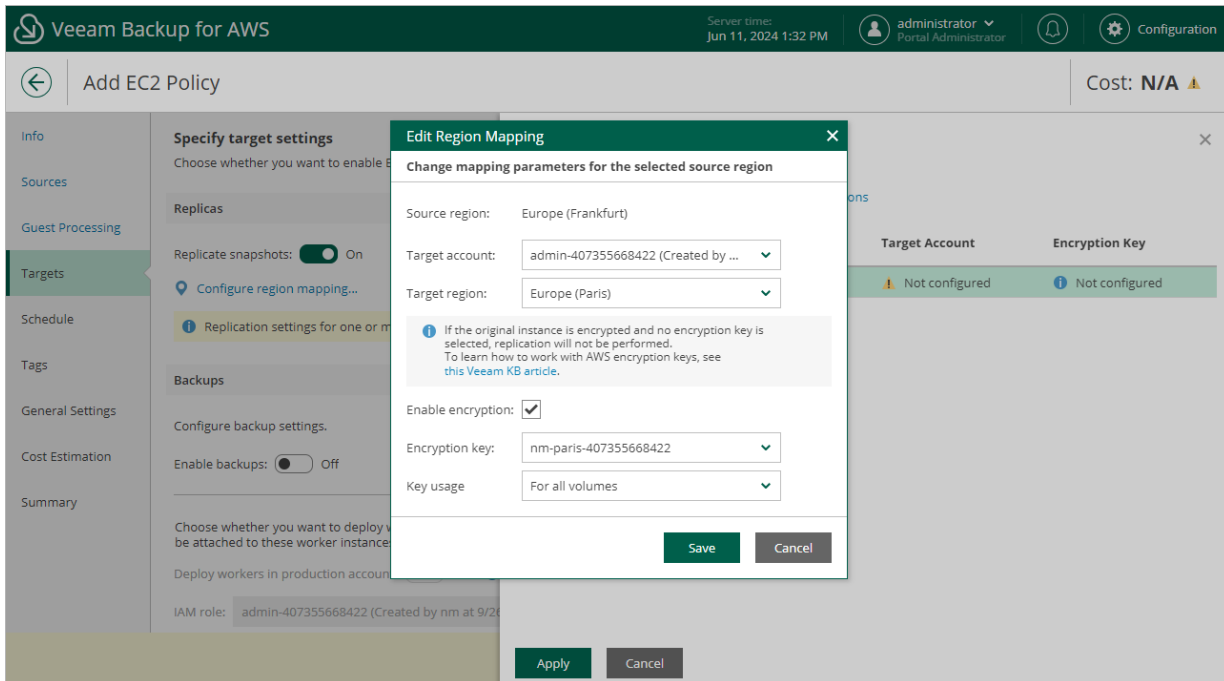
If the original EBS volume is encrypted, you must enable encryption for replicated snapshots, otherwise, the replication process will fail.

- iv. Click **Save**.

TIP

To configure mapping for all source AWS Regions at once, click **Set Mapping for All Regions** and specify settings as described in [step 2.b](#).

c. To save changes made to the backup policy settings, click **Apply**.



Related Resources

[AWS Key Management Service concepts](#)

Configuring Image-Level Backup Settings

In the **Backups** section of the **Targets** step of the wizard, you can instruct Veeam Backup for AWS to create image-level backups of the processed EC2 instances, to copy backups to a long-term archive storage, and to deploy worker instances used for backup operations in the [production account](#).

Configuring Backup Settings

To instruct Veeam Backup for AWS to create image-level backups of the selected EC2 instances, do the following:

1. Set the **Enable backups** toggle to *On*.
2. In the **Repositories** window, select a backup repository where the created image-level backups will be stored, and click **Apply**.

For a backup repository to be displayed in the list of available repositories, it must be added to Veeam Backup for AWS as described in section [Adding Backup Repositories Using Web UI](#). The list shows only backup repositories of the *S3 Standard* storage class.

To learn how Veeam Backup for AWS creates image-level backups, see [EC2 Backup](#).

Configuring Archive Settings

To instruct Veeam Backup for AWS to store backed-up data in a low-cost, long-term archive storage, do the following:

1. Select the **Archives will be stored in** check box.

2. In the **Repositories** window, select a backup repository where the archived data will be stored, and click **Apply**.

For an archive backup repository to be displayed in the list of available repositories, it must be added to Veeam Backup for AWS as described in section [Adding Backup Repositories Using Web UI](#). The list shows only backup repositories of the *S3 Glacier Flexible Retrieval* or *S3 Glacier Deep Archive* storage classes.

For more information on backup archiving, see [Enabling Backup Archiving](#).

IMPORTANT

If you enable the backup archiving, consider that data encryption must be either enabled or disabled for both backup and archive backup repositories. This means that, for example, you cannot select an encrypted standard backup repository and an unencrypted archive backup repository in one backup policy. However, the selected repositories can have different encryption schemes (password and KMS encryption).

Configuring Worker Settings

By default, Veeam Backup for AWS deploys worker instances used to perform backup operations in the [backup account](#). However, you can instruct Veeam Backup for AWS to deploy worker instances in a production account – that is, the same AWS account to which the processed resources belong. To do that, set the **Deploy workers in production account** toggle to *On*, and specify an IAM role that will be attached to the worker instances and used by Veeam Backup for AWS to communicate with these instances. The specified IAM role must belong to the same account to which the IAM role specified to perform the backup operation belongs, and must be assigned the permissions listed in section [Worker Deployment Role Permissions in Production Accounts](#).

For an IAM role to be displayed in the **IAM role** list, it must be added to Veeam Backup for AWS with the *Production worker role* selected as described in section [Adding IAM Roles](#). If you have not added the necessary IAM role to Veeam Backup for AWS beforehand, you can do it without closing the **Add Policy** wizard. To do that, click **Add** and complete the [Add IAM Role](#) wizard.

IMPORTANT

Consider the following:

- If you instruct Veeam Backup for AWS to deploy worker instances in production accounts, you must assign additional permissions to the IAM role used to perform the backup operation. For more information on the required permissions, see [EC2 Backup IAM Role Permissions](#).
- It is recommended that you check whether both the IAM role specified at [step 3.1](#) of the wizard and the IAM role specified in the **Backups** section have the required permissions. If some permissions of the IAM roles are missing, the backup policy will fail to complete successfully. To run the IAM role permission check, click **Check Permissions** and follow the instructions provided in section [Checking IAM Role Permissions](#).
- Veeam Backup for AWS may fail to create image-level backups of EC2 instances with [product codes](#) if the AMIs that were used to launch the instances do not support the type of worker instances deployed for the backup operation. To work around the issue, modify the worker profile to choose another instance type, as described in section [Managing Worker Profiles](#).
- Veeam Backup for AWS does not support backup and restore of EC2 instances with [product codes](#) that have vendor restrictions preventing root EBS volumes from being attached to worker instances as secondary volumes. To learn how Veeam Backup for AWS performs EC2 backup, see [Protecting EC2 Instances](#).

The screenshot displays the 'Add EC2 Policy' configuration page in the Veeam Backup for AWS interface. The page is divided into a left sidebar with navigation options (Info, Sources, Guest Processing, Targets, Schedule, Tags, General Settings, Cost Estimation, Summary) and a main content area. The 'Targets' section is currently selected. The main content area is titled 'Specify target settings' and includes the following sections:

- Replicas:** A section with a toggle for 'Replicate snapshots' set to 'On' and a note indicating 'Mapping for 1 region is configured'.
- Backups:** A section for configuring backup settings, including a toggle for 'Enable backups' set to 'On'. It specifies that backups will be stored in 'v8-rds-metadata-standard' and archives will be stored in 'v8-rds-metadata-glacier'. An information icon indicates that it is recommended to use the S3 Glacier Flexible Retrieval or S3 Glacier Deep Archive storage class for long-term backups.
- Worker Deployment:** A section with a toggle for 'Deploy workers in production account' set to 'On'. It includes a dropdown menu for the 'IAM role' (currently 'admin-407355668422 (Created by nm at 9/26/2023...')') and buttons for '+ Add' and 'Check Permissions'.

At the bottom of the page, there are three buttons: 'Previous', 'Next', and 'Cancel'. The top right corner shows the server time as 'Jun 11, 2024 1:33 PM' and the user as 'administrator Portal Administrator'. The cost is listed as 'N/A'.

Step 6. Specify Policy Scheduling Options

You can instruct Veeam Backup for AWS to start the backup policy automatically according to a specific backup schedule. The backup schedule defines how often data of the instances added to the backup policy will be backed up.

IMPORTANT

If you have selected a standard or an archive backup repository with immutability settings enabled at [step 5](#) of the wizard, you must configure at least one schedule for the backup policy.

To help you implement a comprehensive backup strategy, Veeam Backup for AWS allows you to create schedules of the following types:

- **Daily** – the backup policy will create restore points repeatedly throughout a day on specific days.
- **Weekly** – the backup policy will create restore points once a day on specific days.
- **Monthly** – the backup policy will create restore points once a month on a specific day.
- **Yearly** – the backup policy will create restore points once a year on a specific day.

Combining multiple schedule types together allows you to retain restore points for longer periods of time – for more information, see [Enabling Harmonized Scheduling](#). Combining multiple schedule types together also allows you to archive backups – for more information, see [Enabling Backup Archiving](#).

NOTE

If you do not specify a backup schedule for the backup policy, you will need to start it manually to create EC2 instance snapshots and backups. For information on how to start backup policies, see [Starting and Stopping Policies](#).

Specifying Daily Schedule

To create a daily schedule for the backup policy, at the **Schedule** step of the wizard, do the following:

1. Set the **Daily schedule** toggle to *On* and click **Edit Daily Settings**.
2. In the **Create daily schedule** window, select hours when the backup policy will create cloud-native snapshots, snapshot replicas or image-level backups.

If you want to protect EC2 instance data more frequently, you can instruct the backup policy to create multiple cloud-native snapshots per hour. To do that, click the link to the right of the **Snapshots** hour selection area, and specify the number of cloud-native snapshots that the backup policy will create within an hour.

NOTE

Veeam Backup for AWS does not create snapshot replicas and image-level backups independently from cloud-native snapshots. That is why when you select hours for snapshot replicas and image-level backups, the same hours are automatically selected for cloud-native snapshots. To learn how Veeam Backup for AWS performs backup, see [EC2 Backup](#).

3. Use the **Run at** drop-down list to choose whether you want the backup policy to run everyday, on work days (Monday through Friday) or on specific days.

4. In the **Daily retention** section, configure retention policy settings for the daily schedule:
 - For cloud-native snapshots and snapshot replicas, specify the number of restore points that you want to keep in cloud-native snapshot and snapshot replica chains.

If the restore point limit is exceeded, Veeam Backup for AWS removes the earliest restore point from the chain. For more information, see [EC2 Snapshot Retention](#).

IMPORTANT

To allow the Changed Block Tracking (CBT) mechanism to be used when processing EC2 instance data, you must keep at least one snapshot in the snapshot chain. However, by design, Veeam Backup for AWS permanently retains 2 cloud-native snapshots in the chain due to the CBT mechanism limitations. To learn how the CBT mechanism works, see [Changed Block Tracking](#).

- For image-level backups, specify the number of days (or months) for which you want to keep restore points in a backup chain.

If a restore point is older than the specified time limit, Veeam Backup for AWS removes the restore point from the chain. For more information, see [EC2 Backup Retention](#).

5. To save changes made to the backup policy settings, click **Apply**.

The screenshot shows the 'Add EC2 Policy' wizard in Veeam Backup for AWS. The 'Create daily schedule' window is active, displaying a 24-hour grid for scheduling snapshots, replicas, and backups. The 'Daily retention' section is also visible, with settings for 'Snapshots to keep: 24', 'Replicas to keep: 2', and 'Keep backups for: 14 Days'. The 'Apply' button is highlighted.

Specifying Weekly Schedule

To create a weekly schedule for the backup policy, at the **Schedule** step of the wizard, do the following:

1. Set the **Weekly schedule** toggle to *On* and click **Edit Weekly Settings**.
2. In the **Create weekly schedule** window, select weekdays when the backup policy will create cloud-native snapshots, snapshot replicas or image-level backups.

NOTE

Veeam Backup for AWS does not create snapshot replicas and image-level backups independently from cloud-native snapshots. That is why when you select days to create snapshot replicas and image-level backups, the same days are automatically selected for cloud-native snapshots. To learn how Veeam Backup for AWS performs backup, see [EC2 Backup](#).

3. Use the **Create restore point at** drop-down list to schedule a specific time for the backup policy to run.
4. In the **Weekly retention** section, configure retention policy settings for the weekly schedule:
 - For cloud-native snapshots and snapshot replicas, specify the number of restore points that you want to keep in cloud-native snapshot and snapshot replica chains.

If the restore point limit is exceeded, Veeam Backup for AWS removes the earliest restore point from the chain. For more information, see [EC2 Snapshot Retention](#).

IMPORTANT

To allow the Changed Block Tracking (CBT) mechanism to be used when processing EC2 instance data, you must keep at least one snapshot in the snapshot chain. However, by design, Veeam Backup for AWS permanently retains 2 cloud-native snapshots in the chain due to the CBT mechanism limitations. To learn how the CBT mechanism works, see [Changed Block Tracking](#).

- For image-level backups, specify the number of days (or months) for which you want to keep restore points in a backup chain.

If a restore point is older than the specified time limit, Veeam Backup for AWS removes the restore point from the chain. For more information, see [EC2 Backup Retention](#).

5. To save changes made to the backup policy settings, click **Apply**.

The screenshot displays the Veeam Backup for AWS configuration interface. At the top, the server time is Jun 11, 2024 1:35 PM, and the user is administrator (Portal Administrator). The main heading is "Add EC2 Policy" with a cost of \$12.14. The interface is divided into several sections:

- Specify scheduling options:** Includes "Daily schedule" (On) and "Weekly schedule" (On). It lists settings for Snapshots (24 per day), Replicas (2 per day), Backups (1 per day), and Repository (v8-rds-metadata-standard).
- Create weekly schedule:** A dialog box for specifying the frequency of snapshots, replicas, and backups. It features a calendar grid with columns for Sun through Sat. Snapshots are selected for Sun, Mon, Tue, Wed, Thu, and Fri (Total: 5). Replicas are selected for Mon, Tue, and Thu (Total: 2). Backups are selected for Fri (Total: 1). The "Creation" toggle is set to "On".
- Weekly retention:** A section for configuring retention policies. It includes a note: "Due to a higher cost, snapshots and replicas are best used for short-term retention. For long-term retention, leverage more cost-effective backups." The settings are: Snapshots to keep: 5; Replicas to keep: 4; Keep backups for: 1 Months.

At the bottom of the dialog, there are "Apply" and "Cancel" buttons.

Specifying Monthly Schedule

To create a monthly schedule for the backup policy, at the **Schedule** step of the wizard, do the following:

1. Set the **Monthly schedule** toggle to *On* and click **Edit Monthly Settings**.
2. [This step applies if you enabled backup archiving at the [Targets](#) step of the wizard] In the **Choose monthly backup target** section of the opened window, choose whether you want to store monthly backups in the archive backup repository.

If you set the **Send backups to archive** toggle to *On*, follow the instructions provided in section [Enabling Backup Archiving](#).

3. In the **Create monthly schedule** section, select months when the backup policy will create cloud-native snapshots, snapshot replicas or image-level backups.

NOTE

Veeam Backup for AWS does not create snapshot replicas and image-level backups independently from cloud-native snapshots. That is why when you select months to create snapshot replicas and image-level backups, the same months are automatically selected for cloud-native snapshots. To learn how Veeam Backup for AWS performs backup, see [EC2 Backup](#).

4. Use the **Create restore point at** and **Run on** drop-down lists to schedule a specific time and day for the backup policy to run.

NOTE

Consider the following:

- If you have selected a specific time for the backup policy to run at the **Weekly schedule** section of the **Schedule** step of the wizard, you will not be able to change the time for the monthly schedule unless you select the *On Day* option from the **Run on** drop-down list.
- If you select the **On day** option, [harmonized scheduling](#) cannot be guaranteed. Plus, to support the **On day** option, Veeam Backup for AWS will require to create an additional temporary restore point if there are no other schedules planned to run on that day. However, the temporary restore point will be removed during the *Backup Retention* process from AWS in approximately 24 hours, to reduce unexpected infrastructure charges.

5. In the **Monthly retention** section, configure retention policy settings for the monthly schedule:
 - For cloud-native snapshots and snapshot replicas, specify the number of restore points that you want to keep in cloud-native snapshot and snapshot replica chains.

If the restore point limit is exceeded, Veeam Backup for AWS removes the earliest restore point from each chain. For more information, see [EC2 Snapshot Retention](#).

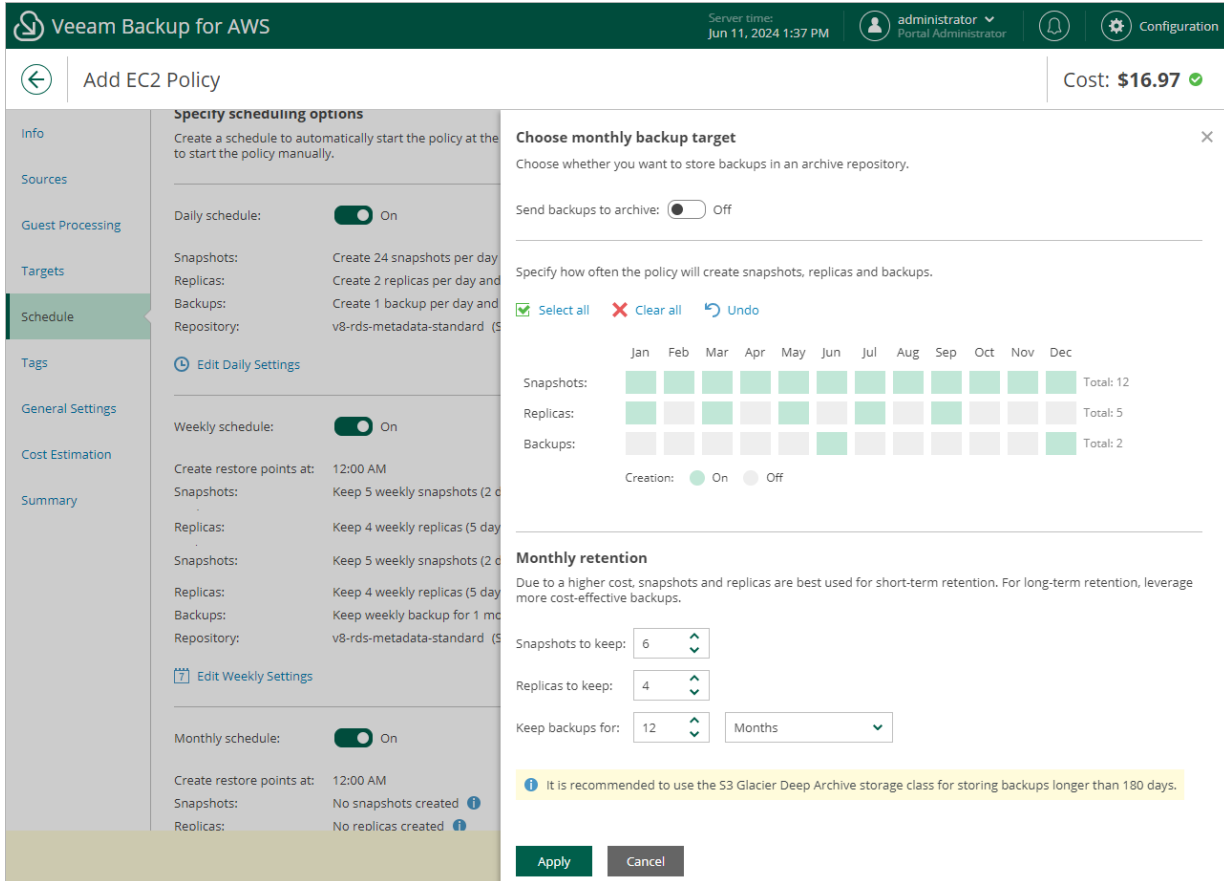
IMPORTANT

To allow the Changed Block Tracking (CBT) mechanism to be used when processing EC2 instance data, you must keep at least one snapshot in the snapshot chain. However, by design, Veeam Backup for AWS permanently retains 2 cloud-native snapshots in the chain due to the CBT mechanism limitations. To learn how the CBT mechanism works, see [Changed Block Tracking](#).

- For image-level backups, specify the number of days (or months) for which you want to keep restore points in a backup chain.

If a restore point is older than the specified time limit, Veeam Backup for AWS removes the restore point from the chain. For more information, see [EC2 Backup Retention](#).

6. To save changes made to the backup policy settings, click **Apply**.



Specifying Yearly Schedule

[This step applies only if you have instructed Veeam Backup for AWS to create image-level backups at the **Targets** step of the wizard]

To create a yearly schedule for the backup policy, at the **Schedule** step of the wizard, do the following:

1. Set the **Yearly schedule** toggle to *On* and click **Edit Yearly Settings**.
2. [This step applies if you enabled backup archiving at the **Targets** step of the wizard] In the **Choose yearly backup target** section of the opened window, choose whether you want to store yearly backups in the archive backup repository.

If you set the **Send backups to archive** toggle to *On*, follow the instructions provided in section [Enabling Backup Archiving](#).

3. In the **Yearly schedule** section, specify a day, month and time when the backup policy will create image-level backups.

For example, if you select *First, Friday, January* and *06:00 PM*, the backup policy will run every first Friday of January at 06:00 PM.

NOTE

Consider the following:

- If you have selected a specific time and day for the backup policy to run at the **Weekly schedule** or **Monthly schedule** sections of the **Schedule** step of the wizard, you will not be able to change the time and day for the yearly schedule unless you select the *On Day* option from the **Create restore point on** drop-down list.
- If you select the *On day* option, **harmonized scheduling** cannot be guaranteed. Plus, to support the *On day* option, Veeam Backup for AWS will require to create an additional temporary restore point if there are no other schedules planned to run on that day. However, the temporary restore point will be removed during the *Backup Retention* process from AWS in approximately 24 hours, to reduce unexpected infrastructure charges.

4. In the **Keep backups for** field, specify the number of years for which you want to keep restore points in a backup chain.

If a restore point is older than the specified time limit, Veeam Backup for AWS removes the restore from the chain. For more information, see [EC2 Backup Retention](#).

5. To save changes made to the backup policy settings, click **Apply**.

The screenshot shows the Veeam Backup for AWS console interface. The main window is titled 'Add EC2 Policy' and shows a sidebar with navigation options like Info, Sources, Guest Processing, Targets, Schedule (highlighted), Tags, General Settings, Cost Estimation, and Summary. The 'Schedule' section is expanded, showing settings for 'Weekly schedule' and 'Monthly schedule'. Both are currently turned 'On'. The 'Weekly schedule' section shows 'Create restore points at: 12:00 AM', 'Snapshots: Keep 5 weekly snapshots (2 d)', 'Replicas: Keep 4 weekly replicas (5 day)', and 'Backups: Keep weekly backup for 1 mo'. The 'Monthly schedule' section shows 'Create restore points at: 12:00 AM', 'Snapshots: Keep 6 monthly snapshots', 'Replicas: Keep 4 monthly replicas (7 m)', and 'Backups: Keep monthly backups for 12 m'. A 'Create yearly schedule' dialog box is open in the foreground, with the following settings: 'Send backups to archive: On', 'Yearly schedule is applied only to image-level backups. Specify for how many years the policy will keep backup files.', 'Create restore points on: First Monday of June at 12:00 AM', and 'Keep archives for: 2 years'. The dialog has 'Apply' and 'Cancel' buttons.

Enabling Harmonized Scheduling

When you combine multiple types of schedules, Veeam Backup for AWS applies the harmonization mechanism that allows you to leverage restore points for long-term retentions instead of taking a new restore point every time. The mechanism simplifies the backup schedule, optimizes the backup performance and reduces the cost of retaining restore points.

With harmonized scheduling, Veeam Backup for AWS can keep restore points created according to a daily, weekly or monthly schedule for longer periods of time:

- Cloud-native snapshots and snapshot replicas can be kept for weeks and months.

- Image-level backups can be kept for weeks, months and years.

For Veeam Backup for AWS to use the harmonization mechanism, there must be specified at least 2 different schedules: one schedule will control the regular creation of restore points, while another schedule will control the process of storing restore points. In terms of harmonized scheduling, Veeam Backup for AWS re-uses restore points created according to a more-frequent schedule (daily, weekly or monthly) to achieve the desired retention for less-frequent schedules (weekly, monthly and yearly). Each restore point is marked with a flag of the related schedule type: the (D) flag is used to mark restore points created according to the daily schedule, (W) – weekly, (M) – monthly, and (Y) – yearly. Veeam Backup for AWS uses these flags to control the retention period for the created restore points. Once a flag of a less-frequent schedule is assigned to a restore point, this restore point can no longer be removed – it is kept for the period defined in the retention settings. When the specified retention period is over, the flag is unassigned from the restore point. If the restore point does not have any other flags assigned, it is removed according to the retention settings of a more-frequent schedule.

NOTE

Restore points created according to a more-frequent schedule and less-frequent schedules compose a single backup or snapshot chain. This means that regardless of flags assigned to restore points, Veeam Backup for AWS adds the restore points to the chain as described in sections [Backup Chain](#) and [Snapshot Chain](#).

Consider the following example. You want a backup policy to create cloud-native snapshots of your critical workloads 3 times a day, to keep 3 daily snapshots in the snapshot chain, and also to keep one of the created snapshots for 2 weeks. In this case, you create 2 schedules when configuring the backup policy settings – daily and weekly:

- In the daily scheduling settings, you select hours and days when snapshots will be created (for example, *7:00 AM, 9:00 AM, and 11:00 AM; Weekdays*), and specify a number of daily restore points to retain (for example, *3*).

Veeam Backup for AWS will propagate these settings to the schedule of a lower frequency (which is the weekly schedule in our example).


The screenshot displays the 'Add EC2 Policy' configuration interface in Veeam Backup for AWS. The top navigation bar includes the Veeam logo, 'Add EC2 Policy', and a cost indicator of '\$2.57'. The left sidebar contains navigation options: Info, Sources, Guest Processing, Targets, Schedule (selected), Tags, General Settings, Cost Estimation, and Summary. The main content area is divided into three sections:

- Specify scheduling options:** Includes toggle switches for 'Daily schedule' (turned on), 'Weekly schedule', 'Monthly schedule', and 'Yearly schedule'. A 'Snapshots: Create' button is also present.
- Create daily schedule:** A section for specifying snapshot frequency. It features a 24-hour clock with AM and PM labels. Three snapshots are selected at 7:00 AM, 9:00 AM, and 11:00 AM. Below the clock, there are 'Snapshots:' and 'Creation: On/Off' options. The 'Run at:' dropdown is set to 'Weekdays'.
- Daily retention:** A section explaining that due to higher costs, snapshots and replicas are best used for short-term retention. It includes a 'Snapshots to keep:' input field set to the value '3'.

At the bottom of the configuration window, there are 'Apply' and 'Cancel' buttons.

- In the weekly scheduling settings, you specify which one of the snapshots created by the daily schedule will be retained for a longer period, and choose for how long you want to keep the selected snapshot.

For example, if you want to keep the daily restore point created at 7:00 AM on Monday for 2 weeks, you select *7:00 AM, Monday* and specify 2 restore points to retain in the weekly schedule settings.




According to the specified scheduling settings, Veeam Backup for AWS will create cloud-native snapshots in the following way:

- On the first work day (Monday), a backup session will start at 7:00 AM to create the first restore point. The restore point will be marked with the (D) flag as it was created according to the daily schedule.


Since *7:00 AM, Monday* is specified in weekly schedule settings, Veeam Backup for AWS will also assign the (W) flag to this restore point. As a result, 2 flags (D,W) will be assigned to the restore point.

- On the same day (Monday), after backup sessions run at 9:00 AM and 11:00 AM, the created restore points will be marked with the (D) flag.

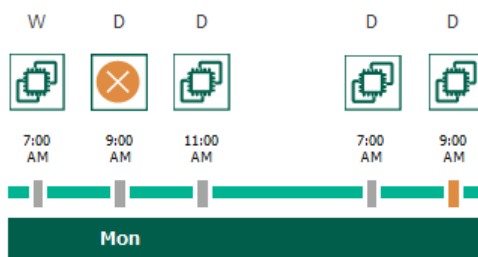


- On the next work day (Tuesday), after a backup session runs at 7:00 AM, the created restore point will be marked with the (D) flag.


By the moment the backup session completes, the number of restore points with the (D) flag will exceed the retention limit specified in the daily scheduling settings. However, Veeam Backup for AWS will not remove the earliest restore point (7:00 AM, Monday) with the (D) flag from the snapshot chain as this restore point is also marked with a flag of a less-frequent schedule. Instead, Veeam Backup for AWS will unassign the (D) flag from the restore point. This restore point will be kept for the retention period specified in the weekly scheduling settings (that is, for 2 weeks).



- On the same day (Tuesday), after a backup session runs at 9:00 AM, the number of restore points with the (D) flag will exceed the retention limit once again. Veeam Backup for AWS will remove from the snapshot chain the restore point created at 9:00 AM on Monday as no flags of a less-frequent schedule are assigned to this restore point.



- Veeam Backup for AWS will continue creating restore points for the next week in the same way as described in steps 1–4.
- On week 3, after a backup session runs at 7:00 AM on Monday, the number of weekly restore points will exceed the retention limit. Veeam Backup for AWS will unassign the (W) flag from the earliest weekly restore point. Since no other flags are assigned to this restore point, Veeam Backup for AWS will remove this restore point from the snapshot chain.



Enabling Backup Archiving

When you combine multiple types of schedules, you can enable the archiving mechanism to instruct Veeam Backup for AWS to store backed-up data in the secure, low-cost and long-term S3 Glacier Flexible Retrieval and S3 Glacier Deep Archive storage classes. The mechanism is the most useful in the following cases:

- Your data retention policy requires that you keep rarely accessed data in an archive.

- You want to reduce data-at-rest costs and to save space in the high-cost, short-term S3 standard storage class.

NOTE

Restoring from an archived backup is longer and more expensive than restoring from a regular backup as it is required to retrieve data from the archive repository. For more information, see [Retrieving EC2 Data From Archive](#).

With backup archiving, Veeam Backup for AWS can retain backup files created according to a daily, weekly or monthly schedule for longer periods of time:

- To enable monthly archiving, you must configure a daily or a weekly schedule (or both).
- To enable yearly archiving, you must configure a daily, a weekly or a monthly schedule (or all three).

For Veeam Backup for AWS to use the archiving mechanism, there must be specified at least 2 different schedules: one schedule will control the regular creation of backup files, while another schedule will control the process of copying backup files to an archive backup repository. Backup chains created according to these two schedules will be completely different – for more information, see [EC2 Backup Chain](#) and [Archive Backup Chain](#).

Consider the following example. You want a backup policy to create image-level backups of your critical workloads once a week, to keep the backed-up data in a standard backup repository for 3 weeks, and also to keep backups created once in 2 months in an archive backup repository for a year. In this case, you create 2 schedules when configuring the backup policy settings – weekly and monthly:

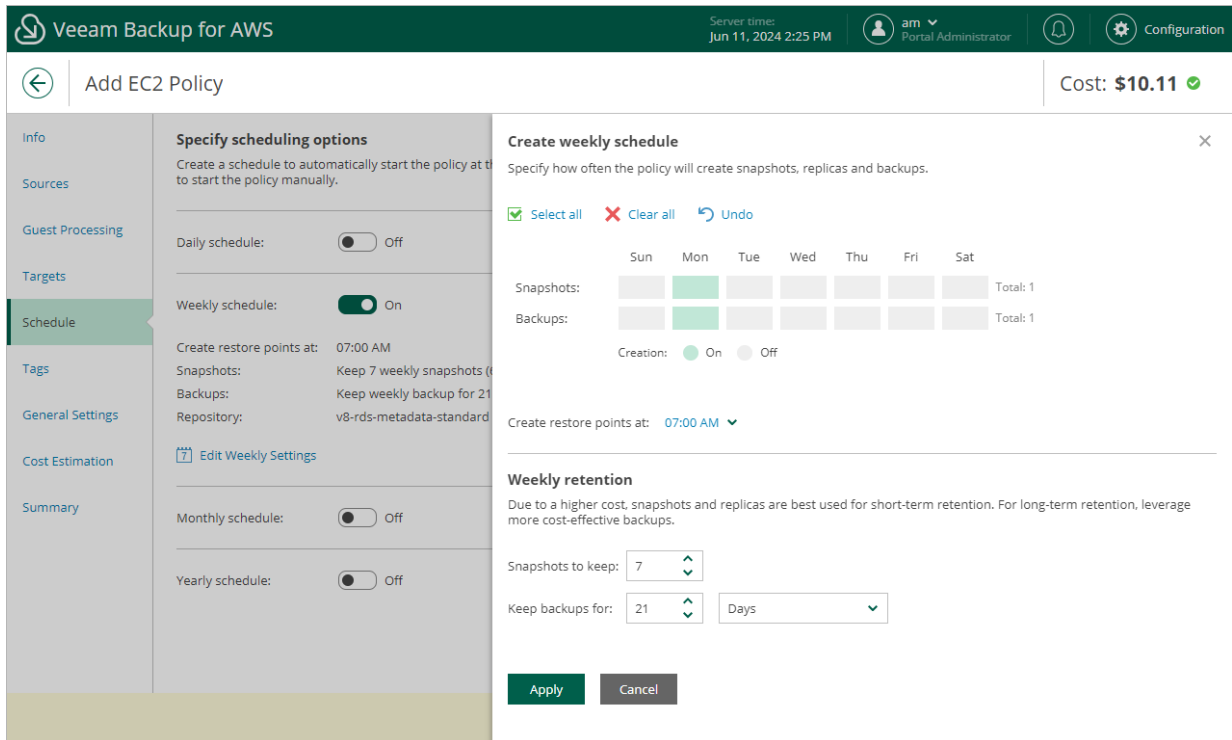
- In the policy target settings, you set the **Enable backups** toggle to *On*, select a backup repository that will store standard backup files, and select an archive backup repository that will store archived data.

The screenshot shows the Veeam Backup for AWS console interface. The main window is titled 'Add EC2 Policy' and shows various configuration tabs like Info, Sources, Guest Processing, Targets, Schedule, Tags, General Settings, Cost Estimation, and Summary. The 'Targets' tab is active, and the 'Backups' section is expanded, showing 'Enable backups' set to 'On'. A modal window titled 'Repositories' is open, displaying a table of available repositories. The 'v8-rds-meta...' repository is highlighted.

Repository	Region	Storage Class	Folder	Description	Immuta...
rds-metadat...	Europe (London)	S3 Glacier Flex...	rds-meta...	Created by...	Disabled
v6-for-ami-i...	Europe (London)	S3 Glacier Flex...	v6-for-a...	Created by...	Disabled
v7-glacier-m...	Europe (London)	S3 Glacier Flex...	v7-glacie...	Created by...	Disabled
v8-glacier-im...	Europe (London)	S3 Glacier Flex...	v8-glacie...	Created by...	Enabled
v8-metadata...	Europe (London)	S3 Glacier Flex...	v8-metad...	Created by...	Disabled
v8-rds-meta...	Europe (London)	S3 Glacier Flex...	v8-rds-m...	Created by...	Disabled

- In the weekly scheduling settings, you select hours and days when backups will be created (for example, *7:00 AM, Monday*), and specify a number of days for which Veeam Backup for AWS will retain backups (for example, *21 days*).

Veeam Backup for AWS will propagate these settings to the archive schedule (which is the monthly schedule in our example).



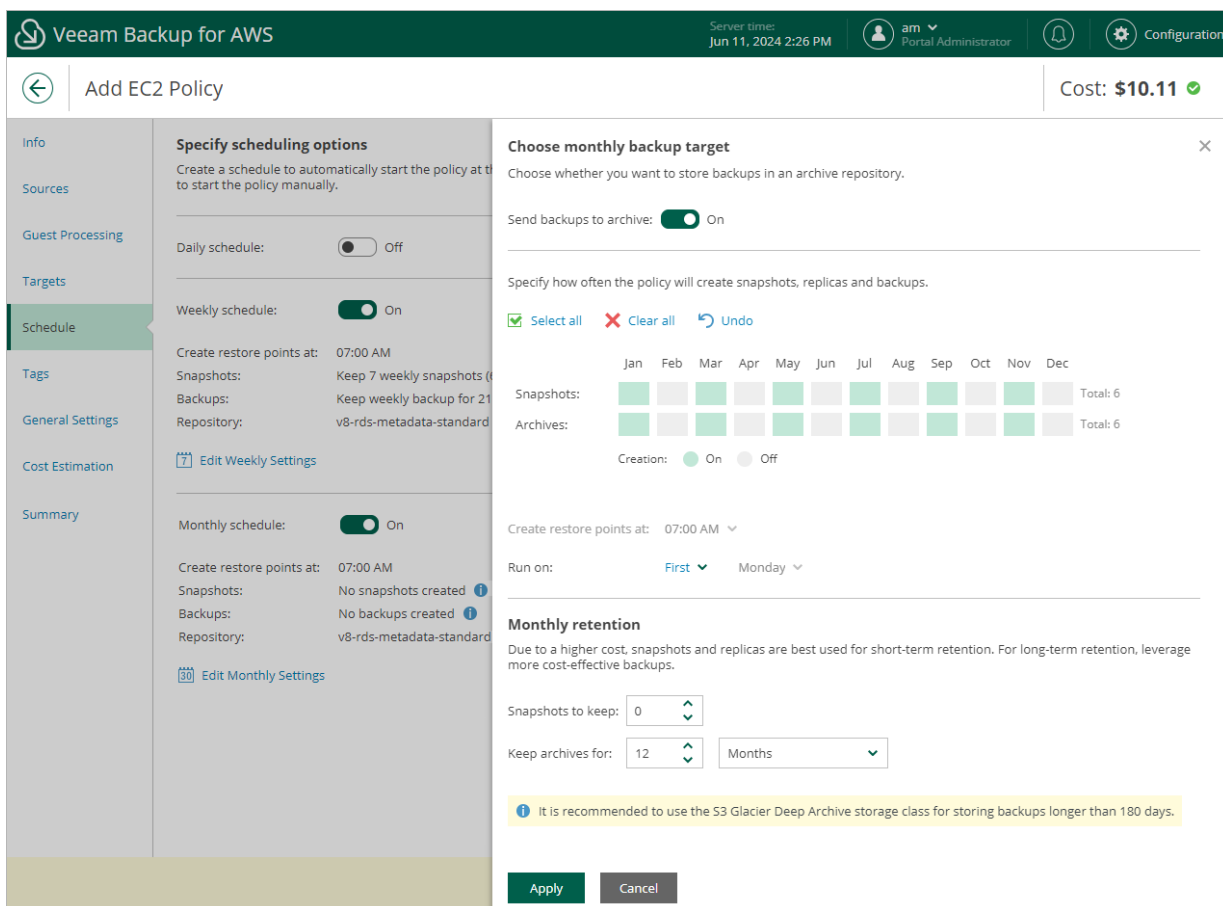
- In the monthly scheduling settings, you enable the archiving mechanism by setting the **Send backups to archive** toggle to *On*, specify when Veeam Backup for AWS will create archive backup files, and choose for how long you want to keep the created backups in the archive backup repository.

For example, *January, March, May, July, September, November, 12 months and First Monday*.

IMPORTANT


Consider the following:

- When you enable backup archiving, you become no longer able to create a schedule of the same frequency for standard backups. By design, these two functionalities are mutually exclusive.
- If you enable backup archiving, it is recommended that you set the **Snapshots to keep** value to *0*, to reduce unexpected snapshot charges.
- If you enable backup archiving, it is recommended that you set the **Keep archives for** value to at least *3 months* (or *90 days*) for the S3 Glacier Flexible Retrieval storage class and at least *6 months* (or *180 days*) for the S3 Glacier Deep Archive storage class. For more information on the minimum storage duration of the Amazon S3 archival storage classes, see [AWS Documentation](#).
- If you select the **On day** option, [harmonized scheduling](#) cannot be guaranteed. Plus, to support the **On day** option, Veeam Backup for AWS will require to create an additional temporary restore point if there are no other schedules planned to run on that day. However, the temporary restore point will be removed during the *Backup Retention* process from AWS in approximately 24 hours, to reduce unexpected infrastructure charges.



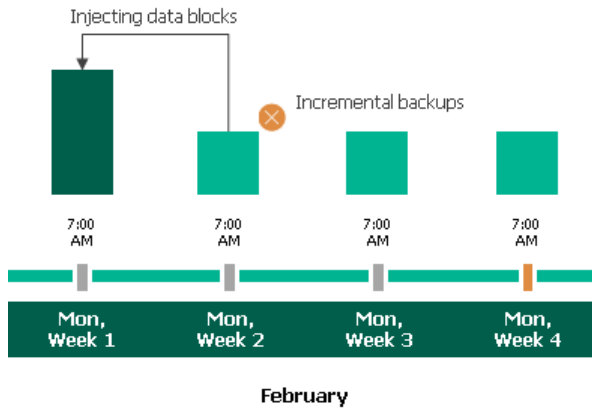
According to the specified scheduling settings, Veeam Backup for AWS will create image-level backups in the following way:

1. On the first Monday of February, a backup session will start at 7:00 AM to create the first restore point in the standard backup chain. Veeam Backup for AWS will store this restore point as a full backup file in the backup repository.
2. On the second and third Mondays of February, Veeam Backup for AWS will create restore points at 7:00 AM and add them to the standard backup chain as incremental backup files in the backup repository.



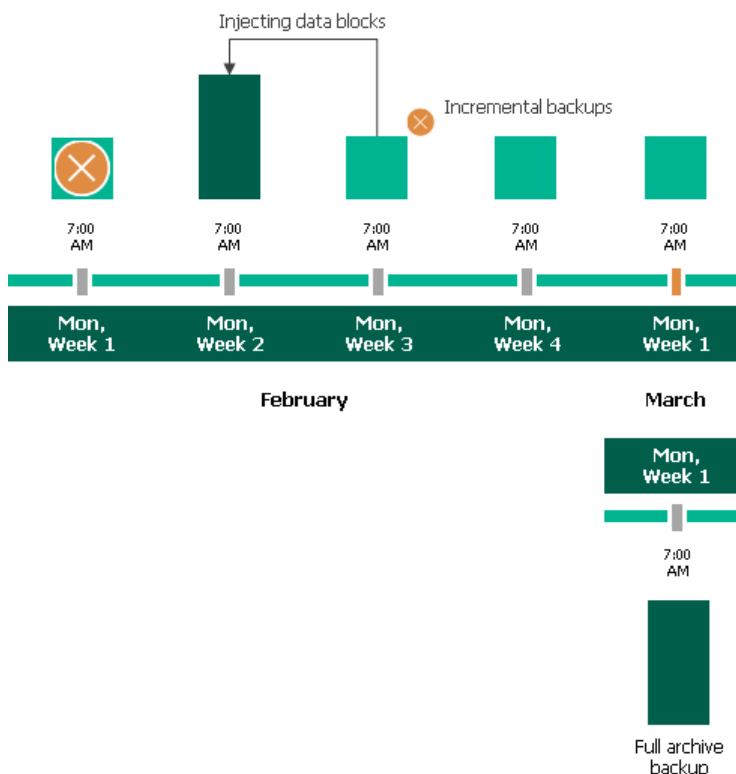
- On the fourth Monday of February, Veeam Backup for AWS will create a new restore point at 7:00 AM. By the moment the ,the earliest restore point in the standard backup chain will get older than the specified retention limit. That is why Veeam Backup for AWS will rebuild the full backup file and remove from the chain the restore point created on the first Monday.

For more information on how Veeam Backup for AWS transforms standard backup chains, see [EC2 Backup Retention](#).




- On the first Monday of March, a backup session will start at 7:00 AM to create another restore point in the standard backup chain. At the same time, the earliest restore point in the standard backup chain will get older than the specified retention limit again. That is why Veeam Backup for AWS will rebuild the full backup file again and remove from the chain the restore point created on the second Monday.

After the backup session completes, an archive session will create a restore point with all data from the standard backup chain. Veeam Backup for AWS will copy this restore point as a full archive backup file to the archive backup repository.



- Up to May, Veeam Backup for AWS will continue adding new restore points to the standard backup chain and deleting outdated backup files from the backup repository, according to the specified weekly scheduling settings.

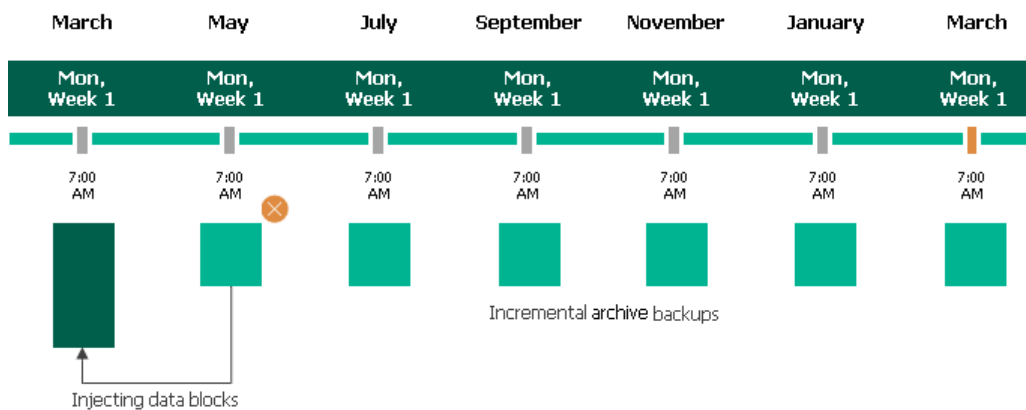
On the first Monday of May, an archive session will create a restore point with only that data that has changed since the previous archive session in March. Veeam Backup for AWS will copy this restore point as an incremental archive backup file to the archive backup repository.



- Up to the first Monday of March of the next year, Veeam Backup for AWS will continue adding new restore points to the standard backup chain and deleting outdated backup files from the backup repository, according to the specified weekly scheduling settings. Veeam Backup for AWS will also continue adding new restore points to the archive backup chain, according to the specified monthly settings.

By the moment the archive session completes, the earliest restore point in the archive backup chain will get older than the specified retention limit. That is why Veeam Backup for AWS will rebuild the full archive backup file and remove from the chain the restore point created on the first Monday of March of the previous year.

For more information on how Veeam Backup for AWS transforms archive backup chains, see [Retention Policy for Archived Backups](#).



Step 7. Enable AWS Tags Assigning

At the **Tags** step of the wizard, you can instruct Veeam Backup for AWS to assign AWS tags to snapshots and snapshots replicas:

1. To assign already existing AWS tags from the EBS volumes of the processed EC2 instance, select the **Copy tags from source volumes** check box.

If you choose to copy tags from the source volumes, Veeam Backup for AWS will first create a cloud-native snapshot or snapshot replica of the EC2 instance and will assign to the created snapshot AWS tags with Veeam metadata, then Veeam Backup for AWS will copy tags from the volumes of the processed instance and, finally, assign the copied AWS tags to the snapshot.

2. To assign your own custom AWS tags, set the **Add custom tags to created snapshots** toggle to *On* and specify the AWS tags explicitly. To do that, use the **Key** and **Value** fields to specify a key and value for the new custom AWS tag, and then click **Add**. Note that you cannot add more than 5 custom AWS tags.

If you choose to add custom tags to the created snapshots, Veeam Backup for AWS will assign the specified tags right after it creates a cloud-native snapshot or snapshot replica.

The screenshot shows the 'Add EC2 Policy' wizard in Veeam Backup for AWS, specifically the 'Tags' step. The interface includes a top navigation bar with the Veeam logo, server time (Jun 11, 2024 1:40 PM), user information (administrator, Portal Administrator), and a Configuration icon. The main content area is titled 'Add EC2 Policy' and shows a cost of '\$42.58'. The 'Tags' step is active, displaying 'Specify tag settings'. The 'Copy tags from source volumes' checkbox is checked. The 'Add custom tags to created snapshots' toggle is set to 'On'. There are two custom tags defined: 'location' with value 'frankfurt' and 'dept: accounting'. A note states 'A maximum of 5 custom tags is allowed.' The bottom of the screen features 'Previous', 'Next', and 'Cancel' buttons.

Step 8. Specify General Settings

At the **GeneralSettings** step of the wizard, you can enable automatic retries, schedule health checks and specify notification settings for the backup policy.

Automatic Retry Settings

To instruct Veeam Backup for AWS to run the backup policy again if it fails on the first try, do the following:

1. In the **Schedule** section of the step, select the **Automatically retry failed policy** check box.
2. In the field to the right of the check box, specify the maximum number of attempts to run the backup policy. The time interval between retries is 60 seconds.

When retrying backup policies, Veeam Backup for AWS processes only those instances that failed to be backed up during the previous attempt.

Health Check Settings

If you have enabled creation of image-level backups at [step 5](#) of the wizard, you can instruct Veeam Backup for AWS to periodically perform a health check for backup restore points created by the policy. During the health check, Veeam Backup for AWS performs an availability check for data blocks in the whole standard backup chain, and a cyclic redundancy check (CRC) for storage metadata to verify its integrity. The health check helps you ensure that the restore points are consistent and that you will be able to restore data using these restore points. For more information on the health check, see [How Health Check Works](#).

NOTE

During a health check, Veeam Backup for AWS does not verify archived restore points created by the policy.

To enable health checks for the backup policy, do the following:

1. In the **Health check** section of the step, set the **Enable health check** toggle to *On*.
2. Use the **Run on** drop-down lists to schedule a specific day for the health check to run.

NOTE

Veeam Backup for AWS performs the health check during the first policy session that runs on the day when the health check is scheduled. If another backup policy session runs on the same day, Veeam Backup for AWS will not perform the health check during that session. For example, if the backup policy is scheduled to run multiple times on Saturday, and the health check is also scheduled to run on Saturday, the health check will only be performed during the first policy session on Saturday.

Email Notification Settings

NOTE

To be able to specify email notification settings for the EC2 Backup policy, you must configure [global notification settings](#) first.

To instruct Veeam Backup for AWS to send email notifications for the backup policy, do the following:

1. In the **Notifications** section of the step, set the **Enabled** toggle to *On*.

If you set the toggle to *Off*, Veeam Backup for AWS will send notifications according to the configured global notification settings.

2. In the **Email** field, specify an email address of a recipient.

Use a semicolon to separate multiple recipient addresses. Do not use spaces after semicolons between the specified email addresses.

3. Use the **Notify on** list to choose whether you want Veeam Backup for AWS to send email notifications in case the backup policy completes successfully, completes with warnings or completes with errors.

4. Select the **Suppress notifications until the last retry** check box to receive a notification about the final backup policy result.

If you do not select the check box, Veeam Backup for AWS will send a notification for every backup policy retry.

NOTE

If you specify the same email recipient in both backup policy notification and [global notification settings](#), Veeam Backup for AWS will override the configured global notification settings and will send each notification to this recipient only once to avoid notification duplicates.

The screenshot displays the Veeam Backup for AWS configuration interface for an EC2 Policy. The top navigation bar includes the Veeam logo, the text "Veeam Backup for AWS", the server time "Jun 11, 2024 1:42 PM", and the user "administrator Portal Administrator". The main content area is titled "Add EC2 Policy" and shows a cost of "\$42.58". The "General Settings" tab is selected in the left sidebar. The "Configure retry and notification settings" section is active, with a sub-header "Specify how many times to retry the policy. You can also enable email notifications to receive policy results." The "Schedule" section shows "Automatically retry failed policy" set to 3 times. A note states "Automatic retry settings are only applicable on a scheduled run of the policy". The "Health check" section has "Enable health check" set to "On". The "Run on" section is set to "First Sunday of every month". The "Notifications" section shows "Enabled" as "On", "Email" as "donna_ortiz@company.mail", and "Notify on" options for "Failure", "Warning", and "Success". The "Suppress notifications until the last retry" checkbox is checked. The bottom bar contains "Previous", "Next", and "Cancel" buttons.

How Health Check Works

When Veeam Backup for AWS saves a new backup restore point to a backup repository, it calculates CRC values for metadata in the backup chain and saves these values to the chain metadata, together with the instance data. When performing a health check, Veeam Backup for AWS verifies the availability of data blocks and uses the saved values to ensure that the restore points being verified are consistent.

If you have enabled health checks for the backup policy, Veeam Backup for AWS performs the following operations at the day scheduled for a health check to run:

1. As soon as a backup policy session completes successfully, Veeam Backup for AWS starts the health check as a new session. For each restore point in the standard backup chain, Veeam Backup for AWS calculates CRC values for backup metadata and compares them to the CRC values that were previously saved to the restore point. Veeam Backup for AWS also checks whether data blocks that are required to rebuild the restore point are available.

If the backup policy session completes with an error, Veeam Backup for AWS tries to run the backup policy again, taking into account the maximum number of retries specified in the [automatic retry settings](#). After the first successful retry (or after the last one out of the maximum number of retries), Veeam Backup for AWS starts the health check.

2. If Veeam Backup for AWS does not detect data inconsistency, the health check session completes successfully. Otherwise, the session completes with an error.

Depending on the detected data inconsistency, Veeam Backup for AWS performs the following operations:

- If the health check detects corrupted metadata in a full or an incremental restore point, Veeam Backup for AWS marks the backup chain as corrupted in the configuration database. During the next backup policy session, Veeam Backup for AWS copies the full instance image, creates a full restore point in the backup repository and starts a new backup chain in the backup repository.

NOTE

Veeam Backup for AWS does not support metadata check for encrypted backup chains.

-
- If the health check detects corrupted disk blocks in a full or an incremental restore point, Veeam Backup for AWS marks the restore point that includes the corrupted data blocks and all subsequent affected incremental restore points as incomplete in the configuration database. During the next backup policy session, Veeam Backup for AWS reads whole data blocks and copies those data blocks that have changed since the previous backup session with corrupted data blocks, and saves these data blocks to the latest restore point that has been created during the current session.

All restore points marked as incomplete will be deleted according to the specified retention policy settings.

Step 9. Review Estimated Cost

[This step applies only if you have created a schedule for the backup policy at the **Schedule** step of the wizard]

At the **Cost Estimation** step of the wizard, review the estimated monthly cost of AWS services and resources that will be consumed to protect the instances added to the backup policy. The total estimated cost includes the following:

- The cost of creating and maintaining cloud-native snapshots of the instances.
For each instance included in the backup policy, Veeam Backup for AWS takes into account the instance type, the number of EBS volumes attached, the number of restore points to be kept in the snapshot chain, and the configured scheduling settings.
- The cost of creating snapshot replicas and maintaining them in the target AWS Region.
For each instance included in the backup policy, Veeam Backup for AWS takes into account the instance type, the number of EBS volumes attached, the number of restore points to be kept in the snapshot chain, and the configured scheduling settings.
- The cost of creating and storing in backup repositories image-level backups of the instances.
For each instance included in the backup policy, Veeam Backup for AWS takes into account the machine type, the number of EBS volumes attached, the number of restore points to be kept in the backup chain, and the configured scheduling settings.
- The cost of creating and storing in archive repositories archived backups of the instances.
For each instance included in the backup policy, Veeam Backup for AWS takes into account the machine type, the number of EBS volumes attached, the number of restore points to be kept in the archive backup chain, and the configured scheduling settings.
- The cost of transferring the instance data between AWS Regions during data protection operations (for example, if a protected instance and the target backup repository reside in different regions).
- The cost of sending API requests to Veeam Backup for AWS during data protection operations.

NOTE

To calculate the estimated cost, Veeam Backup for AWS uses capabilities of the [AWS Pricing Calculator](#) that estimates the cost of services in USD only. This calculator is intended for informational and estimation purposes only.

The estimated cost may occur to be significantly higher due to the backup frequency, cross-region data transfer and snapshot charges. To reduce the cost, you can try the following workarounds:

- To avoid additional costs related to cross-region data transfer, select a backup repository that resides in the same region as instances that you plan to back up.
- To reduce high snapshot charges, adjust the snapshot retention settings to keep less restore points in the snapshot chain.
- To optimize the cost of storing backups, configure the scheduling settings to run the backup policy less frequently, or specify an archive backup repository for long-term retention of restore points.

For more information on cost estimation, see [this Veeam KB article](#).

TIP

You can save the cost estimation as a .CSV or .XML file. To do that, click **Export to** and select the necessary format.

Review cost estimation

The estimated cost takes into account the configured target settings, the specified scheduling options and the number of resources to protect.

Note that Veeam Backup for AWS makes predefined assumptions to calculate the cost, which means that the results should be used only as an approximation. For more information on cost calculation, see [this Veeam KB article](#).

3 instances will be backed up to a different region. This may significantly affect costs. Consider configuring backup policies per region to reduce costs.

\$16.56 Snapshots	\$13.56 Replicas	\$3.16 Backups	\$0.19 Archives	\$9.08 Traffic	\$0.02 Transactions
----------------------	---------------------	-------------------	--------------------	-------------------	------------------------

Estimated monthly cost: \$42.58

Instance Export to...

Instance ↑	Snapshot	Replica	Backup	Archive	Traffic	Transaction	Total
jf_access...	\$7.53	\$6.16	\$1.44	\$0.09	⚠ \$...	\$0.01	\$19.35
jf_vb_vpn	\$9.03	\$7.40	\$1.72	\$0.11	⚠ \$...	\$0.01	\$23.22
pi-WIP-b...	\$0.00	\$0.00	\$0.00	\$0.00	⚠ \$...	\$0.00	\$0.00

Previous Next Cancel

Related Resources

[How AWS Pricing Works](#)

Step 10. Finish Working with Wizard

At the **Summary** step of the wizard, it is recommended that you run the backup policy configuration check before you click **Finish**.

The configuration check will verify whether specified IAM roles have all the required permissions, and networks settings are configured properly to launch worker instances. To run the check, click **Test Configuration**. Veeam Backup for AWS will display the **Test policy configuration** window where you can track the progress and view the results of the check. If the IAM role permissions are insufficient or policy settings are not configured properly, the check will complete with errors, and the list of permissions that must be granted to the IAM role and policy configuration issues will be displayed in the **Test policy configuration** window.

You can grant the missing permissions to the IAM role [using the AWS Management Console](#) or instruct Veeam Backup for AWS to do it.

To let Veeam Backup for AWS grant the missing permissions:

1. In the **Test policy configuration** window, click the **Grant** link.
2. In the **Grant Permissions** window, provide one-time access keys of an IAM user that is authorized to update permissions of IAM roles, and then click **Apply**.

The IAM user must have the following permissions:

```
"iam:AttachRolePolicy",  
"iam:CreatePolicy",  
"iam:CreatePolicyVersion",  
"iam:CreateRole",  
"iam:GetAccountSummary",  
"iam:GetPolicy",  
"iam:GetPolicyVersion",  
"iam:GetRole",  
"iam:ListAttachedRolePolicies",  
"iam:ListPolicyVersions",  
"iam:SimulatePrincipalPolicy",  
"iam:UpdateAssumeRolePolicy"
```

NOTE

Veeam Backup for AWS does not store one-time access keys in the configuration database.

3. After the required permissions are granted, close the **Test policy configuration** window, and then click **Finish** to close the **Add Policy** wizard.

Veeam Backup for AWS will save the configured backup policy.

The screenshot displays the 'Add EC2 Policy' wizard in Veeam Backup for AWS. The top navigation bar shows the server time as Jun 11, 2024 1:44 PM and the user as administrator. The main content area is titled 'Add EC2 Policy' and shows a 'Review configured settings' step with a cost of \$42.58. A yellow banner at the top of the settings area states: 'In order to successfully run this policy, we advise to test the configuration.' Below this, there are buttons for 'Test Configuration' and 'Copy to Clipboard'. The settings are organized into several sections:

- General:** Name: ec2-frankfurt, Description: EC2 instances protection in Frankfurt, Regions: Europe (Frankfurt), Account: admin-407355668422
- Guest processing:** Application-aware snapshots: Enabled, Scripting for Linux instances: Enabled and configured, Scripting for Microsoft Windows instances: Disabled
- Snapshot settings:** Enabled: Yes
- Snapshot schedule:** Daily retention: Create 24 snapshots per day and keep 24 snapshots, Weekly retention: Keep 5 weekly snapshots, Monthly retention: Keep 6 monthly snapshots
- Replication settings:** Enabled: Yes

At the bottom of the wizard, there are three buttons: 'Previous', 'Finish', and 'Cancel'.

Fixing Network Issues

If the backup policy check reveals that network settings are not configured properly, Veeam Backup for AWS will not be able to deploy worker instances and thus perform image-level backup.

To fix network issues:

1. Close the **Test policy configuration** window, and then click **Finish** to close the **Add Policy** wizard. Veeam Backup for AWS will save the configured backup policy.
2. To prevent the backup policy from failing, disable it. For more information, see [Disabling and Enabling Policies](#).
3. Depending on the error message received after the backup policy check, do the following:
 - o Make sure that network settings are configured for each AWS Region selected at [step 3.2](#) of the wizard. For information on how to configure network settings for AWS Regions, see [Managing Worker Configurations](#).
 - o Make sure that VPCs specified in network settings for AWS Regions have access to the required AWS services. The required AWS services are listed in the [Planning and Preparation](#) section.
4. After network issues are fixed, you can enable the backup policy. For more information, see [Disabling and Enabling Policies](#).

Creating EC2 Snapshots Manually

Veeam Backup for AWS allows you to manually create snapshots of EC2 instances. You can instruct Veeam Backup for AWS to store the created snapshots in the same AWS Regions where the processed EC2 instances reside, or in a different AWS Region or AWS account.

NOTE

Veeam Backup for AWS does not include snapshots created manually in the snapshot chain and does not apply the configured retention policy settings to these snapshots. This means that the snapshots are kept in your AWS environment unless you remove them manually, as described in section [Managing Backed-Up EC2 Instance Data](#).

To manually create a cloud-native snapshot of an EC2 instance, do the following:

1. Navigate to **Resources > EC2**.
2. Select the necessary instance and click **Take Snapshot Now**.

For an EC2 instance to be displayed in the list of available instances, an AWS Region where the instance resides must be added to any of [configured EC2 backup policies](#), and the IAM role specified in the backup policy settings must have permissions to access the instance. For more information on required permissions, see [EC2 Backup IAM Role Permissions](#).

3. Complete the **Take Manual Snapshot** wizard:

- a. At the **Account** step of the wizard, specify an IAM role whose permissions Veeam Backup for AWS will use to create the snapshot.

For an IAM role to be displayed in the list, it must be added to Veeam Backup for AWS as described in section [Adding IAM Roles](#).

- b. At the **Snapshot Mode** step of the wizard, choose whether you want to store the snapshot in the same AWS Region where the processed EC2 instance resides, or in another AWS Region or AWS account.
- c. [Applies if you have selected the **New location** option] At the **Settings** step of the wizard, choose an IAM role whose permissions will be used to copy and store the snapshot in a target AWS Region, the target AWS Region and specify whether to encrypt the copied snapshot.
- d. At the **Tags** step of the wizard, choose whether you want to assign AWS tags to the created snapshot.

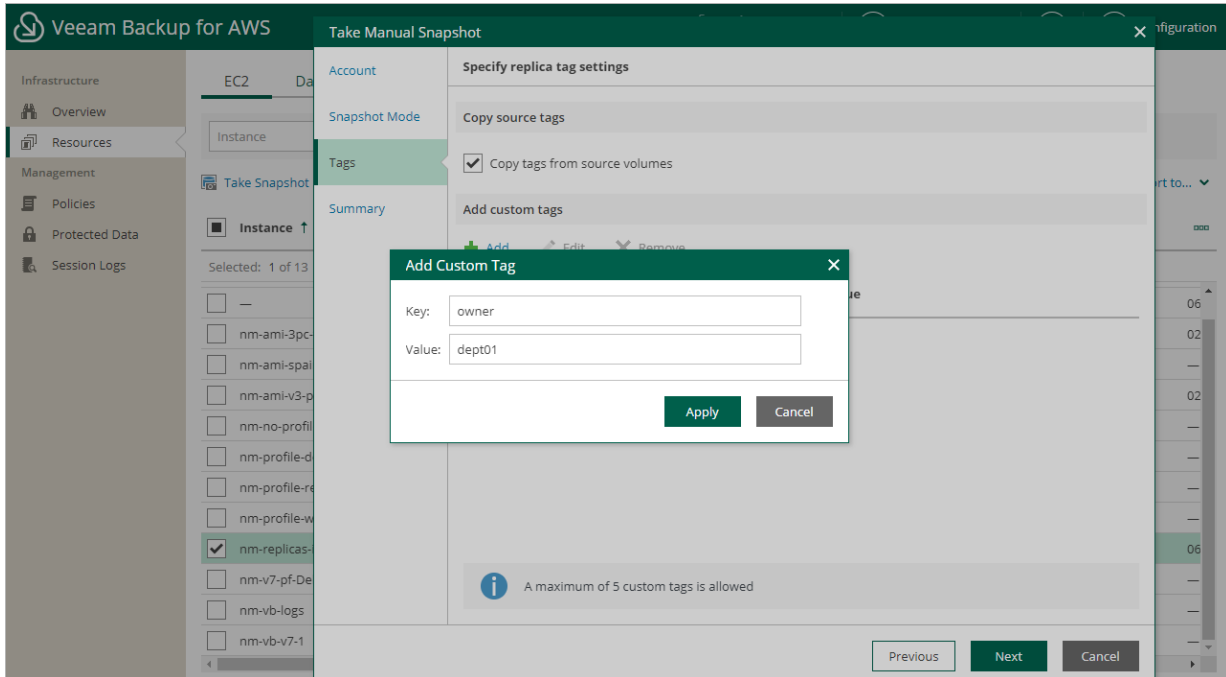
- To assign already existing AWS tags from the EBS volumes of the processed EC2 instance, select the **Copy tags from source volumes** check box.

If you choose to copy tags from source volumes, Veeam Backup for AWS will first create a snapshot of the EC2 instance and assign to the created snapshot AWS tags with Veeam metadata, then Veeam Backup for AWS will copy tags from the volumes of the processed instance and, finally, assign the copied AWS tags to the snapshot.

- To assign your own custom AWS tags, click **Add** and specify the tags explicitly. To do that, in the **Add Custom Tag** window, specify a key and a value for the new AWS tag, and then click **Apply**. Note that you cannot add more than 5 custom AWS tags.

If you choose to add custom tags to created snapshots, Veeam Backup for AWS will assign the specified tags right after it creates a snapshot.

e. At the **Summary** step of the wizard, review summary information and click **Finish**.



Performing RDS Backup

One backup policy can be used to process one or more RDS resources within one AWS account. The scope of data that you can protect in an AWS account is limited by permissions of an IAM role that is specified in the backup policy settings.

Before you create an RDS backup policy, check the following prerequisites:

- If you plan to create image-level backups of RDS resources, backup infrastructure components that will take part in the backup process must be added to the backup infrastructure and configured properly. These include [backup repositories](#) and [worker instances](#).
- If you plan to receive email notifications on RDS backup policy results, configure [email notification settings](#) first.

To schedule data protection tasks to run automatically, [create backup policies](#). For each protected DB instance and Aurora DB cluster, you can also [take a cloud-native snapshot manually](#) when needed.

IMPORTANT

Consider the following:

- Veeam Backup for AWS does not support backup of Oracle DB instances with multi-tenant architecture, as well as backup of PostgreSQL DB clusters with Multi-AZ DB cluster deployment and IBM Db2 DB instances.
- Veeam Backup for AWS allows you to create image-level backups of PostgreSQL DB instances only. For the list of supported PostgreSQL versions, see [Protecting RDS Resources](#).

Creating RDS Backup Policies

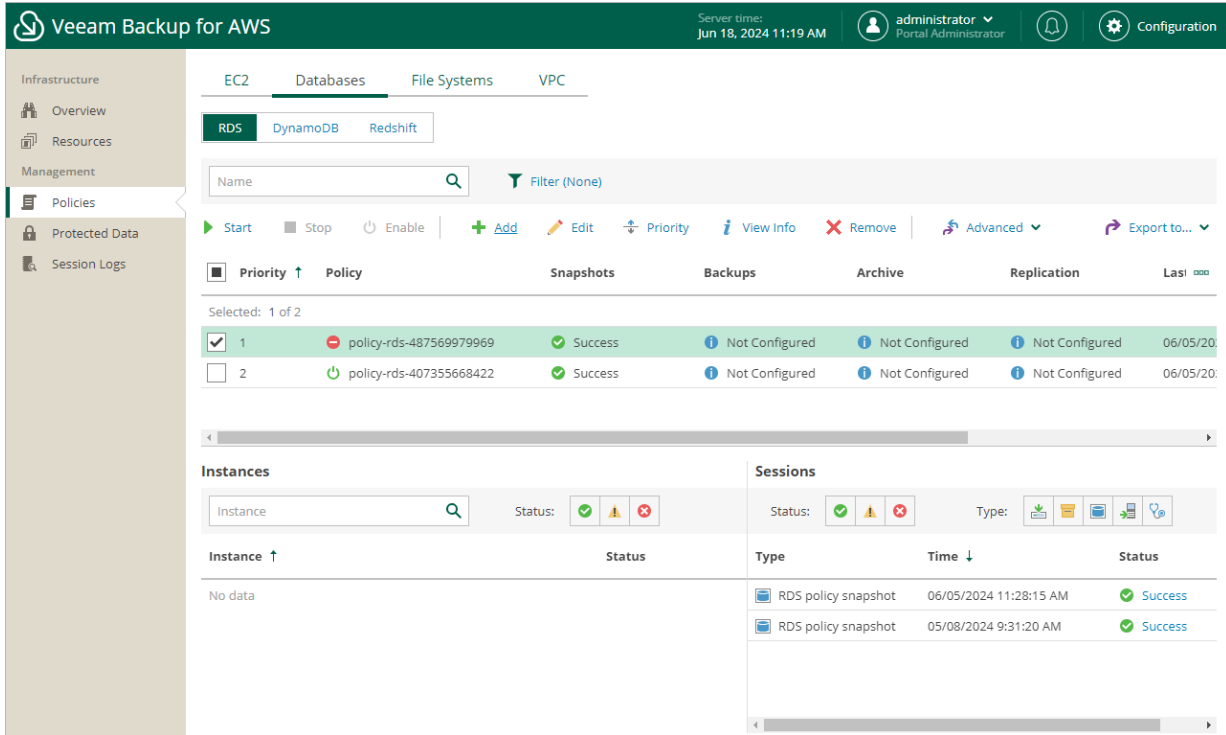
To create a backup policy, do the following:

1. [Launch the Add RDS Policy wizard](#).
2. [Specify a backup policy name and description](#).
3. [Configure backup source settings](#).
4. [Configure backup target settings](#).
5. [Specify processing settings](#).
6. [Specify a schedule for the backup policy](#).
7. [Enable AWS tags assigning](#).
8. [Specify automatic retry, health check and notification settings for the backup policy](#).
9. [Review estimated cost of the selected RDS resources](#).
10. [Finish working with the wizard](#).

Step 1. Launch Add RDS Policy Wizard

To launch the **Add RDS Policy** wizard, do the following:

1. Navigate to **Policies > Databases > RDS**.
2. Click **Add**.



Step 2. Specify Policy Name and Description

At the **Info** step of the wizard, use the **Name** and **Description** fields to specify a name for the new backup policy and to provide a description for future reference. The name must be unique in Veeam Backup for AWS; the maximum length of the name is 127 characters, the maximum length of the description is 255 characters.

The screenshot shows the 'Add RDS Policy' wizard in Veeam Backup for AWS. The interface is in the 'Info' step, where the user specifies the policy name and description. The 'Name' field contains 'rds-policy-eu' and the 'Description' field contains 'Created by administrator at 6/18/2024 11:19 AM'. The 'Cost' is shown as 'N/A' with a warning icon. The 'Next' button is highlighted in green, and the 'Cancel' button is in grey. The left sidebar shows the wizard steps: Info, Sources, Targets, Schedule, Tags, General Settings, Cost Estimation, and Summary. The top navigation bar includes the Veeam logo, server time (Jun 18, 2024 11:19 AM), user (administrator), and configuration options.

Veeam Backup for AWS

Server time: Jun 18, 2024 11:19 AM

administrator Portal Administrator

Configuration

← Add RDS Policy Cost: N/A ⚠

Info

Specify policy name and description

Enter a name and description for the policy.

Name:

rds-policy-eu

Description:

Created by administrator at 6/18/2024 11:19 AM

Next Cancel

Step 3. Configure Backup Source Settings

At the **Sources** step of the wizard, specify backup source settings:

1. [Select an IAM role whose permissions will be used to perform RDS backup.](#)
2. [Select AWS Regions where RDS resources that you plan to back up reside.](#)
3. [Select DB instances and Aurora DB clusters to back up.](#)

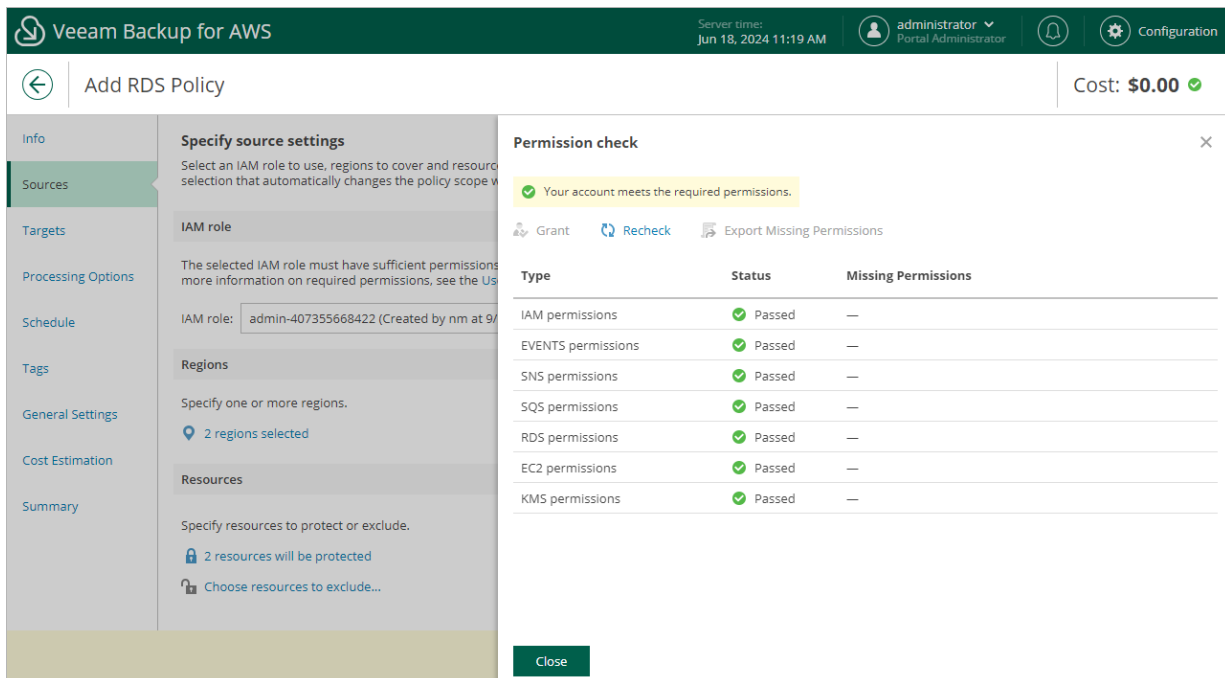
Step 3.1 Specify IAM Role

In the **IAM role** section of the **Sources** step of the wizard, specify an IAM role whose permissions will be used to access AWS services and resources, and to create cloud-native snapshots of DB instances and Aurora DB clusters. The specified IAM role must belong to the AWS account in which the RDS resources that you want to protect reside, and must be assigned the permissions listed in section [RDS Backup IAM Role Permissions](#).

For an IAM role to be displayed in the **IAM role** list, it must be added to Veeam Backup for AWS with the *Amazon RDS Snapshot* operation selected as described in section [Adding IAM Roles](#). If you have not added the necessary IAM role to Veeam Backup for AWS beforehand, you can do it without closing the **Add RDS Policy** wizard. To do that, click **Add** and complete the [Add IAM Role](#) wizard.

IMPORTANT

It is recommended that you check whether the selected IAM role has all the permissions required to perform the operation. If some permissions of the IAM role are missing, the backup policy will fail to complete successfully. To run the IAM role permission check, click **Check Permissions** and follow the instructions provided in section [Checking IAM Role Permissions](#).



The screenshot shows the 'Add RDS Policy' wizard in Veeam Backup for AWS. The 'Sources' step is selected, and the 'IAM role' section is expanded. The 'IAM role' field contains 'admin-407355668422 (Created by nm at 9/18/2024)'. The 'Regions' section shows '2 regions selected'. The 'Resources' section shows '2 resources will be protected'. A 'Permission check' dialog is open, displaying a table of permissions that have all passed.

Type	Status	Missing Permissions
IAM permissions	Passed	—
EVENTS permissions	Passed	—
SNS permissions	Passed	—
SQS permissions	Passed	—
RDS permissions	Passed	—
EC2 permissions	Passed	—
KMS permissions	Passed	—

Step 3.2 Select AWS Regions

In the **Regions** section of the **Sources** step of the wizard, choose AWS Regions where RDS resources that you plan to back up reside:

1. Click **Choose regions**.
2. In the **Choose regions** window, select the necessary regions and click **Add** to include them in the backup policy.
3. To save changes made to the backup policy settings, click **Apply**.

The screenshot shows the 'Add RDS Policy' wizard in Veeam Backup for AWS. The 'Sources' step is active, and the 'Regions' section is expanded. The 'Choose regions' dialog is open, showing a list of available regions and a 'Selected Regions' list. The 'Europe (Ireland)' region is selected in the available list, and 'Europe (Paris)' is in the selected list. The 'Apply' button is highlighted.

Veeam Backup for AWS | Server time: Jun 18, 2024 11:22 AM | administrator Portal Administrator | Configuration

Add RDS Policy | Cost: **N/A** ⚠

Specify source settings
Select an IAM role to use, regions to cover and resource selection that automatically changes the policy scope

IAM role
The selected IAM role must have sufficient permissions more information on required permissions, see the US
IAM role: admin-407355668422 (Created by nm at 9/

Regions
Specify one or more regions.
7 regions selected

Resources
Specify resources to protect or exclude.
Choose resources to protect...
Choose resources to exclude...

Choose regions

Available Regions (24)

- Asia Pacific (Mumbai)
- Asia Pacific (Osaka)
- Asia Pacific (Seoul)
- Asia Pacific (Singapore)
- Asia Pacific (Sydney)
- Asia Pacific (Tokyo)
- Canada (Central)
- Canada West (Calgary)
- Europe (Frankfurt)
- Europe (Ireland)**
- Europe (London)
- Europe (Spain)
- Europe (Stockholm)
- Europe (Zurich)

Selected Regions (1)

- Europe (Paris)

Buttons: Add, Remove, Apply, Cancel

Step 3.3 Select RDS Resources

In the **Resources** section of the **Sources** step of the wizard, specify the backup scope – select DB instances and Aurora DB clusters that Veeam Backup for AWS will back up:

1. Click **Choose resources to protect**.
2. In the **Choose resources to protect** window, choose whether you want to back up all RDS resources from AWS Regions selected at [step 3.2](#) of the wizard or only specific RDS resources.

If you select the **All resources** option, Veeam Backup for AWS will regularly check for new DB instances and Aurora DB clusters launched in the selected regions and automatically update the backup policy settings to include these resources into the backup scope.

If you select the **Protect only following resources** option, you must also specify the resources explicitly:

- a. Use the **Type** drop-down list to choose whether you want to add individual RDS resources or AWS tags to the backup scope.

If you select the **Tag** option, Veeam Backup for AWS will back up only those resources from the selected AWS Regions that are assigned specific tags.

- b. Use the search field to the right of the **Type** list to find the necessary resource, and then click **Protect** to add the resource to the backup scope.

For a resource to be displayed in the list of available resources, it must reside in an AWS Region that has ever been specified in any backup policy. Otherwise, the only option to discover the available resources is to click **Browse to select specific resources from the global list** and to wait for Veeam Backup for AWS to populate the resource list.

TIP

You can simultaneously add multiple resources to the backup scope. To do that, click **Browse to select specific sources from the global list**, select check boxes next to the necessary RDS resources or AWS tags in the list of available resources, and then click **Protect**.

If the list does not show the resources that you want to back up, click **Rescan** to launch the data collection process. As soon as the process is over, Veeam Backup for AWS will update the resource list.

If you add an AWS tag to the backup scope, Veeam Backup for AWS will regularly check for new RDS resources assigned the added AWS tag and automatically update the backup policy settings to include these resources in the scope. However, this applies only to DB instances and Aurora DB clusters from the AWS Regions selected at [step 3.2](#) of the wizard. If you select an AWS tag assigned to RDS resources from other AWS Regions, these resources will not be protected by the backup policy. To work around the issue, either go back to [step 3.2](#) and add the missing AWS Regions, or create a new backup policy.

3. To save changes made to the backup policy settings, click **Apply**.

TIP

As an alternative to selecting the **Protect only following resources** option and specifying the resources explicitly, you can select the **All resources** option and exclude a number of resources from the backup scope. To do that, click **Choose resources to exclude** and specify the resources or tags that you do not want to protect – the procedure is the same as described for including resources in the backup scope.

Note that if a resource appears both in the list of included and excluded resources, Veeam Backup for AWS will still not process the resource because the list of excluded resources has a higher priority.

The screenshot shows the 'Add RDS Policy' configuration window in Veeam Backup for AWS. The 'Choose resources to protect' dialog is open, showing the 'Protect only following resources' option selected. The dialog includes a search bar, a list of protected resources, and buttons for 'Apply' and 'Cancel'.

Item ↑	ID	Value	Region
bd-paris-rds-postgres...	db-jp2vuww5ohncvoyjstu...	—	Europe (Paris)

Step 4. Configure Backup Target Settings

By default, backup policies create only cloud-native snapshots of processed instances. At the **Targets** step of the wizard, you can enable the following additional data protection scenarios:

- [Instruct Veeam Backup for AWS to replicate cloud-native snapshots to other AWS accounts or AWS Regions.](#)
- [Instruct Veeam Backup for AWS to create image-level backups.](#)

IMPORTANT

Creating image-level backups is supported for PostgreSQL DB instances only. For the list of supported PostgreSQL versions, see [Protecting RDS Resources](#).

Configuring Snapshot Replica Settings

If you want to replicate cloud-native snapshots to other AWS accounts or regions, do the following:

1. In the **Snapshots** section of the **Targets** step of the wizard, set the **Replicate snapshots** toggle to *On*.
2. In the **Replication settings** window, configure the following mapping settings for each AWS Region where source instances reside:

IMPORTANT

If DB engine versions of the processed Aurora DB clusters are not supported in the target AWS Region, the replication operation will fail. For the list of supported DB engine versions in AWS Regions, see [AWS Documentation](#).

- a. Select a source AWS Region from the list and click **Edit Region Mapping**.
- b. In the **Edit Region Mapping** window, specify the following settings:
 - i. From the **Target account** drop-down list, select an IAM role whose permissions will be used to copy and store cloud-native snapshots in a target AWS Region.

If you select an IAM role created in another AWS account, the cloud-native snapshots will be copied to the target AWS Region in that AWS account.
 - ii. From the **Target region** drop-down list, select the target AWS Region to which Veeam Backup for AWS must copy cloud-native snapshots.
 - iii. If you want to encrypt the cloud-native snapshots copied to the target AWS Region, select the **Enable encryption** check box and choose the necessary KMS key from the **Encryption key** drop-down list. For a KMS key to be displayed in the list of available encryption keys, it must be stored in the AWS Region selected at [step 3.2](#) of the wizard and the IAM role specified for the backup operation must have permissions to access the key. For more information on KMS keys, see [AWS Documentation](#).

Then, use the **Key usage** drop-down list to choose whether you want to encrypt snapshots for all resources or only snapshots of the encrypted resources.

NOTE

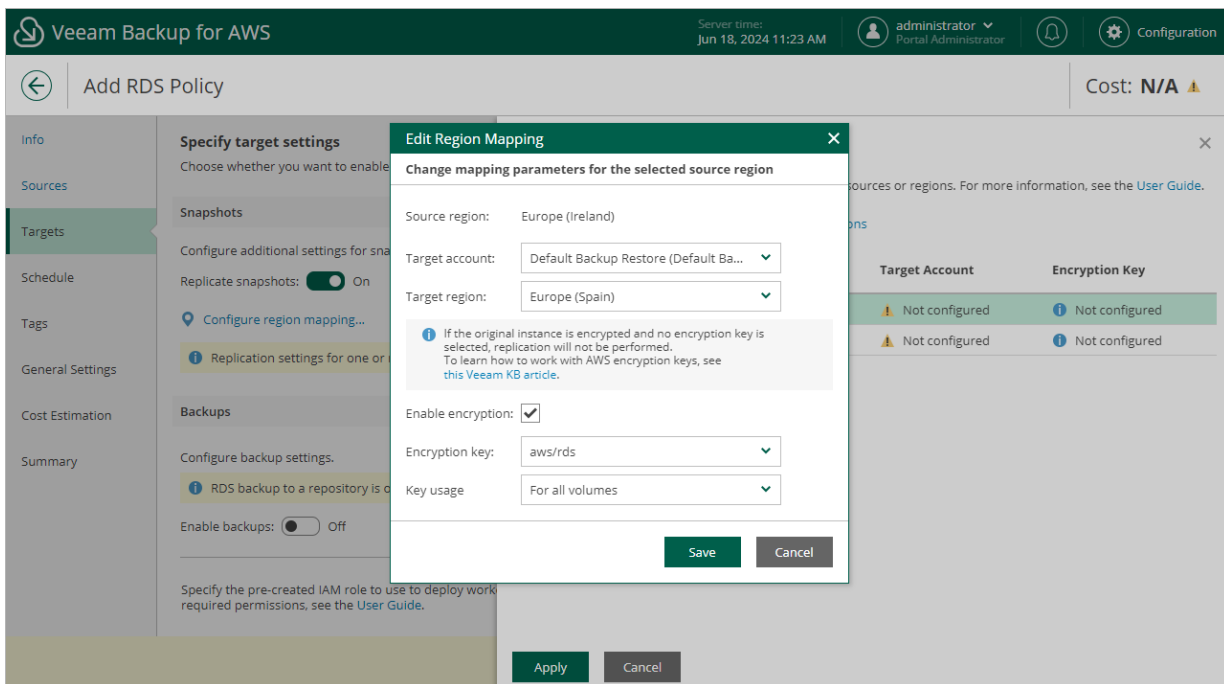
Consider the following:

- If the source DB instances or Aurora DB clusters are encrypted, you must enable encryption for replicated snapshots. Otherwise the replication process will fail.
- If the source Aurora DB cluster is unencrypted, the encryption must be disabled for replicated snapshots. Otherwise the replication process will fail.

iv. Click **Save**.

To configure mapping for all source AWS Regions at once, click **Set Mapping for All Regions** and specify settings as described at [step 2.b](#).

c. To save changes made to the backup policy settings, click **Apply**.



Configuring Image-Level Backup Settings

In the **Backups** section of the **Targets** step of the wizard, you can instruct Veeam Backup for AWS to create image-level backups of the processed DB instances and to copy backups to a long-term archive storage.

NOTE

To create RDS image-level backups, Veeam Backup for AWS deploys worker instances in a production account – that is, the same AWS account to which the processed resources belong. For more information, see [Worker Deployment Options](#).

Configuring Backup Settings

To instruct Veeam Backup for AWS to create image-level backups of the selected RDS resources, do the following:

1. Set the **Enable backups** toggle to *On*.

2. In the **Repositories** window, select a backup repository where the created image-level backups will be stored, and click **Apply**.

For a backup repository to be displayed in the list of available repositories, it must be added to Veeam Backup for AWS as described in section [Adding Backup Repositories](#). The list shows only backup repositories of the *S3 Standard* storage class.

To learn how Veeam Backup for AWS creates image-level backups, see [RDS Backup](#).

Configuring Archive Settings

To instruct Veeam Backup for AWS to store backed-up data in a low-cost, long-term archive storage, do the following:

1. Select the **Archives will be stored in** check box.
2. In the **Repositories** window, select a backup repository where the archived data will be stored, and click **Apply**.

For an archive backup repository to be displayed in the list of available repositories, it must be added to Veeam Backup for AWS as described in section [Adding Backup Repositories](#). The list shows only backup repositories of the *S3 Glacier Flexible Retrieval* or *S3 Glacier Deep Archive* storage classes.

For more information on backup archiving, see [Enabling Backup Archiving](#).

IMPORTANT

If you enable the backup archiving, consider that data encryption must be either enabled or disabled for both backup and archive backup repositories. This means that, for example, you cannot select an encrypted standard backup repository and an unencrypted archive backup repository in one backup policy. However, the selected repositories can have different encryption schemes (password and KMS encryption).

Configuring Worker Settings

From the **IAM role** drop-downlist, select an IAM role that will be attached to the worker instances and used by Veeam Backup for AWS to communicate with these instances. The role must belong to the same account to which the IAM role specified to perform the backup operation belongs and must be assigned permissions listed in section [Worker Deployment Role Permissions in Production Accounts](#).

For an IAM role to be displayed in the **IAM role** list, it must be added to Veeam Backup for AWS with the *Production worker role* selected as described in section [Adding IAM Roles](#). If you have not added the necessary IAM role to Veeam Backup for AWS beforehand, you can do it without closing the **Add RDS Policy** wizard. To do that, click **Add** and complete the **Add IAM Role** wizard.

IMPORTANT

Consider the following:

- For Veeam Backup for AWS to deploy worker instances in production accounts, you must assign additional permissions to the IAM role used to perform the backup operation. For more information on the required permissions, see section [RDS Backup IAM Role Permissions](#).
- It is recommended that you check whether both the IAM role specified at [step 3.1](#) of the wizard and the IAM role specified in the **Backups** section have the required permissions. If some permissions of the IAM role are missing, the backup policy will fail to complete successfully. To run the IAM role permission check, click **Check Permissions** and follow the instructions provided in section [Checking IAM Role Permissions](#).

Worker Instance Requirements

To perform RDS image-level backups, Veeam Backup for AWS deploys worker instances in production accounts in the same AWS Regions and VPCs in which processed PostgreSQL DB instances reside. By default, Veeam Backup for AWS uses the most appropriate network settings of AWS Regions in production accounts to deploy worker instances. However, you can add [specific worker configurations](#) to specify network settings for each region in which worker instances will be deployed.

If no [specific worker configurations](#) are added to Veeam Backup for AWS, the most appropriate network settings of AWS Regions are used to deploy worker instances for the RDS backup operation. For Veeam Backup for AWS to be able to deploy a worker instance used to create an image-level backup:

- The DNS resolution settings option must be enabled for the VPC. For more information, see [AWS Documentation](#).
- As Veeam Backup for AWS uses public access to communicate with worker instances, the [public IPv4 addressing](#) attribute must be enabled at least for one subnet in the Availability Zone where the DB instance resides and the VPC to which the subnet belongs must have an [internet gateway attached](#). VPC and subnet route tables must have routes that direct internet-bound traffic to this internet gateway.

If you want worker instances to operate in a private network, enable the [private network deployment](#) functionality and configure [specific VPC endpoints](#) for the subnet to let Veeam Backup for AWS use private IPv4 addresses. Alternatively, configure VPC interface endpoints as described in section [Appendix C. Configuring Endpoints in AWS](#).

NOTE

During RDS image-level backup operations, Veeam Backup for AWS creates 2 additional security groups that are further associated with the source DB instances and worker instances to allow direct network traffic between them. To learn how RDS resource backup works, see [RDS Backup](#).

The screenshot displays the Veeam Backup for AWS console interface. The main window is titled 'Add RDS Policy' and shows a 'Permission check' dialog box. The dialog box contains the following information:

- Message: "Your account does not meet the required permissions."
- Buttons: Grant, Recheck, Export Missing Permissions
- Table:

Type	Status	Missing Permissions
Checking backup policy role ...	Passed	—
Checking worker role permis...	Failed	Trust Policy: The following service must be added to t...

Selected: 1 of 2

The 'Grant Permissions' dialog box is overlaid on top, with the following details:

- Title: Grant Permissions
- Section: Provide temporary credentials
- Message: "You can grant permissions manually in the AWS Management Console or automatically using the form below. These keys are not saved or stored. For more information on how to assign missing permissions to an IAM role, see the User Guide."
- Access key: AKIAI4ZWOU4WMVRAGEVN
- Secret key: [Redacted]
- Buttons: Apply, Cancel

Step 5. Specify Processing Settings

[This step applies only if you have enabled image-level backups at the **Targets** step of the wizard]

At the **Processing Options** step of the wizard, select a database account whose credentials will be used to authenticate against databases of the DB instances added to the backup scope. For an account to be displayed in the list of available accounts, it must be added to Veeam Backup for AWS as described in section [Adding Database Accounts](#). If you have not added the necessary account to Veeam Backup for AWS beforehand, you can do it without closing the **Add RDS Policy** wizard. To do that, click **Add** and complete the **Add Account** wizard.

By default, the selected account will be used to access all databases of the DB instances added to the backup policy. You can also granularly specify credentials that Veeam Backup for AWS will use to connect to specific DB instances. To do that, set the **Customize credentials** toggle to *On*, choose a DB instance for which you want to specify the credentials and click **Edit Credentials**.

IMPORTANT

For Veeam Backup for AWS to be able to protect the DB instance added to the backup policy, the selected account must exist on this instance.

The screenshot shows the 'Add RDS Policy' wizard in Veeam Backup for AWS. The 'Processing Options' step is active, showing 'Specify database processing settings'. A yellow warning banner states: 'RDS backup to a repository is only supported for PostgreSQL.' Below this, the 'Default credentials' are listed as 'database_acc_dept01 (admin)'. A note explains that these default credentials will be used for all protected databases unless customized. The 'Customize credentials' toggle is turned 'On'. There are links for 'Edit Credentials' and 'Use Default'. A table lists two PostgreSQL instances: 'rds-postgres' (selected) and 'ma-external'. At the bottom, there are 'Previous', 'Next', and 'Cancel' buttons.

Item	Engine	Credentials
<input checked="" type="checkbox"/> rds-postgres	PostgreSQL	Default
<input type="checkbox"/> ma-external	PostgreSQL	Default

Step 6. Specify Policy Scheduling Options

You can instruct Veeam Backup for AWS to start the backup policy automatically according to a specific backup schedule. The backup schedule defines how often data of the instances added to the backup policy must be backed up.

IMPORTANT

If you have selected a standard or an archive backup repository with immutability settings enabled at [step 4](#) of the wizard, you must configure at least one schedule for the backup policy.

To help you implement a comprehensive backup strategy, Veeam Backup for AWS allows you to create schedules of the following types:

- **Daily** – the backup policy will create restore points repeatedly throughout a day on specific days.
- **Weekly** – the backup policy will create restore points once a day on specific days.
- **Monthly** – the backup policy will create restore points once a month on a specific day.
- **Yearly** – the backup policy will create restore points once a year on a specific day.

Combining multiple schedule types together allows you to retain restore points for longer periods of time. For more information, see [Enabling Harmonized Scheduling](#).

NOTE

If you do not specify the backup schedule after you configure the backup policy, you will need to start it manually to create RDS snapshots and backups. For information on how to start backup policies, see [Starting and Stopping Policies](#).

Specifying Daily Schedule

To create a daily schedule for the backup policy, at the **Schedule** step of the wizard, do the following:

1. Set the **Daily schedule** toggle to *On* and click **Edit Daily Settings**.
2. In the **Create daily schedule** window, select hours when the backup policy must create cloud-native snapshots, snapshot replicas or image-level backups.

If you want to protect RDS resources data more frequently, you can instruct the backup policy to create multiple cloud-native snapshots per hour. To do that, click the link to the right of the **Snapshots** hour selection area, and specify the number of cloud-native snapshots that the backup policy must create within an hour.

NOTE

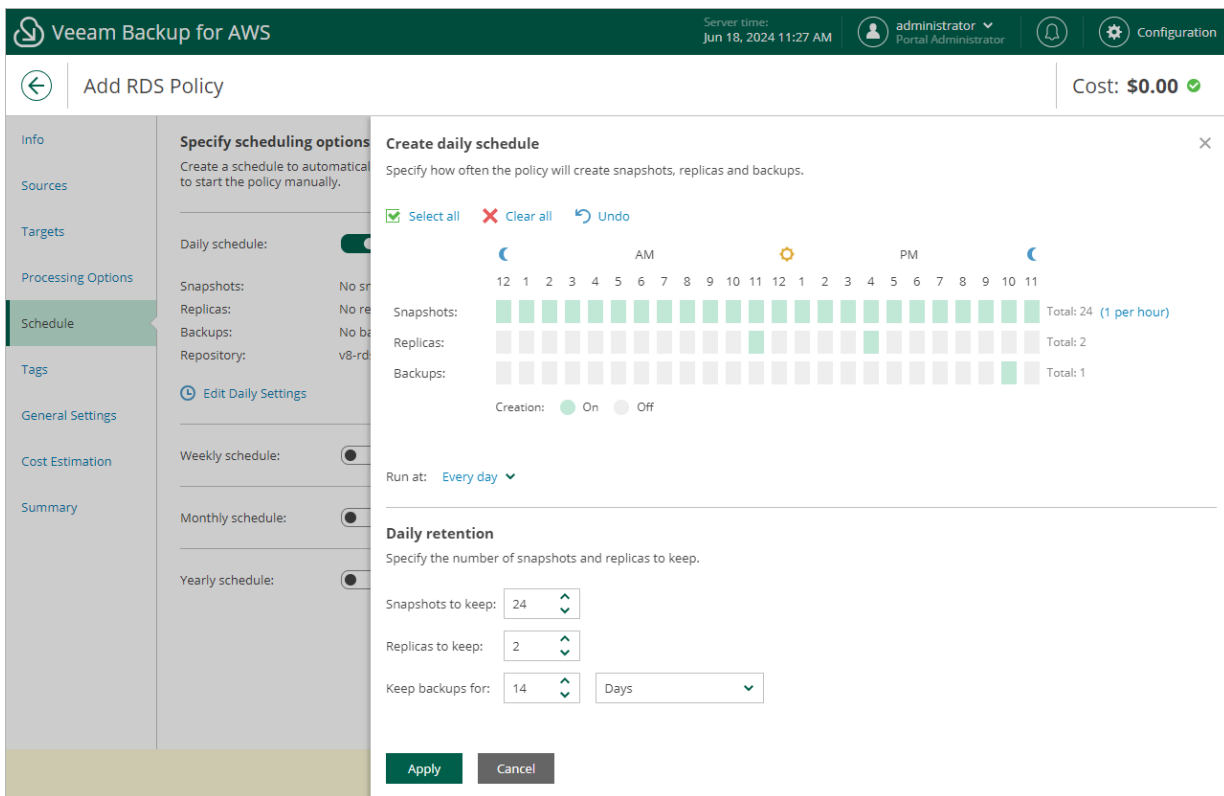
Veeam Backup for AWS does not create snapshot replicas and image-level backups independently from cloud-native snapshots. That is why when you select hours to create snapshot replicas and image-level backups, the same hours are automatically selected for cloud-native snapshots. To learn how Veeam Backup for AWS performs backup, see [RDS Backup](#).

3. Use the **Run at** drop-down list to choose whether you want the backup policy to run everyday, on work days (Monday through Friday) or on specific days.

4. In the **Daily retention** section, configure retention policy settings for the daily schedule:
 - For cloud-native snapshots and snapshot replicas, specify the number of restore points that you want to keep in cloud-native snapshot and snapshot replica chains.

If the restore point limit is exceeded, Veeam Backup for AWS removes the earliest restore point from the chain. For more information, see [RDS Snapshot Retention](#).
 - For image-level backups, specify the number of days (or months) for which you want to keep restore points in a backup chain.

If a restore point is older than the specified time limit, Veeam Backup for AWS removes the restore point from the chain. For more information, see [RDS Backup Retention](#).
5. To save changes made to the backup policy settings, click **Apply**.



Specifying Weekly Schedule

To create a weekly schedule for the backup policy, at the **Schedule** step of the wizard, do the following:

1. Set the **Weekly schedule** toggle to *On* and click **Edit Weekly Settings**.
2. In the **Create weekly schedule** window, select weekdays when the backup policy must create cloud-native snapshots, snapshot replicas or image-level backups.

NOTE

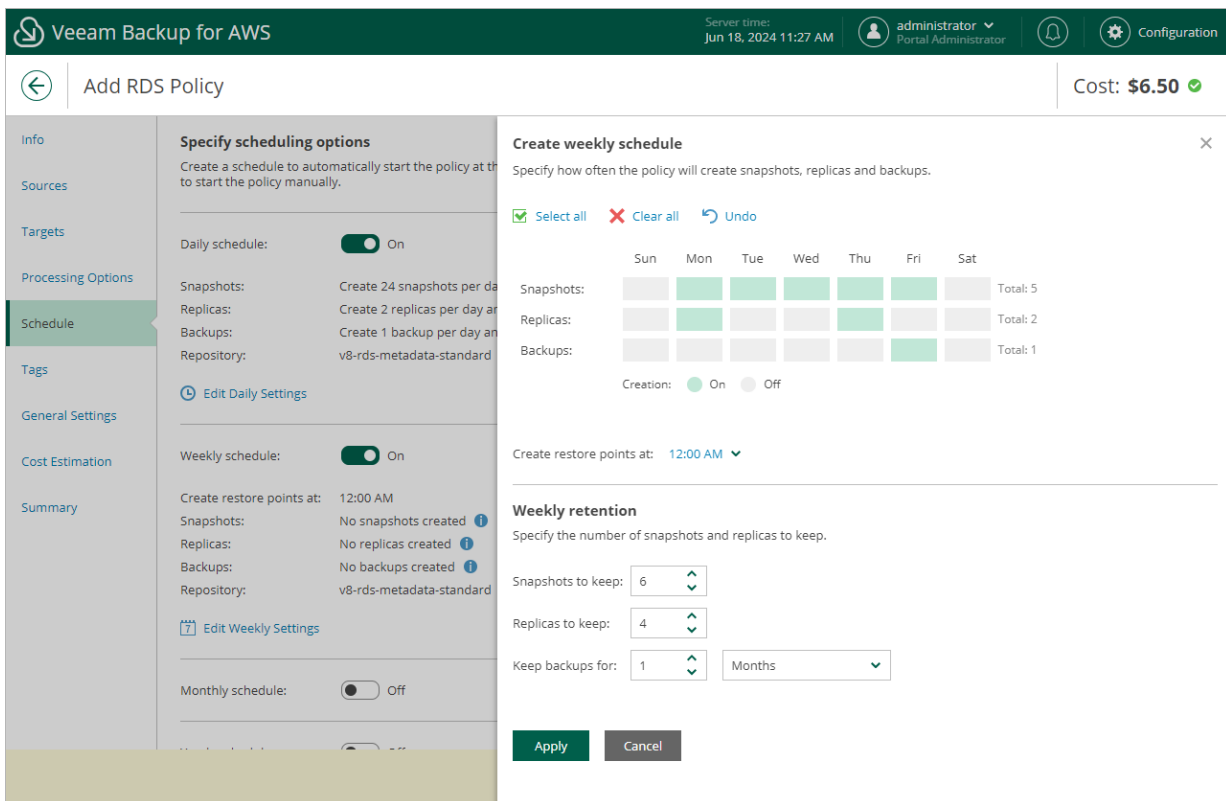
Veeam Backup for AWS does not create snapshot replicas and image-level backups independently from cloud-native snapshots. That is why when you select days to create snapshot replicas and image-level backups, the same days are automatically selected for cloud-native snapshots. To learn how Veeam Backup for AWS performs backup, see [RDS Backup](#).

3. Use the **Create restore point at** drop-down list to schedule a specific time for the backup policy to run.

4. In the **Weekly retention** section, configure retention policy settings for the weekly schedule:
 - For cloud-native snapshots and snapshot replicas, specify the number of restore points that you want to keep in cloud-native snapshot and snapshot replica chains.

If the restore point limit is exceeded, Veeam Backup for AWS removes the earliest restore point from the chain. For more information, see [RDS Snapshot Retention](#).
 - For image-level backups, specify the number of days (or months) for which you want to keep restore points in a backup chain.

If a restore point is older than the specified time limit, Veeam Backup for AWS removes the restore point from the chain. For more information, see [RDS Backup Retention](#).
5. To save changes made to the backup policy settings, click **Apply**.



Specifying Monthly Schedule

To create a monthly schedule for the backup policy, at the **Schedule** step of the wizard, do the following:

1. Set the **Monthly schedule** toggle to *On* and click **Edit Monthly Settings**.
2. [This step applies if you have enabled backup archiving at the [Targets](#) step of the wizard] In the **Create monthly schedule** section of the opened window, choose whether you want to store monthly backups in the archive repository.

If you set the **Send backups to archive** toggle to *On*, follow the instructions provided in section [Enabling Backup Archiving](#).
3. In the **Create monthly schedule** window, select months when the backup policy must create cloud-native snapshots, snapshot replicas or image-level backups.

NOTE

Veeam Backup for AWS does not create snapshot replicas and image-level backups independently from cloud-native snapshots. That is why when you select months to create snapshot replicas and image-level backups, the same months are automatically selected for cloud-native snapshots. To learn how Veeam Backup for AWS performs backup, see [RDS Backup](#).

4. Use the **Create restore point at** and **Run on** drop-down lists to schedule a specific time and day for the backup policy to run.

NOTE

Consider the following:


- If you have selected a specific time for the backup policy to run at the **Weekly schedule** section of the **Schedule** step of the wizard, you will not be able to change the time for the monthly schedule unless you select the **On Day** option from the **Run on** drop-down list.
- If you select the **On day** option, [harmonized scheduling](#) cannot be guaranteed. Plus, to support the **On day** option, Veeam Backup for AWS will require to create an additional temporary restore point if there are no other schedules planned to run on that day. However, the temporary restore point will be removed by the *Backup Retention* process from AWS within approximately 24 hours, to reduce unexpected infrastructure charges.

5. In the **Monthly retention** section, configure retention policy settings for the monthly schedule:
 - For cloud-native snapshots and snapshot replicas, specify the number of restore points that you want to keep in cloud-native snapshot and snapshot replica chains.

If the restore point limit is exceeded, Veeam Backup for AWS removes the earliest restore point from each chain. For more information, see [RDS Snapshot Retention](#).
 - For image-level backups, specify the number of days (or months) for which you want to keep restore points in a backup chain.

If a restore point is older than the specified time limit, Veeam Backup for AWS removes the restore point from the chain. For more information, see [RDS Backup Retention](#).

6. To save changes made to the backup policy settings, click **Apply**.



Specifying Yearly Schedule

[This step applies only if you have instructed Veeam Backup for AWS to create image-level backups at the **Targets** step of the wizard]

To create a yearly schedule for the backup policy, at the **Schedule** step of the wizard, do the following:

1. Set the **Yearly schedule** toggle to *On* and click **Edit Yearly Settings**.
2. [This step applies if you have enabled backup archiving at the **Targets** step of the wizard] In the **Create yearly schedule** section of the opened window, choose whether you want to store yearly backups in the archive backup repository.

If you set the **Send backups to archive** toggle to *On*, follow the instructions provided in section [Enabling Backup Archiving](#).

3. In the **Yearly schedule** section, specify a day, month and time when the backup policy will create image-level backups.

For example, if you select *First, Friday, January* and *06:00 PM*, the backup policy will run every first Friday of January at 06:00 PM.

NOTE

Consider the following:

- If you have selected a specific time and day for the backup policy to run at the **Weekly schedule** or **Monthly schedule** sections of the **Schedule** step of the wizard, you will not be able to change the time and day for the yearly schedule unless you select the *On Day* option from the **Create restore point on** drop-down list.
- If you select the *On day* option, **harmonized scheduling** cannot be guaranteed. Plus, to support the **On day** option, Veeam Backup for AWS will require to create an additional temporary restore point if there are no other schedules planned to run on that day. However, the temporary restore point will be removed by the *Backup Retention* process from AWS within approximately 24 hours, to reduce unexpected infrastructure charges.

4. In the **Keep backups for** field, specify the number of years for which you want to keep restore points in a backup chain.

If a restore point is older than the specified time limit, Veeam Backup for AWS removes the restore from the chain. For more information, see [RDS Backup Retention](#).

5. To save changes made to the backup policy settings, click **Apply**.

The screenshot shows the 'Add RDS Policy' configuration interface. The 'Schedule' section is highlighted in the left sidebar. The main content area shows the 'Create yearly schedule' dialog box. The dialog box has the following settings:

- Send backups to archive: On
- Create restore points on: First of July at 12:00 AM
- Keep archives for: 2 years

The dialog box has 'Apply' and 'Cancel' buttons. The top navigation bar shows 'Veeam Backup for AWS', 'Server time: Jun 18, 2024 11:28 AM', and 'administrator Portal Administrator'. The top right corner shows 'Cost: \$33.98' with a green checkmark.

Enabling Harmonized Scheduling

When you combine multiple types of schedules, Veeam Backup for AWS applies the harmonization mechanism that allows you to leverage restore points for long-term retentions instead of taking a new restore point every time. The mechanism simplifies the backup schedule, optimizes the backup performance and reduces the cost of retaining restore points.


With harmonized scheduling, Veeam Backup for AWS can keep restore points created according to a daily or weekly schedule for longer periods of time: cloud-native snapshots and snapshot replicas can be kept for weeks and months.

For Veeam Backup for AWS to use the harmonization mechanism, there must be specified at least 2 different schedules: one schedule will control the regular creation of restore points, while another schedule will control the process of storing restore points. In terms of harmonized scheduling, Veeam Backup for AWS re-uses restore points created according to a more-frequent schedule (daily or weekly) to achieve the desired retention for less-frequent schedules (weekly and monthly). Each restore point is marked with a flag of the related schedule type: the (D) flag is used to mark restore points created daily, (W) – weekly, and (M) – monthly. Veeam Backup for AWS uses these flags to control the retention period for the created restore points. Once a flag of a less-frequent schedule is assigned to a restore point, this restore point can no longer be removed – it is kept for the period defined in the retention settings. When the specified retention period is over, the flag is unassigned from the restore point. If the restore point does not have any other flags assigned, it is removed according to the retention settings of a more-frequent schedule.

Consider the following example. You want a backup policy to create cloud-native snapshots of your critical workloads 3 times a day, to keep 3 daily snapshots in the snapshot chain, and also to keep one of the created snapshots for 2 weeks. In this case, you create 2 schedules when configuring the backup policy settings – daily and weekly:

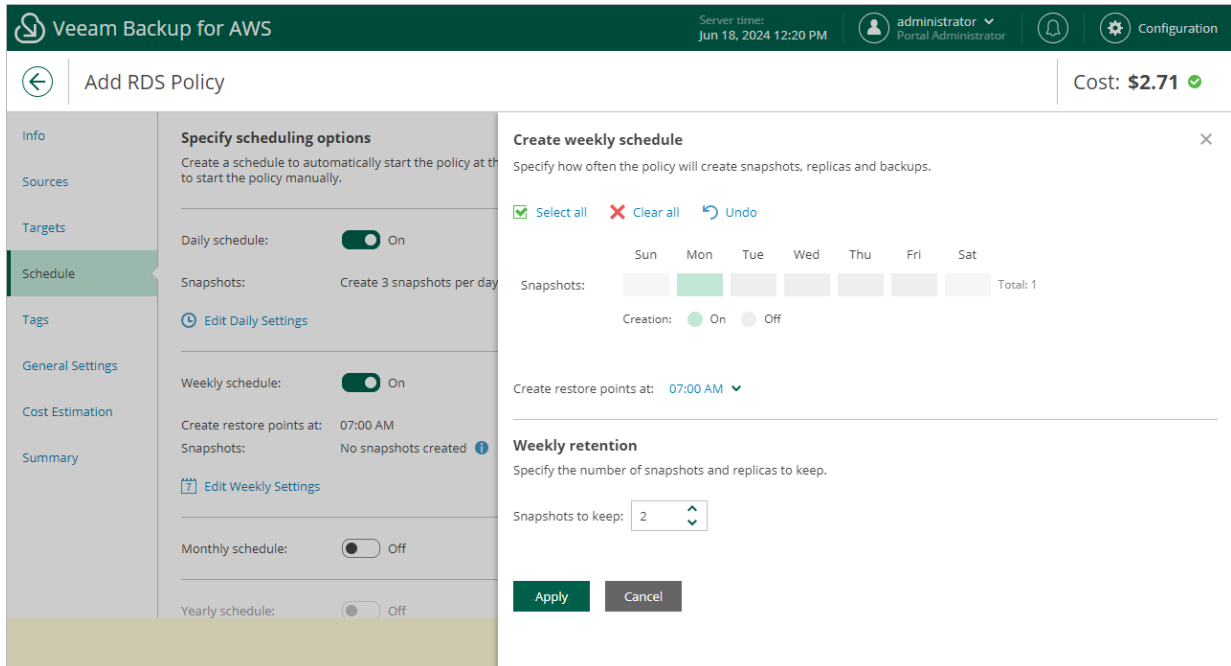
- In the daily scheduling settings, you select hours and days when snapshots will be created (for example, 7:00 AM, 9:00 AM, and 11:00 AM; Weekdays), and specify a number of daily restore points to retain (for example, 3).

Veeam Backup for AWS will propagate these settings to the schedule of a lower frequency (which is the weekly schedule in our example).



- In the weekly scheduling settings, you specify which one of the snapshots created by the daily schedule will be retained for a longer period, and choose for how long you want to keep the selected snapshot.

For example, if you want to keep the daily restore point created at 7:00 AM on Monday for 2 weeks, you select *7:00 AM, Monday* and specify 2 restore points to retain in the weekly schedule settings.




According to the specified scheduling settings, Veeam Backup for AWS will create cloud-native snapshots in the following way:

- On the first work day (Monday), a backup session will start at 7:00 AM to create the first restore point. The restore point will be marked with the (D) flag as it was created according to the daily schedule.

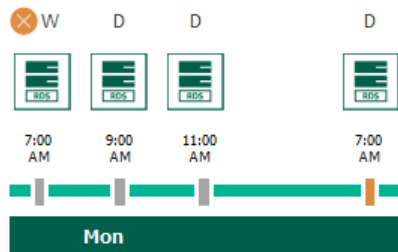
Since *7:00 AM, Monday* is specified in the weekly schedule settings, Veeam Backup for AWS will also assign the (W) flag to this restore point. As a result, 2 flags (D,W) will be assigned to the restore point.

- On the same day (Monday), after backup sessions run at 9:00 AM and 11:00 AM, the created restore points will be marked with the (D) flag.




- On the next work day (Tuesday), after a backup session runs at 7:00 AM, the created restore point will be marked with the (D) flag.

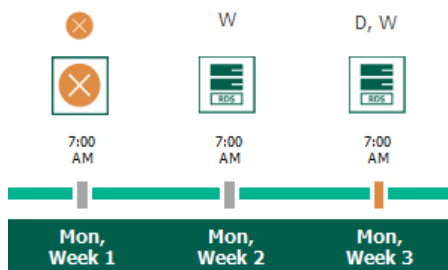
By the moment the backup session completes, the number of restore points with the (D) flag will exceed the retention limit specified in the daily schedule settings. However, Veeam Backup for AWS will not remove the earliest restore point (7:00 AM, Monday) with the (D) flag from the snapshot chain as this restore point is also marked with a flag of a less-frequent schedule. Instead, Veeam Backup for AWS will unassign the (D) flag from the restore point. This restore point will be kept for the retention period specified in the weekly schedule settings (that is, for 2 weeks).



- On the same day (Tuesday), after a backup session runs at 9:00 AM, the number of restore points with the (D) flag will exceed the retention limit once again. Veeam Backup for AWS will remove from the snapshot chain the restore point created at 9:00 AM on Monday as no flags of a less-frequent schedule are assigned to this restore point.



- Veeam Backup for AWS will continue creating restore points for the next week in the same way as described in steps 1–4.
- On week 3, after a backup session runs at 7:00 AM on Monday, the number of weekly restore points will exceed the retention limit. Veeam Backup for AWS will unassign the (W) flag from the earliest weekly restore point. Since no other flags are assigned to this restore point, Veeam Backup for AWS will remove this restore point from the snapshot chain.



Enabling Backup Archiving

When you combine multiple types of schedules, you can enable the archiving mechanism to instruct Veeam Backup for AWS to store backed-up data in the secure, low-cost and long-term S3 Glacier Flexible Retrieval and S3 Glacier Deep Archive storage classes. The mechanism is the most useful in the following cases:

- Your data retention policy requires that you keep rarely accessed data in an archive.

- You want to reduce data-at-rest costs and to save space in the high-cost, short-term S3 standard storage class.

NOTE

Restoring from an archived backup is longer and more expensive than restoring from a regular backup as it is required to retrieve data from the archive repository. For more information, see [Performing Database Restore](#).

With backup archiving, Veeam Backup for AWS can retain backup files created according to a daily, weekly or monthly schedule for longer periods of time:

- To enable monthly archiving, you must configure a daily or a weekly schedule (or both).
- To enable yearly archiving, you must configure a daily, a weekly or a monthly schedule (or all three).

For Veeam Backup for AWS to use the archiving mechanism, there must be specified at least 2 different schedules: one schedule will control the regular creation of backup files, while another schedule will control the process of copying backup files to an archive backup repository. Backup chains created according to these two schedules will be completely different – for more information, see [RDS Backup Chain](#) and [Archive Backup Chain](#).

Consider the following example. You want a backup policy to create image-level backups of your critical workloads once a week, to keep the backed-up data in a standard backup repository for 3 weeks, and also to keep backups created once in 2 months in an archive backup repository for a year. In this case, you create 2 schedules when configuring the backup policy settings – weekly and monthly:


- In the policy target settings, you set the **Enable backups** toggle to *On*, select a backup repository that will store standard backup files, and select an archive backup repository that will store archived data.

The screenshot shows the 'Add RDS Policy' configuration page in the Veeam Backup for AWS console. The 'Targets' tab is selected, and the 'Backups' section is expanded. The 'Enable backups' toggle is turned on. A table of repositories is displayed, with 'v8-glacier-im...' selected as the archive repository.

Repository	Region	Storage Class	Folder	Description	Immuta...
rds-metadat...	Europe (London)	S3 Glacier Flex...	rds-meta...	Created by...	Disabled
v6-for-ami-I...	Europe (London)	S3 Glacier Flex...	v6-for-a...	Created by...	Disabled
v7-glacier-m...	Europe (London)	S3 Glacier Flex...	v7-glacie...	Created by...	Disabled
v8-glacier-im...	Europe (London)	S3 Glacier Flex...	v8-glacie...	Created by...	Enabled
v8-metadata...	Europe (London)	S3 Glacier Flex...	v8-metad...	Created by...	Disabled
v8-rds-meta...	Europe (London)	S3 Glacier Flex...	v8-rds-m...	Created by...	Disabled

- In the weekly scheduling settings, you select hours and days when backups will be created (for example, *7:00 AM, Monday*), and specify a number of days for which Veeam Backup for AWS will retain backups (for example, *21 days*).

Veeam Backup for AWS will propagate these settings to the archive schedule (which is the monthly schedule in our example).



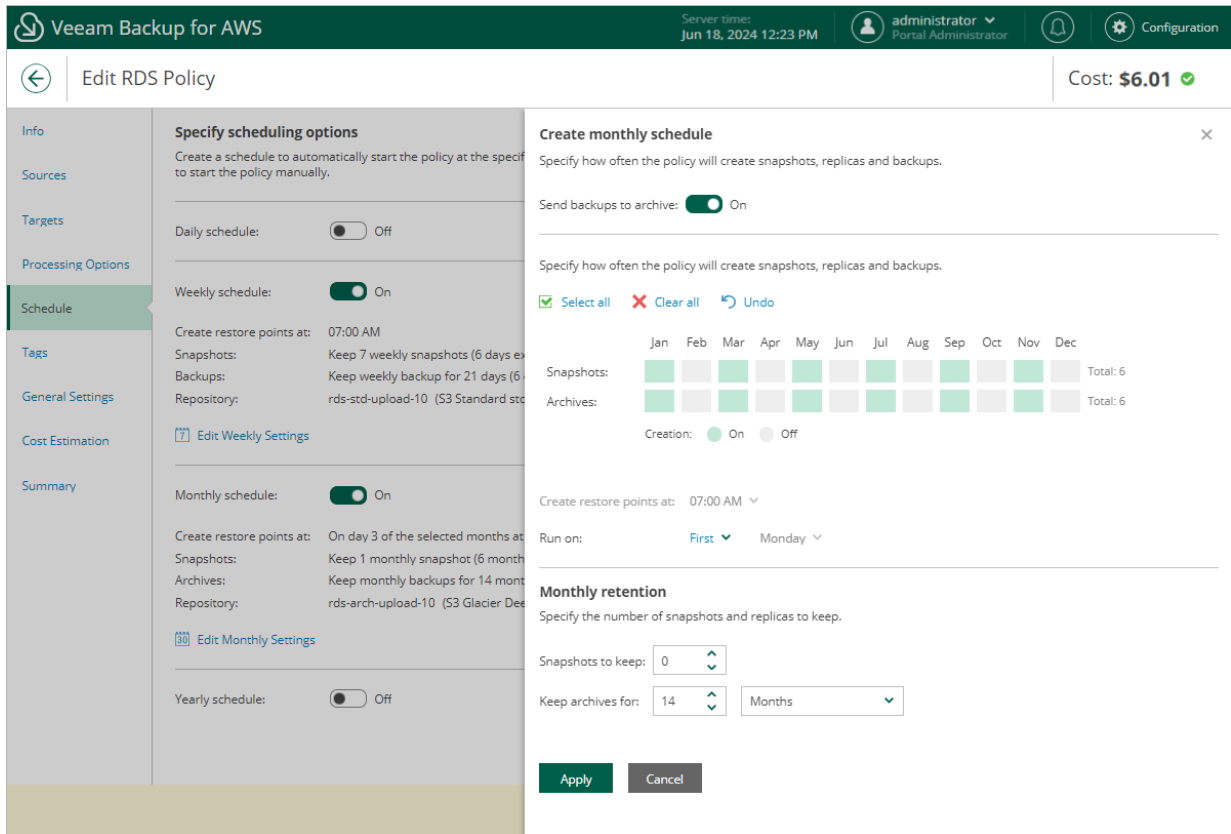
- In the monthly scheduling settings, you enable the archiving mechanism by setting the **Send backups to archive** toggle to *On*, specify when Veeam Backup for AWS will create archive backup files, and choose for how long you want to keep the created backups in the archive backup repository.

For example, *January, March, May, July, September, November, 12 months and First Monday*.

IMPORTANT

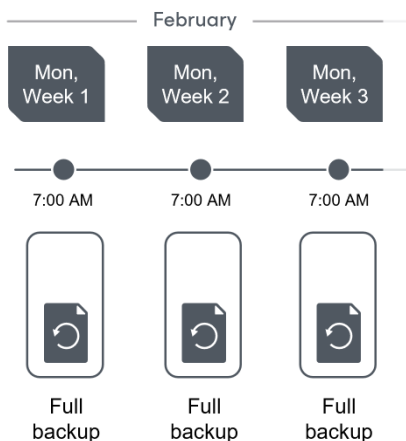
Consider the following:

- When you enable backup archiving, you become no longer able to create a schedule of the same frequency for standard backups. By design, these two functionalities are mutually exclusive.
- If you enable backup archiving, it is recommended that you set the **Snapshots to keep** value to *0*, to reduce unexpected snapshot charges.
- If you enable backup archiving, it is recommended that you set the **Keep archives for** value to at least *3 months* (or *90 days*) for the S3 Glacier Flexible Retrieval storage class and at least *6 months* (or *180 days*) for the S3 Glacier Deep Archive storage class. For more information on the minimum storage duration of the Amazon S3 archival storage classes, see [AWS Documentation](#).
- If you select the **On day** option, **harmonized scheduling** cannot be guaranteed. Plus, to support the **On day** option, Veeam Backup for AWS will require to create an additional temporary restore point if there are no other schedules planned to run on that day. However, the temporary restore point will be removed during the *Backup Retention* process from AWS in approximately 24 hours, to reduce unexpected infrastructure charges.



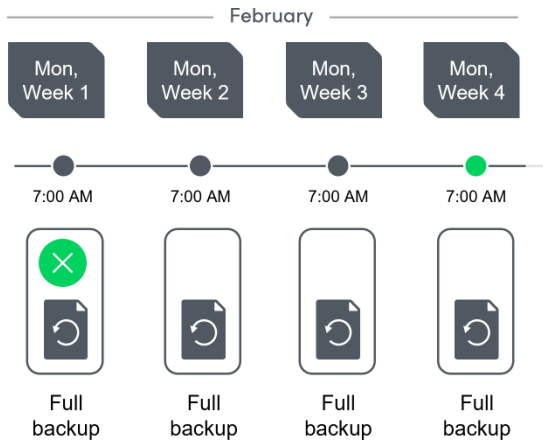
According to the specified scheduling settings, Veeam Backup for AWS will create image-level backups in the following way:

1. On the first Monday of February, a backup session will start at 7:00 AM to create the first restore point in the standard backup chain. Veeam Backup for AWS will store this restore point as a full backup file in the backup repository.
2. On the second and third Mondays of February, Veeam Backup for AWS will create restore points at 7:00 AM and add them to the standard backup chain as a full backup file in the backup repository.



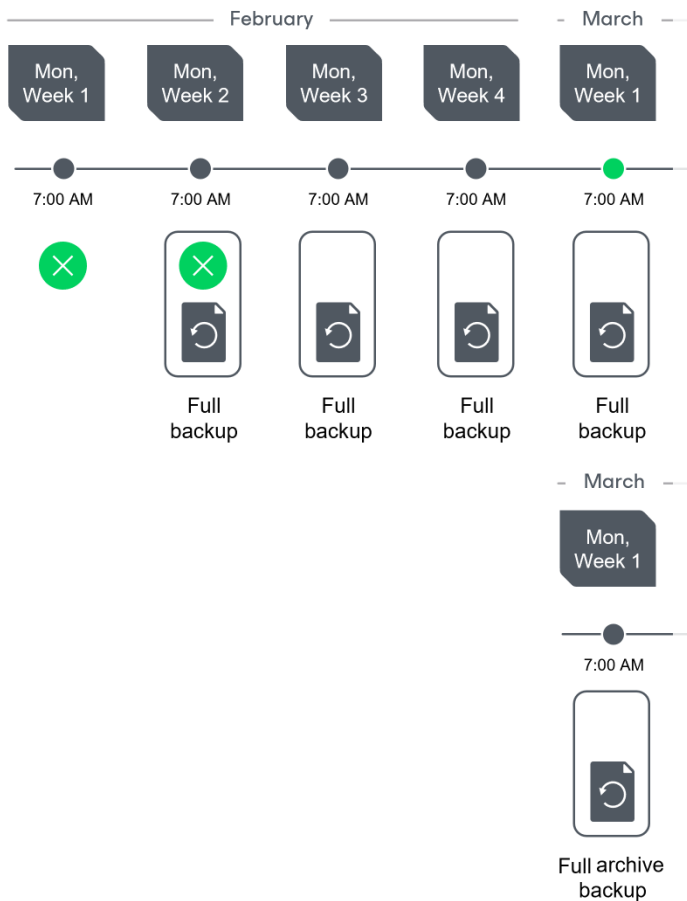
- On the fourth Monday of February, Veeam Backup for AWS will create a new restore point at 7:00 AM. By the moment the backup session completes, the earliest restore point in the standard backup chain will get older than the specified retention limit. That is why Veeam Backup for AWS will remove from the chain the restore point created on the first Monday.

For more information on how Veeam Backup for AWS transforms standard backup chains, see [RDS Backup Retention](#).



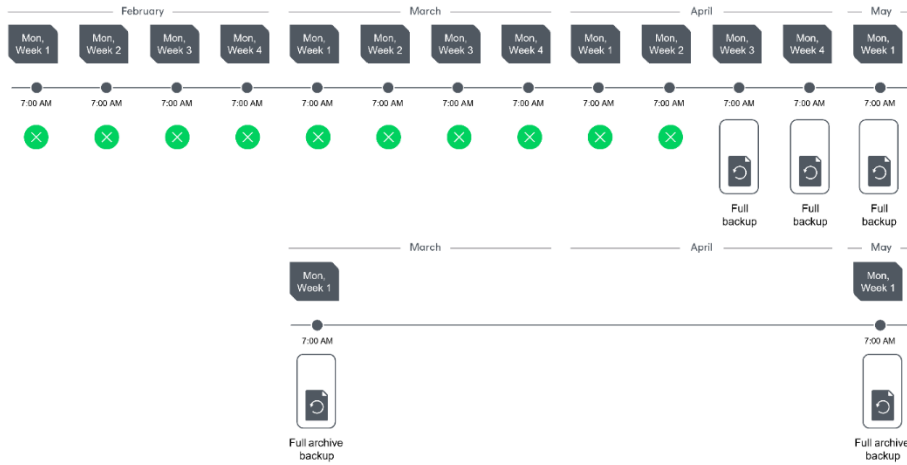
- On the first Monday of March, a backup session will start at 7:00 AM to create another restore point in the standard backup chain. At the same time, the earliest restore point in the standard backup chain will get older than the specified retention limit again. That is why Veeam Backup for AWS will remove from the chain the restore point created on the second Monday.

After the backup session completes, an archive session will create a restore point with all data from the standard backup chain. Veeam Backup for AWS will copy this restore point as a full archive backup file to the archive backup repository.



- Up to May, Veeam Backup for AWS will continue adding new restore points to the standard backup chain and deleting outdated backup files from the backup repository, according to the specified weekly scheduling settings.


On the first Monday of May, an archive session will create a restore point. Veeam Backup for AWS will copy this restore point as a full archive backup file to the archive backup repository.



- Up to the first Monday of March of the next year, Veeam Backup for AWS will continue adding new restore points to the standard backup chain and deleting outdated backup files from the backup repository, according to the specified weekly scheduling settings. Veeam Backup for AWS will also continue adding new restore points to the archive backup chain, according to the specified monthly settings.

By the moment the archive session completes, the earliest restore point in the archive backup chain will get older than the specified retention limit. That is why Veeam Backup for AWS remove from the chain the restore point created on the first Monday of March of the previous year.

For more information on how Veeam Backup for AWS transforms archive backup chains, see [Retention Policy for Archived Backups](#).



Step 7. Enable AWS Tags Assigning

At the **Tags** step of the wizard, you can instruct Veeam Backup for AWS to assign AWS tags to snapshots and snapshots replicas:

1. To assign already existing AWS tags from the processed RDS resources, select the **Copy tags from source RDS instances** check box.

If you choose to copy tags from the source instances, Veeam Backup for AWS will first create a cloud-native snapshot or snapshot replica of the DB instance or Aurora DB cluster and assign to the created snapshot AWS tags with Veeam metadata, then Veeam Backup for AWS will copy tags from the processed instance and finally assign the copied AWS tags to the snapshot.

2. To assign your own custom AWS tags, set the **Add custom tags to created snapshots** toggle to *On* and specify the AWS tags explicitly. To do that, use the **Key** and **Value** fields to specify a key and value for the new custom AWS tag, and then click **Add**. Note that you cannot add more than 5 custom AWS tags.

If you choose to add custom tags to the created snapshots, Veeam Backup for AWS will assign the specified tags right after it creates a cloud-native snapshot or snapshot replica.

The screenshot shows the 'Add RDS Policy' wizard in the Veeam Backup for AWS interface. The 'Tags' step is active, showing 'Specify tag settings'. The interface includes a top navigation bar with the Veeam logo, server time (Jun 18, 2024 12:16 PM), and user information (administrator, Portal Administrator). A sidebar on the left lists navigation options: Info, Sources, Targets, Processing Options, Schedule, Tags (selected), General Settings, Cost Estimation, and Summary. The main content area displays the 'Specify tag settings' section with a description: 'You can copy tags from source RDS instances and additionally assign up to 5 custom tags to snapshots created by the policy. Tags can help you manage, identify, organize, search for, and filter resources.' There are two main sections: 'Copy tags from source RDS instances' (checked) and 'Add custom tags to created snapshots' (toggle set to 'On'). Under the custom tags section, there are input fields for 'Key' (containing 'user') and 'Value' (containing 'donna_ortiz'), with an '+ Add' button. Below this, a tag 'owner: dept01' is shown with a close icon. A note states 'A maximum of 5 custom tags is allowed.' At the bottom, there are 'Previous', 'Next', and 'Cancel' buttons. A cost indicator in the top right shows 'Cost: \$33.98' with a green checkmark.

Step 8. Specify General Settings

At the **General Settings** step of the wizard, you can enable automatic retries and specify notification settings for the backup policy.

Automatic Retry Settings

To instruct Veeam Backup for AWS to run the backup policy again if it fails on the first try, do the following:

1. In the **Schedule** section of the step, select the **Automatically retry failed policy** check box.
2. In the field to the right of the check box, specify the maximum number of attempts to run the backup policy. The time interval between retries is 60 seconds.

When retrying backup policies, Veeam Backup for AWS processes only those instances that failed to be backed up during the previous attempt.

Health Check Settings

If you have enabled creation of image-level backups at [step 4](#) of the wizard, you can instruct Veeam Backup for AWS to periodically perform a health check for backup restore points created by the policy. During the health check, Veeam Backup for AWS performs an availability check for data blocks in the whole standard backup chain, and a cyclic redundancy check (CRC) for storage metadata to verify its integrity. The health check helps you ensure that the restore points are consistent and that you will be able to restore data using these restore points. For more information on the health check, see [How Health Check Works](#).

NOTE

During a health check, Veeam Backup for AWS does not verify archived restore points created by the policy.

To enable health checks for the backup policy, do the following:

1. In the **Health check** section of the step, set the **Enable health check** toggle to *On*.
2. Use the **Run on** drop-down lists to schedule a specific day for the health check to run.

NOTE

Veeam Backup for AWS performs the health check during the first policy session that runs on the day when the health check is scheduled. If another backup policy session runs on the same day, Veeam Backup for AWS will not perform the health check during that session. For example, if the backup policy is scheduled to run multiple times on Saturday, and the health check is also scheduled to run on Saturday, the health check will only be performed during the first policy session on Saturday.

Email Notification Settings

NOTE

To be able to specify email notification settings for the RDS Backup policy, you must configure [global notification settings](#) first.

To instruct Veeam Backup for AWS to send email notifications for the backup policy, do the following:

1. In the **Notifications** section of the step, set the **Enabled** toggle to *On*.

If you set the toggle to *Off*, Veeam Backup for AWS will send notifications according to the configured global notification settings.

2. In the **Email** field, specify an email address of a recipient.

Use a semicolon to separate multiple recipient addresses. Do not use spaces after semicolons between the specified email addresses.

3. Use the **Notify on** list to choose whether you want Veeam Backup for AWS to send email notifications in case the backup policy completes successfully, completes with warnings or completes with errors.

4. Select the **Suppress notifications until the last retry** check box to receive a notification about the final backup policy result.

If you do not select the check box, Veeam Backup for AWS will send a notification for every backup policy retry.

NOTE

If you specify the same email recipient in both backup policy notification and [global notification settings](#), Veeam Backup for AWS will override the configured global notification settings and will send each notification to this recipient only once to avoid notification duplicates.

The screenshot shows the 'Add RDS Policy' configuration page in Veeam Backup for AWS. The 'General Settings' tab is active. The 'Notifications' section is expanded, showing the following settings:

- Enabled:** On
- Email:** donna_ortiz@companymail.com
- Notify on:**
 - Failure
 - Warning
 - Success
- Suppress notifications until the last retry

Other visible settings include:

- Automatically retry failed policy:** 3 times
- Health check:** On
- Run on:** First of Sunday of every month
- Cost:** \$33.98

How Health Check Works

When Veeam Backup for AWS saves a new backup restore point to a backup repository, it calculates CRC values for metadata in the backup chain and saves these values to the chain metadata, together with the instance data. When performing a health check, Veeam Backup for AWS verifies the availability of data blocks and uses the saved values to ensure that the restore points being verified are consistent.

If you have enabled health checks for the backup policy, Veeam Backup for AWS performs the following operations at the day scheduled for a health check to run:

1. As soon as a backup policy session completes successfully, Veeam Backup for AWS starts the health check as a new session. For each restore point in the standard backup chain, Veeam Backup for AWS calculates CRC values for backup metadata and compares them to the CRC values that were previously saved to the restore point. Veeam Backup for AWS also checks whether data blocks that are required to rebuild the restore point are available.

If the backup policy session completes with an error, Veeam Backup for AWS tries to run the backup policy again, taking into account the maximum number of retries specified in the [automatic retry settings](#). After the first successful retry (or after the last one out of the maximum number of retries), Veeam Backup for AWS starts the health check.

2. If Veeam Backup for AWS does not detect data inconsistency, the health check session completes successfully. Otherwise, the session completes with an error.

Depending on the detected data inconsistency, Veeam Backup for AWS performs the following operations:

- If the health check detects corrupted metadata in a full or an incremental restore point, Veeam Backup for AWS marks the backup chain as corrupted in the configuration database. During the next backup policy session, Veeam Backup for AWS copies the full instance image, creates a full restore point in the backup repository and starts a new backup chain in the backup repository.

NOTE

Veeam Backup for AWS does not support metadata check for encrypted backup chains.

- If the health check detects corrupted disk blocks in a full or an incremental restore point, Veeam Backup for AWS marks the restore point that includes the corrupted data blocks and all subsequent affected incremental restore points as incomplete in the configuration database. During the next backup policy session, Veeam Backup for AWS copies not only those data blocks that have changed since the previous backup session but also data blocks that have been corrupted, and saves these data blocks to the latest restore point that has been created during the current session.

Step 9. Review Estimated Cost

[This step applies only if you have created a schedule for the backup policy at the **Schedule** step of the wizard]

At the **Cost Estimation** step of the wizard, review the estimated monthly cost of AWS services and resources that will be consumed to protect the instances added to the backup policy. The total estimated cost includes the following:

- The cost of creating and maintaining cloud-native snapshots of the instances.
For each instance included in the backup policy, Veeam Backup for AWS takes into account the number of restore points to be kept in the snapshot chain, and the configured scheduling settings.
- The cost of creating snapshot replicas and maintaining them in the target AWS Region.
For each instance included in the backup policy, Veeam Backup for AWS takes into account the number of restore points to be kept in the snapshot chain, and the configured scheduling settings.
- The cost of sending API requests to Veeam Backup for AWS during data protection operations.

To calculate the estimated cost, Veeam Backup for AWS uses capabilities of the [AWS Pricing Calculator](#) that estimates the cost of services in USD only. This calculator is intended for informational and estimation purposes only.

IMPORTANT

Veeam Backup for AWS does not support calculating the cost of creating and maintaining snapshots of Aurora DB clusters.

The estimated cost may occur to be significantly higher due to the snapshot charges. To reduce the cost, you can try to adjust the snapshot retention settings to keep less restore points in the snapshot chain. For more information on cost estimation, see [this Veeam KB article](#).

TIP

You can save the cost estimation as a .CSV or .XML file. To do that, click **Export to** and select the necessary format.

Review cost estimation

The estimated cost takes into account the configured target settings, the specified scheduling options and the number of resources to protect.

Note that Veeam Backup for AWS makes predefined assumptions to calculate the cost, which means that the results should be used only as an approximation. For more information on cost calculation, see [this Veeam KB article](#).

\$17.77 Snapshots	\$14.46 Replicas	\$1.75 Traffic	\$0.00 Transactions
-----------------------------	----------------------------	--------------------------	-------------------------------

Estimated monthly cost: \$33.98

Instance ↑	Snapshot	Replica	Traffic	Transact...	Total
bd-paris-rds-po...	\$8.89	\$7.23	\$0.88	\$0.00	\$16.99
pl-ireland-psql	\$8.89	\$7.23	\$0.88	\$0.00	\$16.99

Related Resources

[How AWS Pricing Works](#)

Step 10. Finish Working with Wizard

At the **Summary** step of the wizard, it is recommended that you run the backup policy configuration check before you click **Finish**.

The configuration check will verify whether specified IAM roles have all the required permissions. To run the check, click **Test Configuration**. Veeam Backup for AWS will display the **Test policy configuration** window where you can track the progress and view the results of the check. If some permissions of the IAM role are missing or policy settings are not configured properly, the check will complete with errors. You can grant the missing permissions to the IAM role as described in section [Checking IAM Role Permissions](#).

After the required permissions are granted, close the **Test policy configuration** window, and then click **Finish** to close the **Add Policy** wizard.

Veeam Backup for AWS will save the configured backup policy.

The screenshot shows the Veeam Backup for AWS interface. The main window is titled 'Add RDS Policy' and shows a 'Summary' tab. A 'Test policy configuration' window is open, displaying a table of test results. The table has columns for 'Type', 'Status', 'Action', and 'Result'. The results show four successful checks: 'Checking policy configuration...', 'Checking policy role permissions', 'Checking replication role permissions', and 'Session finished at 06/18/2024 12:18:25 PM.'. A 'Recheck' button is visible above the table. The 'Close' button is at the bottom of the window.

Type	Status	Action	Result
Checking policy configuration...	Success	—	—
Checking policy role permissions	Success	—	—
Checking replication role permissions	Success	—	—
Session finished at 06/18/2024 12:18:25 PM.	Success	—	—

Creating RDS Snapshots Manually

Veeam Backup for AWS allows you to manually create snapshots of RDS resources. You can instruct Veeam Backup for AWS to store the created snapshots in the same AWS Regions where the processed DB instances and DB clusters reside, or in a different AWS Region or AWS account.

NOTE

Veeam Backup for AWS does not include snapshots created manually in the snapshot chain and does not apply the configured retention policy settings to these snapshots. This means that the snapshots are kept in your AWS environment unless you remove them manually, as described in section [Managing Backed-Up RDS Data](#).

To manually create a cloud-native snapshot of a DB instance or an Aurora DB cluster, do the following:

1. Navigate to **Resources > Databases > RDS**.
2. Select the necessary instance and click **Take Snapshot Now**.

For an RDS resource to be displayed in the list of available instances, an AWS Region where the instance resides must be added to any of [configured RDS backup policies](#), and the IAM role specified in the backup policy settings must have permissions to access the instance. For more information on required permissions, see [RDS Backup IAM Role Permissions](#).

3. Complete the **Take Manual Snapshot** wizard:

- a. At the **Account** step of the wizard, specify an IAM role whose permissions Veeam Backup for AWS will use to create the snapshot.

For an IAM role to be displayed in the list, it must be added to Veeam Backup for AWS as described in section [Adding IAM Roles](#).

- b. At the **Snapshot Mode** step of the wizard, choose whether you want to store the snapshot in the same AWS Region where the processed RDS resource resides, or in another AWS Region or AWS account.
- c. [Applies if you have selected the **New location** option] At the **Settings** step of the wizard, choose an IAM role whose permissions will be used to copy and store the snapshot in the target AWS Region and specify whether to encrypt the copied snapshot.
- d. At the **Tags** step of the wizard, choose whether you want to assign AWS tags to the created snapshot.

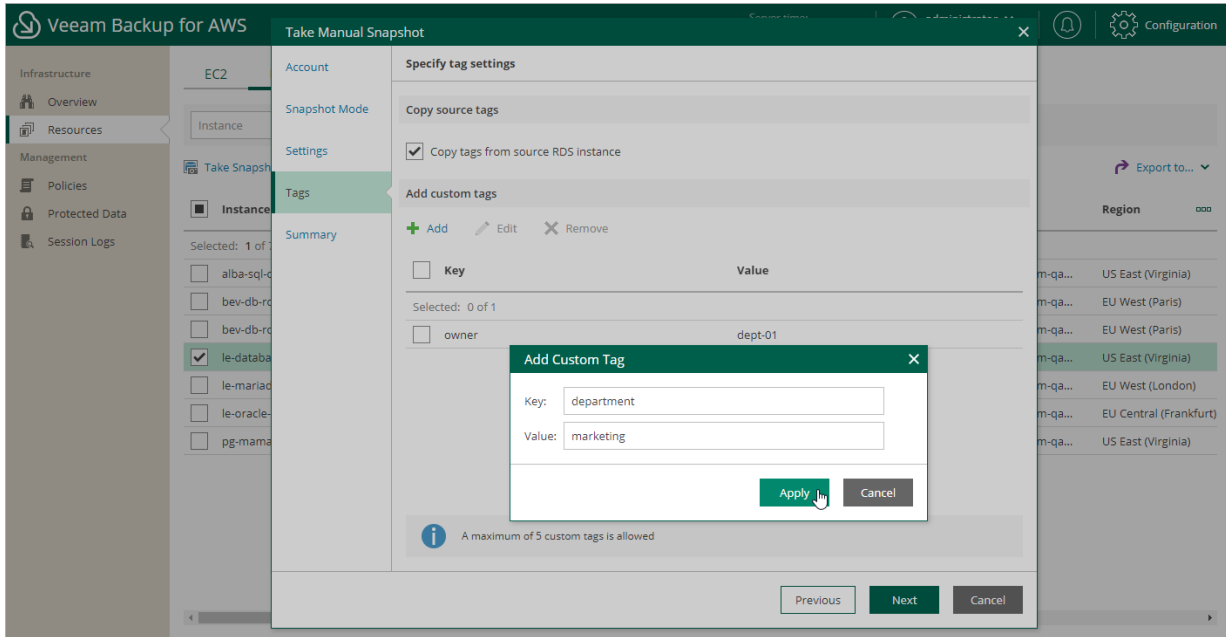
- To assign already existing AWS tags from the source DB instance and Aurora DB cluster, select the **Copy tags from source RDS instance** check box.

If you choose to copy tags from the source RDS resource, Veeam Backup for AWS will first create a snapshot of the DB instance or Aurora DB cluster and assign to the created snapshot AWS tags with Veeam metadata. Then, Veeam Backup for AWS will copy tags from the processed resource and assign the copied AWS tags to the snapshot.

- To assign your own custom AWS tags, click **Add** and specify the tags explicitly. To do that, in the **Add Custom Tag** window, specify a key and a value for the new AWS tag, and then click **Apply**. Note that you cannot add more than 5 custom AWS tags.

If you choose to add custom tags to the created snapshots, Veeam Backup for AWS will assign the specified tags right after it creates a snapshot.

e. At the **Summary** step of the wizard, review summary information and click **Finish**.



Performing DynamoDB Backup

One backup policy can be used to process one or more DynamoDB tables within one AWS account. The scope of data that you can protect in an AWS account is limited by permissions of an IAM role that is specified in the backup policy settings.

NOTE

If you plan to receive email notifications on backup policy results, configure email notification settings before creating a DynamoDB backup policy. For more information, see [Configuring Global Notification Settings](#).

To schedule data protection tasks to run automatically, [create backup policies](#). For each protected DynamoDB table, you can also [take a backup manually](#) when needed.

IMPORTANT

Consider the following:

- You can back up DynamoDB tables only to the same AWS accounts where the source tables belong.
- You can back up only those DynamoDB table properties that are described in section [Protecting DynamoDB Tables](#).

Creating DynamoDB Backup Policies


To create a backup policy, do the following:

1. [Launch the Add DynamoDB Policy wizard](#).
2. [Specify a backup policy name and description](#).
3. [Configure backup source settings](#).
4. [Configure backup target settings](#).
5. [Specify a schedule for the backup policy](#).
6. [Enable AWS tags assigning](#).
7. [Specify automatic retry settings and notification settings for the backup policy](#).
8. [Review estimated cost of the selected DynamoDB tables](#).
9. [Finish working with the wizard](#).

Step 1. Launch Add DynamoDB Policy Wizard

To launch the **Add DynamoDB Policy** wizard, do the following:

1. Navigate to **Policies > Databases > DynamoDB**.
2. Click **Add**.



Step 2. Specify Policy Name and Description

At the **Info** step of the wizard, use the **Name** and **Description** fields to specify a name for the new backup policy and to provide a description for future reference. The name must be unique in Veeam Backup for AWS; the maximum length of the name is 127 characters, the maximum length of the description is 255 characters.

The screenshot shows the 'Add DynamoDB Policy' wizard in the Veeam Backup for AWS console. The interface is in the 'Info' step, where the user is prompted to 'Specify policy name and description'. The 'Name' field contains 'DynamoDB backup policy 02' and the 'Description' field contains 'Created by administrator at 6/19/2024 12:14 AM'. The 'Cost' is shown as 'N/A' with a warning icon. The 'Next' button is highlighted in green, and the 'Cancel' button is greyed out. The left sidebar shows the navigation menu with 'Info' selected. The top header includes the Veeam logo, server time (Jun 19, 2024 12:15 PM), and user information (administrator, Portal Administrator).

Veeam Backup for AWS

Server time: Jun 19, 2024 12:15 PM

administrator Portal Administrator

Configuration

← Add DynamoDB Policy Cost: N/A ⚠

Info

Sources

Targets

Schedule

Tags

General Settings

Cost Estimation

Summary

Specify policy name and description

Enter a name and description for the policy.

Name:

DynamoDB backup policy 02

Description:

Created by administrator at 6/19/2024 12:14 AM

Next Cancel

Step 3. Configure Backup Source Settings

At the **Sources** step of the wizard, specify backup source settings:

1. [Select an IAM role whose permissions will be used to perform DynamoDB backup.](#)
2. [Select AWS Regions where DynamoDB tables that you plan to back up reside.](#)
3. [Select DynamoDB tables to back up.](#)

Step 3.1 Specify IAM Role

In the **IAM role** section of the **Sources** step of the wizard, specify an IAM role whose permissions will be used to access AWS services and resources, and to create backups of DynamoDB tables. The specified IAM role must belong to the AWS account in which the DynamoDB tables that you want to protect reside, and must be assigned the permissions listed in section [DynamoDB Backup IAM Role Permissions](#).

For an IAM role to be displayed in the **IAM role** list, it must be added to Veeam Backup for AWS with the *Amazon DynamoDB Backup* operation selected for the role as described in section [Adding IAM Roles](#). If you have not added the necessary IAM role to Veeam Backup for AWS beforehand, you can do it without closing the **Add DynamoDB Policy** wizard. To do that, click **Add** and complete the [Add IAM Role](#) wizard.

IMPORTANT

It is recommended that you check whether the selected IAM role has all the permissions required to perform the operation. If some permissions of the IAM role are missing, the backup policy will fail to complete successfully. To run the IAM role permission check, click **Check Permissions** and follow the instructions provided in section [Checking IAM Role Permissions](#).

The screenshot shows the 'Add DynamoDB Policy' wizard in the Veeam Backup for AWS interface. The top navigation bar includes the Veeam logo, 'Veeam Backup for AWS', server time 'Jun 19, 2024 12:16 PM', user 'administrator Portal Administrator', and a 'Configuration' link. The main title is 'Add DynamoDB Policy' with a 'Cost: N/A' indicator. A left sidebar lists steps: Info, Sources (selected), Targets, Schedule, Tags, General Settings, Cost Estimation, and Summary. The main content area is titled 'Specify source settings' and contains instructions: 'Select an IAM role to use, regions to cover and resources to process by the policy. Using tags provides dynamic selection that automatically changes the policy scope when tags are assigned to tables.' Below this are three sections: 'IAM role' with a dropdown menu showing 'dynamo-backup-487569979969 (Created by nm at ...)' and buttons for '+ Add' and 'Check Permissions'; 'Regions' with the instruction 'Specify one or more regions.' and a 'Choose regions...' button; and 'Resources' with the instruction 'Specify resources to protect or exclude.' and buttons for 'Choose resources to protect...' and 'Choose resources to exclude...'. At the bottom, there are 'Previous', 'Next', and 'Cancel' buttons.

Step 3.2 Select AWS Regions

In the **Regions** section of the **Sources** step of the wizard, choose AWS Regions where DynamoDB tables that you plan to back up reside:

1. Click **Choose regions**.
2. In the **Choose regions** window, select the necessary regions, and click **Add** to include them in the backup policy.
3. To save changes made to the backup policy settings, click **Apply**.

The screenshot shows the Veeam Backup for AWS interface. The main window is titled "Add DynamoDB Policy" and has a "Cost: N/A" indicator. On the left, a navigation pane shows "Sources" selected. The main content area is divided into "Specify source settings" and "Choose regions".

Specify source settings:

- IAM role:** dynamo-backup-487569979969 (Create)
- Regions:** Specify one or more regions. 1 region selected.
- Resources:** Specify resources to protect or exclude. 1 resource will be protected. Choose resources to exclude...

Choose regions dialog box:

- Available Regions (24):** Asia Pacific (Seoul), Asia Pacific (Singapore), Asia Pacific (Sydney), Asia Pacific (Tokyo), Canada (Central), Canada West (Calgary), Europe (Frankfurt), Europe (Ireland), Europe (London), Europe (Milan), Europe (Paris) (highlighted), Europe (Spain), Europe (Stockholm), Israel (Tel Aviv).
- Buttons:** Add, Remove.
- Selected Regions (0):** No data.
- Footer:** Apply, Cancel.

Step 3.3 Select DynamoDB Tables

In the **Resources** section of the **Sources** step of the wizard, specify the backup scope – select DynamoDB tables that Veeam Backup for AWS will back up:

1. Click **Choose resources to protect**.
2. In the **Choose resources to protect** window, choose whether you want to back up all DynamoDB tables from AWS Regions selected at [step 3.2](#) of the wizard, or only specific DynamoDB tables.

If you select the **All resources** option, Veeam Backup for AWS will regularly check for new DynamoDB tables launched in the selected regions and automatically update the backup policy settings to include these tables into the backup scope.

If you select the **Protect only following resources** option, you must also specify the tables explicitly:

- a. Use the **Type** drop-down list to choose whether you want to add individual DynamoDB tables or AWS tags to the backup scope.

If you select the **Tag** option, Veeam Backup for AWS will back up only those DynamoDB tables from the selected AWS Regions that are assigned specific tags.

- b. Use the search field to the right of the **Type** list to find the necessary resource, and then click **Protect** to add the resource to the backup scope.

For a resource to be displayed in the list of available resources, it must reside in an AWS Region that has ever been specified in any backup policy. Otherwise, the only option to discover the available resources is to click **Browse to select specific resources from the global list** and to wait for Veeam Backup for AWS to populate the resource list.

TIP

You can simultaneously add multiple resources to the backup scope. To do that, click **Browse to select specific sources from the global list**, select check boxes next to the necessary DynamoDB tables or AWS tags in the list of available resources, and then click **Protect**.

If the list does not show the resources that you want to back up, click **Rescan** to launch the data collection process. As soon as the process is over, Veeam Backup for AWS will update the resource list.

If you add an AWS tag to the backup scope, Veeam Backup for AWS will regularly check for new DynamoDB tables assigned the added AWS tag and automatically update the backup policy settings to include these resources in the scope. However, this applies only to DynamoDB tables from the AWS Regions selected at [step 3.2](#) of the wizard. If you select an AWS tag assigned to DynamoDB tables from other AWS Regions, these tables will not be protected by the backup policy. To work around the issue, either go back to [step 3.2](#) and add the missing AWS Regions, or create a new backup policy.

3. To save changes made to the backup policy settings, click **Apply**.

TIP

As an alternative to selecting the **Protect only following resources** option and specifying the resources explicitly, you can select the **All resources** option and exclude a number of resources from the backup scope. To do that, click **Choose resources to exclude** and specify the tables or tags that you do not want to protect – the procedure is the same as described for including resources in the backup scope.

Note that if a resource appears both in the list of included and excluded resources, Veeam Backup for AWS will still not process the resource because the list of excluded resources has a higher priority.

The screenshot shows the Veeam Backup for AWS interface. The main window is titled "Add DynamoDB Policy" and has a "Cost: N/A" indicator. The left sidebar contains navigation options: Info, Sources, Targets, Schedule, Tags, General Settings, Cost Estimation, and Summary. The "Sources" section is active, showing "Specify source settings" with an IAM role of "dynamo-backup-487569979969 (Created)". The "Regions" section shows "1 region selected". The "Resources" section shows "1 resource will be protected" and a link to "Choose resources to exclude...".

The "Choose resources to protect" dialog is open, showing the following options:

- All resources
- Protect only following resources

The "Type" is set to "Table" and the "Name" is empty. A "Protect" button is visible. Below this, there is a link to "Browse to select specific resources from the global list...".

The "Protected resources (1)" section shows a search bar and a "Remove" button. Below this is a table with the following data:

Item ↑	ID	Value	Region
Selected: 0 of 1			
DataTable	3ba66e2c-7e72-408c-908...	—	Europe (Paris)

At the bottom of the dialog are "Apply" and "Cancel" buttons.

Step 4. Configure Backup Target Settings

By default, backup policies create only backups of processed DynamoDB tables. At the **Targets** step of the wizard, you can specify the following backup target settings:

- [Specify backup vaults where Veeam Backup for AWS will store DynamoDB backups.](#)
- [Instruct Veeam Backup for AWS to copy DynamoDB backups to other AWS Regions.](#)
- [Instruct Veeam Backup for AWS to store DynamoDB backups in a cold storage tier.](#)

Configuring Backup Settings

To specify backup vaults that will be used to store backups of the selected DynamoDB tables, do the following:

1. In the **Backups** section of the **Targets** step of the wizard, click **Choose backup vaults**.
2. In the **Choose backup vaults** window, for each AWS Region included in the policy, specify a backup vault to save and organize table backups. To do that:
 - a. Select an AWS Region and click **Edit**.
 - b. In the **Edit Backup Vault** window, from the **Backup vault** drop-down list, select the necessary backup vault.

For a backup vault to be displayed in the **Backup vault** list, it must be created in the AWS Backup console as described in [AWS Documentation](#). If no custom backup vaults exist in the selected AWS Region, the list will contain the default backup vault only.


IMPORTANT

Consider the following:

- Veeam Backup for AWS does not support storing backups in [logically air gapped vaults](#).
- Make sure policies assigned to the selected backup vault allow Veeam Backup for AWS to access vault resources and to perform backup, backup copy and restore operations. For more information on vault access policies, see [AWS Documentation](#).
- For Veeam Backup for AWS to be able to back up DynamoDB tables, you must configure the AWS Backup settings to enable both the Opt-in service and the advanced features for Amazon DynamoDB backups. Otherwise, Veeam Backup for AWS will automatically enable these settings for each AWS Region specified in the **Backups** section in your AWS account while performing backup operations. For more information on advanced DynamoDB backup, see [AWS Documentation](#).

- c. Click **Save**.

3. To save changes made to the backup policy settings, click **Apply**.



Enabling Additional Backup Copy

If you want to copy DynamoDB backups to other AWS Regions, do the following:

1. In the **Backup copies** section of the **Targets** step of the wizard, set the **Create backup copies** toggle to *On*.
2. In the **Choose backup vaults** window, configure the following mapping settings for each AWS Region where original tables reside:
 - a. Select a source AWS Region in the list and click **Edit Region Mapping**.
 - b. In the **Edit Region Mapping** window, specify the following settings:
 - i. From the **Target region** drop-down list, select the target AWS Region to which Veeam Backup for AWS must copy created backups of the selected tables.
 - ii. From the **Backup vault** drop-down list, select a backup vault that will be used to store the copied backups.

For a backup vault to be displayed in the **Backup vault** list, it must be created in the AWS Backup console as described in [AWS Documentation](#). If you have not created a backup vault for the selected AWS Region, Veeam Backup for AWS will display only the default backup vault existing in this region.

IMPORTANT


Consider the following:

- Veeam Backup for AWS does not support storing backups in [logically air gapped vaults](#).
- Make sure policies assigned to the selected backup vault allow Veeam Backup for AWS to access vault resources and to perform backup, backup copy and restore operations. For more information on vault access policies, see [AWS Documentation](#).
- For Veeam Backup for AWS to be able to back up DynamoDB tables, you must configure the AWS Backup settings to enable both the Opt-in service and the advanced features for Amazon DynamoDB backups. Otherwise, Veeam Backup for AWS will automatically enable these settings for each AWS Region specified in the **Backup copies** section in your AWS account while performing backup operations. For more information on advanced DynamoDB backup, see [AWS Documentation](#).

iii. Click **Save**.

To configure mapping for all source AWS Regions at once, click **Set Mapping for All Regions** and specify settings as described in [step 2.b](#).

c. To save changes made to the backup policy settings, click **Apply**.



Configuring Archive Settings

If you want to reduce the cost of storing backups that you plan to access infrequently, you can instruct Veeam Backup for AWS to move backups from a high-available warm storage tier to a low-cost cold storage tier:

1. In the **Archive** section of the **Targets** step of the wizard, set the **Enable cold storage** toggle to *On*.
Note that after you enable the archiving mechanism, you must configure the [retention policy settings](#).
2. In the **Move backups after** field, specify the number of days for which you want to keep backups in a warm storage tier before moving them to a cold storage tier (the minimum value is 1; the maximum value is 36,135). As soon as the specified period is over, the backups will be moved to the cold storage tier and will be stored there according to the configured retention policy settings.

Keep in mind that once moved to a cold storage tier in an AWS Region, backups can only be used to restore tables to the same AWS Region. For more information, see [DynamoDB Restore](#).

IMPORTANT

Consider the following:

- It is recommended that you keep backups in a cold storage tier for at least 90 days since there is a [limitation on the AWS Backup service side](#) – it will still charge you for 90 days even if your backups are stored for less than 90 days.
- The configured archive settings apply to all restore points (both backups and backup copies) that will be created by this backup policy.

Veeam Backup for AWS

Server time: Jun 19, 2024 12:28 PM

administrator Portal Administrator

Configuration

← Add DynamoDB Policy Cost: N/A ⚠

Info

Sources

Targets

Schedule

Tags

General Settings

Cost Estimation

Summary

Specify target settings

Specify a location to store backups created by the policy, and choose whether you want to enable backup copying and archiving.

Backups

Configure backup settings.

Backup vault for 1 region is selected

Backup copies

Enable backup copies to create an additional backup in another vault or region.

Create backup copies: On

Backup vault for 1 region is selected

Archive

Specify if you want to use cold storage for backups and backup copies. Backups transitioned to cold storage must be stored in cold storage for a minimum of 90 days.

Enable cold storage: On

Move backups after: 30

i Your backups will be stored in warm tier for 30 days. The minimum optimal recommended retention value for your archived backups is 120. For more information, see the User Guide.

Previous Next Cancel

Step 5. Specify Policy Scheduling Options

You can instruct Veeam Backup for AWS to start the backup policy automatically according to a specific backup schedule. The backup schedule defines how often data of the tables added to the backup policy must be backed up.

IMPORTANT

If you have instructed Veeam Backup for AWS to move backups to the cold storage tier at [step 4](#) of the wizard, you must configure at least one schedule for the backup policy.

To help you implement a comprehensive backup strategy, Veeam Backup for AWS allows you to create schedules of the following types:

- **Daily** – the backup policy will create restore points repeatedly throughout a day on specific days.
- **Weekly** – the backup policy will create restore points once a day on specific days.
- **Monthly** – the backup policy will create restore points once a month on a specific day.
- **Yearly** – the backup policy will create restore points once a year on a specific day.

Combining multiple schedule types together allows you to retain restore points for longer periods of time. For more information, see [Enabling Harmonized Scheduling](#).

NOTE

If you do not specify the backup schedule, after you configure the backup policy, you will need to start it manually to create DynamoDB table backups. To learn how to start backup policies, see [Starting and Stopping Policies](#).

Specifying Daily Schedule

To create a daily schedule for the backup policy, at the **Schedule** step of the wizard, do the following:

1. Set the **Daily schedule** toggle to *On* and click **Edit Daily Settings**.
2. In the **Create daily schedule** window, select hours when the backup policy must create table backups and backup copies.

If you want to protect table data more frequently, you can instruct the backup policy to create multiple backups per hour. To do that, click the link to the right of the **Backups** hour selection area, and specify the number of backups that the backup policy must create within an hour.

NOTE


Veeam Backup for AWS does not create backup copies independently from table backups. That is why when you select hours for backup copies, the same hours are automatically selected for backups. To learn how Veeam Backup for AWS performs backup, see [DynamoDB Backup](#).

3. Use the **Run at** drop-down list to choose whether you want the backup policy to run everyday, on work days (Monday through Friday) or on specific days.

- In the **Daily retention** section, configure retention policy settings for the daily schedule. For backups and backup copies, specify the number of days (or months) for which you want to keep restore points in a backup chain.

If a restore point is older than the specified time limit, Veeam Backup for AWS removes the restore point from the chain. For more information, see [DynamoDB Backup Retention](#).

- To save changes made to the backup policy settings, click **Apply**.



Specifying Weekly Schedule

To create a weekly schedule for the backup policy, at the **Schedule** step of the wizard, do the following:

- Set the **Weekly schedule** toggle to *On* and click **Edit Weekly Settings**.
- In the **Create weekly schedule** window, select weekdays when the backup policy must create table backups and backup copies.


NOTE

Veeam Backup for AWS does not create backup copies independently from table backups. That is why when you select days to create backup copies, the same days are automatically selected for backups. To learn how Veeam Backup for AWS performs backup, see [DynamoDB Backup](#).

- Use the **Create restore point at** drop-down list to schedule a specific time for the backup policy to run.
- In the **Weekly retention** section, configure retention policy settings for the weekly schedule. For backups and backup copies, specify the number of days (or months) for which you want to keep restore points in a backup chain.

If a restore point is older than the specified time limit, Veeam Backup for AWS removes the restore point from the chain. For more information, see [DynamoDB Backup Retention](#).

5. To save changes made to the backup policy settings, click **Apply**.



Specifying Monthly Schedule

To create a monthly schedule for the backup policy, at the **Schedule** step of the wizard, do the following:

1. Set the **Monthly schedule** toggle to *On* and click **Edit Monthly Settings**.
2. In the **Create monthly schedule** window, select months when the backup policy must create table backups and backup copies.

NOTE

Veeam Backup for AWS does not create backup copies independently from table backups. That is why when you select hours for backup copies, the same hours are automatically selected for backups. To learn how Veeam Backup for AWS performs backup, see [DynamoDB Backup](#).

3. Use the **Create restore point at** and **Run on** drop-down lists to schedule a specific time and day for the backup policy to run.

NOTE

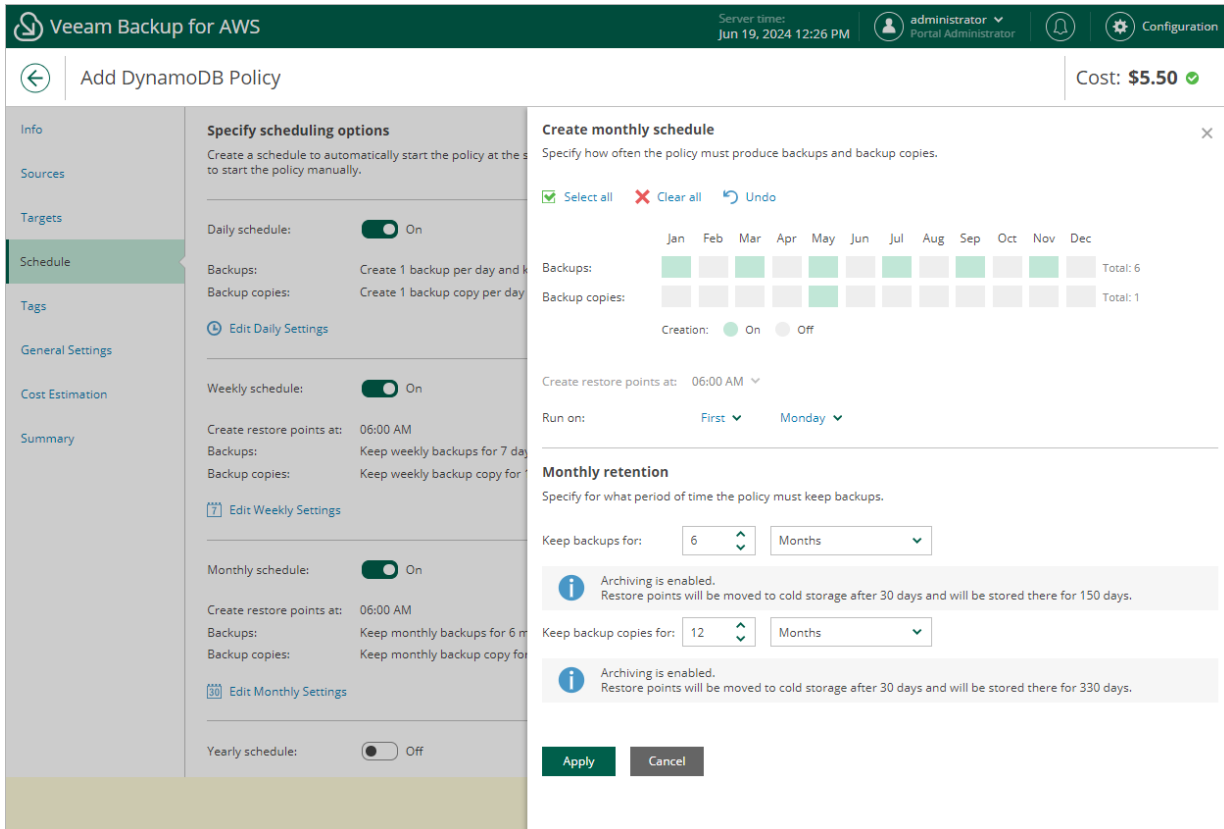
Consider the following:

- If you have selected a specific time for the backup policy to run at the **Weekly schedule** section of the **Schedule** step of the wizard, you will not be able to change the time for the monthly schedule unless you select the *On Day* option from the **Run on** drop-down list.
- If you select the **On day** option, [harmonized scheduling](#) cannot be guaranteed.

- In the **Monthly retention** section, configure retention policy settings for the monthly schedule. For backups and backup copies, specify the number of days (or months) for which you want to keep restore points in a backup chain.

If a restore point is older than the specified time limit, Veeam Backup for AWS removes the restore point from the chain. For more information, see [DynamoDB Backup Retention](#).

- To save changes made to the backup policy settings, click **Apply**.



Specifying Yearly Schedule

The yearly schedule is applied only to DynamoDB backups, no backup copies are created according to this schedule.

To create a yearly schedule for the backup policy, at the **Schedule** step of the wizard, do the following:

- Set the **Yearly schedule** toggle to *On* and click **Edit Yearly Settings**.
- In the **Create yearly schedule** window, specify a day, month and time when the backup policy must create table backups.

For example, if you select *First, Friday, January* and *06:00 PM*, the backup policy will run every first Friday of January at 06:00 PM.

NOTE

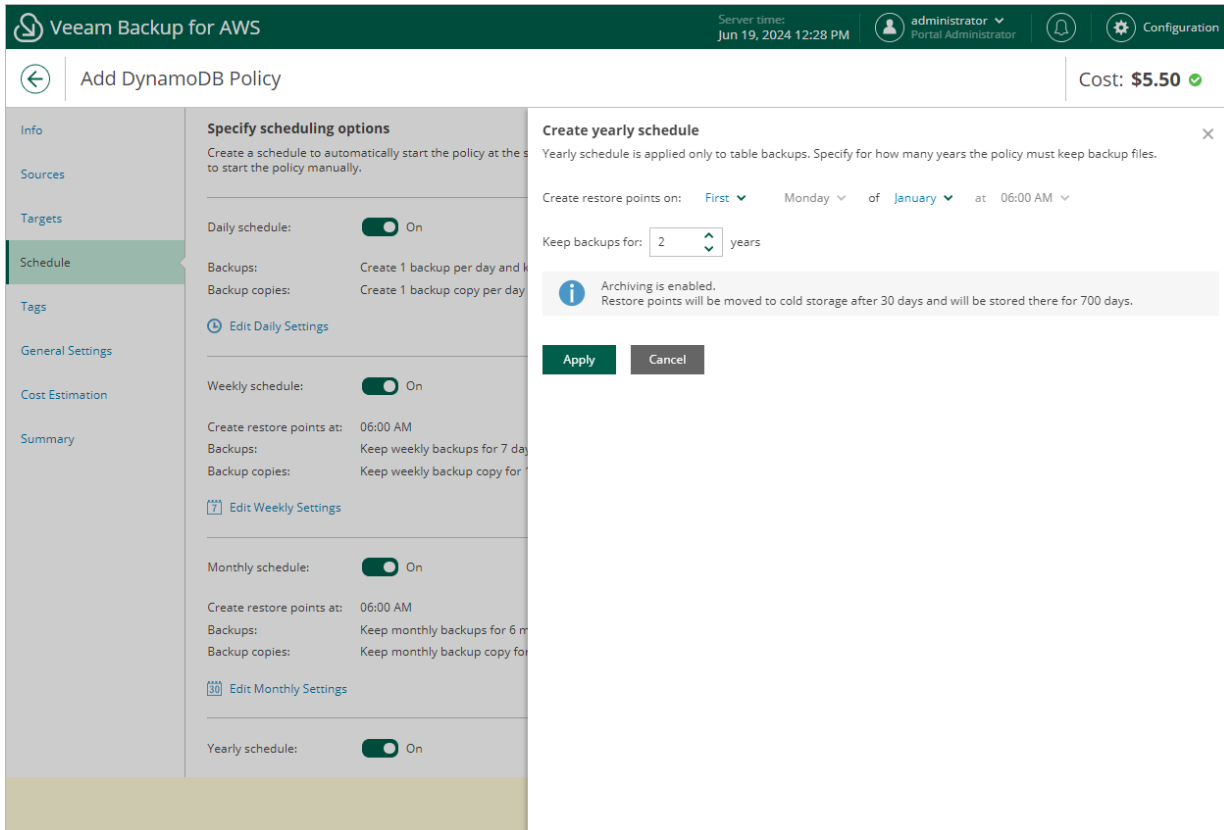
Consider the following:

- If you have selected a specific time and day for the backup policy to run at the **Weekly schedule** or **Monthly schedule** sections of the **Schedule** step of the wizard, you will not be able to change the time and day for the yearly schedule unless you select the *On Day* option from the **Create restore point on** drop-down list.
- If you select the *On day* option, [harmonized scheduling](#) cannot be guaranteed.

3. In the **Keep backups for** field, specify the number of years for which you want to keep restore points in a backup chain.

If a restore point is older than the specified time limit, Veeam Backup for AWS removes the restore from the chain. For more information, see [DynamoDB Backup Retention](#).

4. To save changes made to the backup policy settings, click **Apply**.



Enabling Harmonized Scheduling

When you combine multiple types of schedules, Veeam Backup for AWS applies the harmonization mechanism that allows you to leverage restore points for long-term retentions instead of taking a new restore point every time. The mechanism simplifies the backup schedule, optimizes the backup performance and reduces the cost of retaining restore points.

With harmonized scheduling, Veeam Backup for AWS can keep restore points created according to a daily, weekly or monthly schedule for longer periods of time: DynamoDB backups and backup copies can be kept for weeks, months and years.

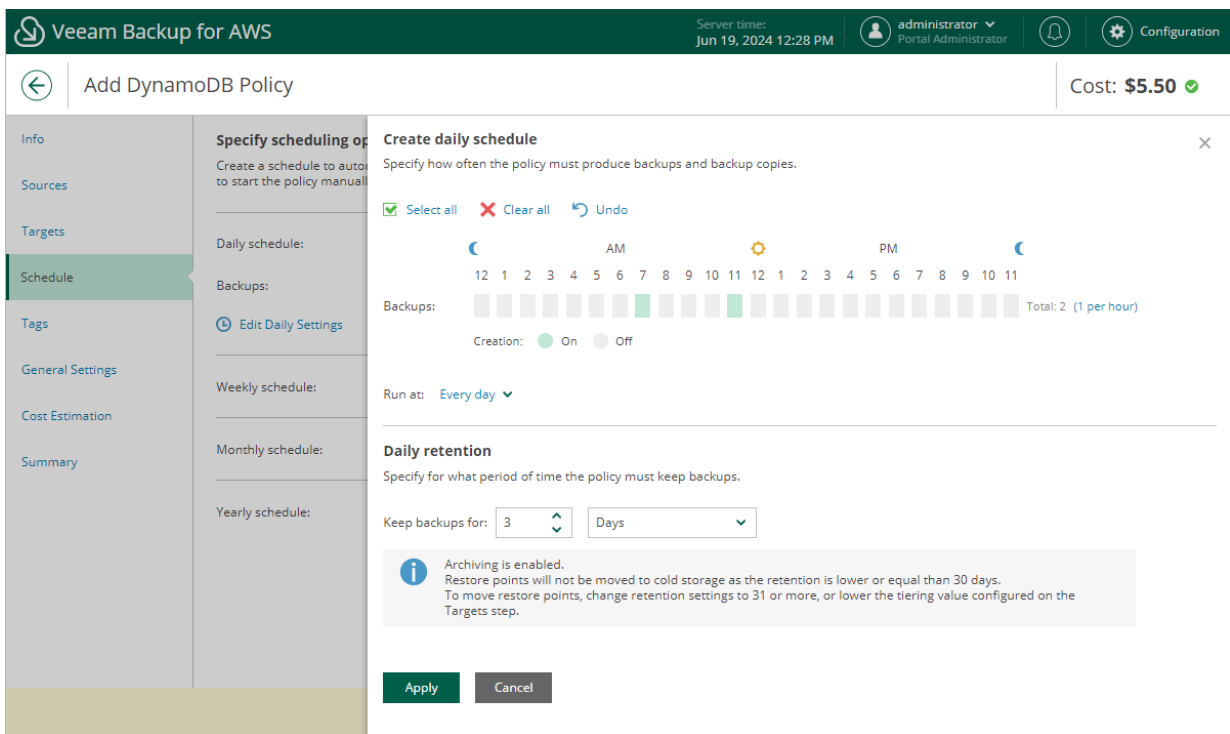
For Veeam Backup for AWS to use the harmonization mechanism, there must be specified at least 2 different schedules: one schedule will control the regular creation of restore points, while another schedule will control the process of storing restore points. In terms of harmonized scheduling, Veeam Backup for AWS re-uses restore points created according to a more-frequent schedule (daily, weekly or monthly) to achieve the desired retention for less-frequent schedules (weekly, monthly and yearly). Each restore point is marked with a flag of the related schedule type: the (D) flag is used to mark restore points created daily, (W) – weekly, (M) – monthly, and (Y) – yearly. Veeam Backup for AWS uses these flags to control the retention period for the created restore points. Once a flag of a less-frequent schedule is assigned to a restore point, this restore point can no longer be removed – it is kept for the period defined in the retention settings. When the specified retention period is over, the flag is unassigned from the restore point. If the restore point does not have any other flags assigned, it is removed according to the retention settings of a more-frequent schedule.

Consider the following example. You want a backup policy to create backups of a DynamoDB table 2 times a day, to keep daily backups in the backup chain for 3 days, and also to retain one of the created backups for 2 weeks. Since you plan to access the weekly backups infrequently, you want to move one of these backups to a cold storage tier and retain it there for 6 months. In this case, you create 3 schedules when configuring the backup policy settings – daily, weekly and monthly:

Consider the following example. You want a backup policy to create backups of a DynamoDB table 2 times a day, to keep daily backups in the backup chain for 3 days, and also to retain one of the created backups for 2 weeks. Since you plan to access the weekly backups infrequently, you want to move one of these backups to a cold storage tier and retain it for 6 months. In this case, you create 3 schedules when configuring the backup policy settings – daily, weekly and monthly:

1. In the policy target settings, set the **Enable cold storage** toggle to *On* and instruct Veeam Backup for AWS to keep backups in a warm storage tier for 30 days before moving them to the cold storage tier. During this period, you will be able to perform restore from a backup stored in the high-available warm storage.
2. In the daily scheduling settings, you select hours and days when backups will be created (for example, *7:00 AM and 11:00 AM; Every Day*), and specify a number of days for which you want to keep daily restore points in a backup chain (for example, *3*). Veeam Backup for AWS will propagate these settings to the less-frequent schedules (which are the weekly and monthly schedules in our example).

Since you want to retain backups in the backup chain for only 3 days while instructing Veeam Backup for AWS to move them to the cold storage tier after 30 days, the restore points created by the daily schedule will not be moved from the warm storage tier.



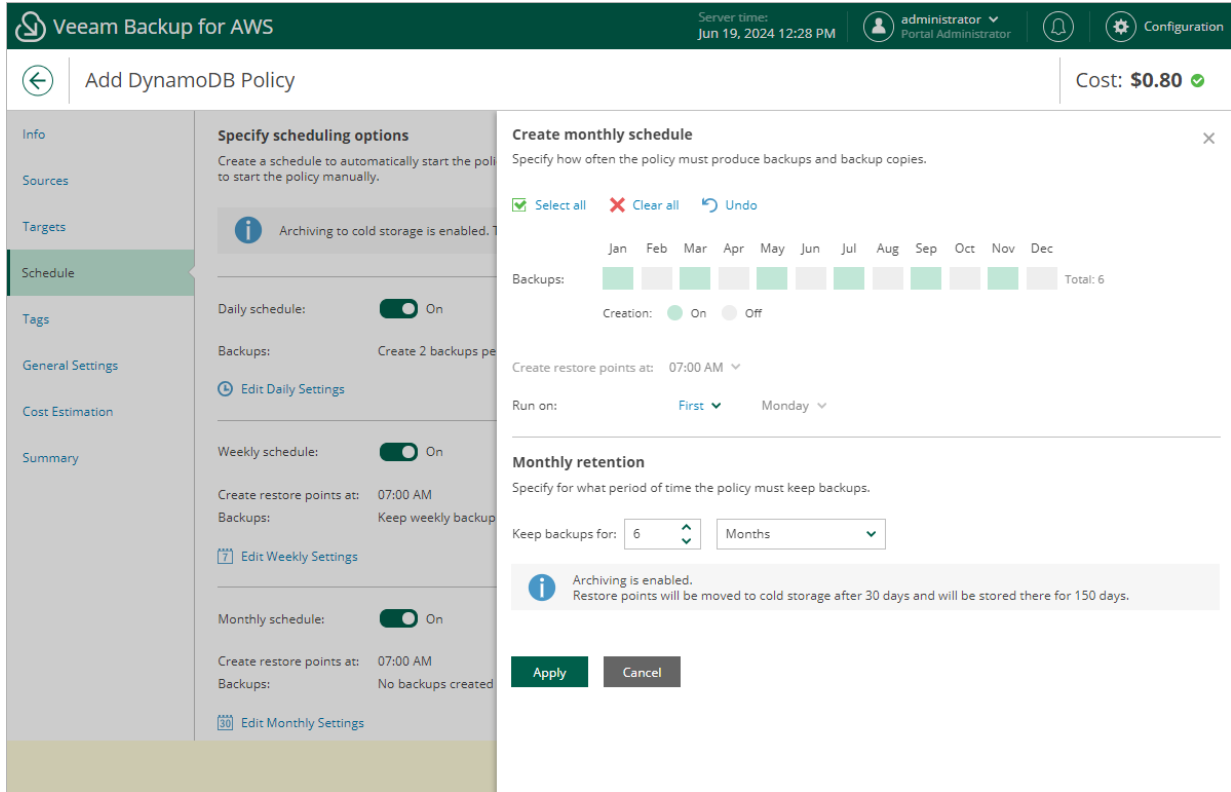
- In the weekly scheduling settings, you specify which one of the backups created by the daily schedule will be retained for a longer period, and choose for how long you want to keep the selected backup. For example, if you want to keep the daily restore point created on Monday for 2 weeks, you select *7:00 AM, Monday* and specify 14 days to keep in the weekly schedule settings.

Since you want to retain backups in the backup chain for only 14 days while instructing Veeam Backup for AWS to move them to the cold storage tier after 30 days, the restore points created by the weekly schedule will not be moved from the warm storage tier.

The screenshot displays the Veeam Backup for AWS configuration interface for adding a DynamoDB policy. The main panel is titled "Add DynamoDB Policy" with a cost of \$5.50. The left sidebar shows navigation options like Info, Sources, Targets, Schedule, Tags, General Settings, Cost Estimation, and Summary. The "Specify scheduling options" section is active, showing "Daily schedule" and "Weekly schedule" both turned on. The "Create weekly schedule" dialog is open, allowing selection of days for backups (Monday is selected) and setting the restore point time to 07:00 AM. The "Weekly retention" section is set to keep backups for 14 days. A warning message indicates that because archiving is enabled and the retention period is less than 30 days, restore points will remain in warm storage.

- In the monthly scheduling settings, you specify which one of the backups created by the weekly schedule will be retained for a longer period, and choose for how long you want to keep the selected backup. For example, *January, March, May, July, September, November, 6 months* and *First Monday*.

Since you want to retain backups in the backup chain for the full 6 months while instructing Veeam Backup for AWS to move them to the cold storage tier after 30 days, the restore points created by the monthly schedule will be moved from the warm storage tier.




According to the specified scheduling settings, Veeam Backup for AWS will create DynamoDB backups in the following way:

- On the first work day (Monday), a backup session will start at 7:00 AM to create the first restore point. The restore point will be marked with the (D) flag as it was created according to the daily schedule.


Since *7:00 AM, Monday* is specified in weekly and monthly schedule settings, Veeam Backup for AWS will also assign the (W, M) flags to this restore point. As a result, 3 flags (D, W, M) will be assigned to the restore point.

- On the same week, after starting the next backup sessions, the created restore points will be marked with the (D) flag.



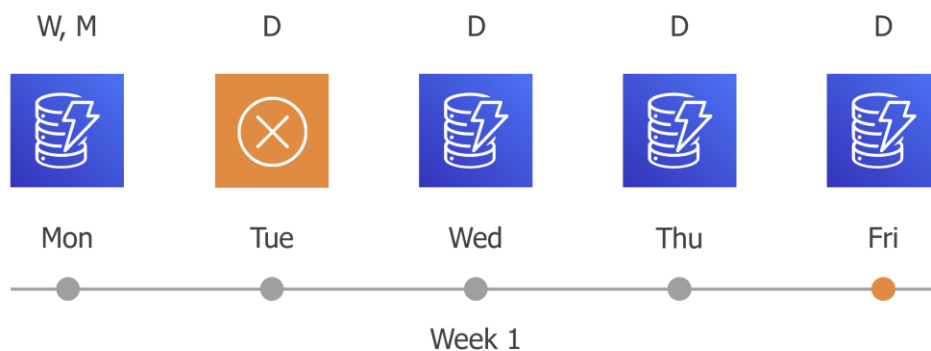
- On the fourth work day (Thursday), after a backup session runs at 7:00 AM, the created restore point will be marked with the (D) flag.

By this moment, the earliest restore point in the backup chain will get older than the specified retention limit. However, Veeam Backup for AWS will not remove the earliest restore point (7:00 AM, Monday) with the (D) flag from the backup chain as this restore point is also marked with a flag of a less-frequent schedule. Instead, Veeam Backup for AWS will unassign the (D) flag from the restore point. This restore point will be kept for the retention period specified in the weekly scheduling settings (that is, for 2 weeks).



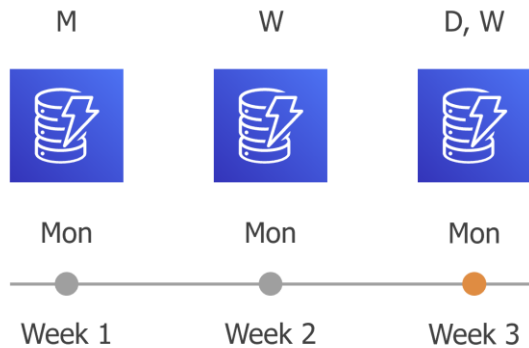
- On the fifth working day (Friday), after a backup session runs at 7:00 AM, the created restore point will be marked with the (D) flag.

By this moment, the restore point created on Tuesday with the (D) flag will get older than the specified retention limit. Veeam Backup for AWS will remove from the backup chain the restore point created at 7:00 AM on Tuesday as no flags of a less-frequent schedule are assigned to this restore point.



- Veeam Backup for AWS will continue creating restore points for the next week in the same way as described in steps 1-4.

- On week 3, after a backup session runs at 7:00 AM on Monday, the earliest weekly restore point in the backup chain will get older than the specified retention limit. However, Veeam Backup for AWS will not remove the earliest restore point (7:00 AM, Monday) with the (W) flag from the backup chain as this restore point is also marked with a flag of a less-frequent schedule. Instead, Veeam Backup for AWS will unassign the (W) flag from the restore point. This restore point will be kept for the retention period specified in the monthly scheduling settings (that is, for 6 months).



- On month 7, after a backup session runs at 7:00 AM on Monday, the earliest monthly restore point in the backup chain will get older than the specified retention limit. Veeam Backup for AWS will unassign the (M) flag from the earliest monthly restore point. Since no other flags are assigned to this restore point, Veeam Backup for AWS will remove this restore point from the backup chain.

Step 6. Enable AWS Tags Assigning

At the **Tags** step of the wizard, choose whether you want to assign AWS tags to backups and backup copies.

- To assign already existing AWS tags from the processed DynamoDB tables, select the **Copy tags from source tables** check box.

If you choose to copy tags from the source tables, Veeam Backup for AWS will first create a backup or backup copy of the DynamoDB table and assign to the created backup AWS tags with Veeam metadata, then Veeam Backup for AWS will copy tags from the processed table and, finally, assign the copied AWS tags to the backup.

- To assign your own custom AWS tags, set the **Add custom tags to created backups** toggle to *On* and specify the AWS tags explicitly. To do that, use the **Key** and **Value** fields to specify a key and value for the new custom AWS tag, and then click **Add**. Note that you cannot add more than 5 custom AWS tags.

If you choose to add custom tags to the created backups, Veeam Backup for AWS will assign the specified tags right after it creates a backup or backup copy.

The screenshot shows the 'Add DynamoDB Policy' wizard in Veeam Backup for AWS, specifically the 'Tags' step. The interface includes a top navigation bar with the Veeam logo, server time (Jun 19, 2024 12:30 PM), and user information (administrator, Portal Administrator). The main content area is titled 'Add DynamoDB Policy' and shows a cost of '\$5.50'. The 'Tags' step is active, displaying 'Specify tag settings' with instructions: 'You can copy tags from source tables and additionally assign up to 5 custom tags to backups and backup copies created by the policy. Tags can help you manage, identify, organize, search for, and filter resources.' The 'Copy tags from source tables' checkbox is checked. The 'Add custom tags to created backups' toggle is set to 'On'. Below this, there are input fields for 'Key' (containing 'owner') and 'Value' (containing 'dept01'), with an '+ Add' button. A list of existing tags shows 'department: accounting' with a close icon. A note states 'A maximum of 5 custom tags is allowed.' The bottom of the wizard features 'Previous', 'Next', and 'Cancel' buttons.

Step 7. Specify General Settings

At the **General Settings** step of the wizard, you can enable automatic retries and specify notification settings for the backup policy.

Automatic Retry Settings

To instruct Veeam Backup for AWS to run the backup policy again if it fails on the first try, do the following:

1. In the **Schedule** section of the step, select the **Automatically retry failed policy** check box.
2. In the field to the right of the check box, specify the maximum number of attempts to run the backup policy. The time interval between retries is 60 seconds.

When retrying backup policies, Veeam Backup for AWS processes only those tables that failed to be backed up during the previous attempt.

Email Notification Settings

NOTE

To be able to specify email notification settings for the DynamoDB Backup policy, you must configure [global notification settings](#) first.

To instruct Veeam Backup for AWS to send email notifications for the backup policy, do the following:

1. In the **Notifications** section of the step, set the **Enabled** toggle to *On*.
If you set the toggle to *Off*, Veeam Backup for AWS will send notifications according to the configured global notification settings.
2. In the **Email** field, specify an email address of a recipient.
Use a semicolon to separate multiple recipient addresses. Do not use spaces after semicolons between the specified email addresses.
3. Use the **Notify on** list to choose whether you want Veeam Backup for AWS to send email notifications in case the backup policy completes successfully, completes with warnings or completes with errors.
4. Select the **Suppress notifications until the last retry** check box to receive a notification about the final backup policy result.
If you do not select the check box, Veeam Backup for AWS will send a notification for every backup policy retry.

NOTE

If you specify the same email recipient in both backup policy notification and [global notification settings](#), Veeam Backup for AWS will override the configured global notification settings and will send each notification to this recipient only once to avoid notification duplicates.

The screenshot displays the 'Add DynamoDB Policy' configuration interface in Veeam Backup for AWS. The top navigation bar includes the Veeam logo, the product name, the server time (Jun 19, 2024 12:30 PM), the user role (administrator Portal Administrator), and a Configuration icon. The main header shows a back arrow, the title 'Add DynamoDB Policy', and a cost indicator of '\$5.50' with a green checkmark. A left-hand sidebar lists navigation options: Info, Sources, Targets, Schedule, Tags, General Settings (highlighted), Cost Estimation, and Summary. The main content area is titled 'Policy settings' and includes a sub-header 'Specify retry times for the policy and e-mail notifications'. It is divided into two sections: 'Schedule' and 'Notifications'. In the 'Schedule' section, the 'Automatically retry failed policy' checkbox is checked, and the value is set to 3 times. A yellow information banner below states: 'Automatic retry settings are only applicable on a scheduled run of the policy'. The 'Notifications' section has an 'Enabled' toggle switch turned 'On'. The 'Email' field contains 'donna_ortiz@company.com'. Under 'Notify on', the checkboxes for 'Failure', 'Warning', and 'Success' are all checked. At the bottom, there is a 'Suppress notifications until the last retry' checkbox, which is also checked. At the bottom right of the configuration area, there are three buttons: 'Previous', 'Next', and 'Cancel'.

Step 8. Review Estimated Cost

[This step applies only if you have created a schedule for the backup policy at the **Schedule** step of the wizard]

At the **Cost Estimation** step of the wizard, review the estimated monthly cost of AWS services and resources that will be consumed to protect the tables added to the backup policy. The total estimated cost includes the following:

- The cost of creating backups of the DynamoDB tables.
For each table included in the backup policy, Veeam Backup for AWS takes into account the number of restore points to be kept in the backup chain and the configured scheduling settings.
- The cost of creating backup copies and maintaining them in the target AWS Region.
For each table included in the backup policy, Veeam Backup for AWS takes into account the number of restore points to be kept in the backup chain and the configured scheduling settings.

NOTE

To calculate the estimated cost, Veeam Backup for AWS uses capabilities of the [AWS Pricing Calculator](#) that estimates the cost of services in USD only. This calculator is intended for informational and estimation purposes only.

The estimated cost may occur to be significantly higher due to the backup frequency, cross-region data transfer and AWS backup charges. To reduce the cost, you can try the following workarounds:

- To reduce high AWS backup charges, adjust the backup retention settings to keep less restore points in the backup chain.
- To optimize the cost of storing backups, configure the scheduling settings to run the backup policy less frequently, or instruct Veeam Backup for AWS to move backups from a high-available warm storage tier to a low-cost cold storage tier.

TIP

You can save the cost estimation as a .CSV or .XML file. To do that, click **Export to** and select the necessary format.

The screenshot shows the 'Add DynamoDB Policy' configuration page in Veeam Backup for AWS. The top navigation bar includes the Veeam logo, 'Veeam Backup for AWS', server time 'Jun 19, 2024 12:30 PM', user 'administrator', and 'Configuration' settings. The main content area is titled 'Add DynamoDB Policy' and shows a total cost of '\$5.50'. A 'Review cost estimation' section explains that the cost is based on configured settings and provides a note about predefined assumptions. Below this, three cost components are listed: Backups (\$1.79), Backup Copies (\$2.34), and Traffic (\$1.37). A large green box displays the 'Estimated monthly cost: \$5.50'. An 'Export to...' button is visible next to a search bar. At the bottom, there are 'Previous', 'Next', and 'Cancel' buttons.

Table ↑	Backup	Backup Copy	Traffic
DataTable	\$1.79	\$2.34	\$1.37

Related Resources

[How AWS Pricing Works](#)

Step 9. Finish Working with Wizard

At the **Summary** step of the wizard, it is recommended that you run the backup policy configuration check before you click **Finish**.

The configuration check will verify whether specified IAM roles have all the required permissions. To run the check, click **Test Configuration**. Veeam Backup for AWS will display the **Test policy configuration** window where you can track the progress and view the results of the check. If the IAM role permissions are insufficient or policy settings are not configured properly, the check will complete with errors, and the list of permissions that must be granted to the IAM role and policy configuration issues will be displayed in the **Test policy configuration** window.

You can grant the missing permissions to the IAM role [using the AWS Management Console](#) or instruct Veeam Backup for AWS to do it.

To let Veeam Backup for AWS grant the missing permissions:

1. In the **Test policy configuration** window, click the **Grant** link.
2. In the **Grant Permissions** window, provide one-time access keys of an IAM user that is authorized to update permissions of IAM roles, and then click **Apply**.

The IAM user must have the following permissions:

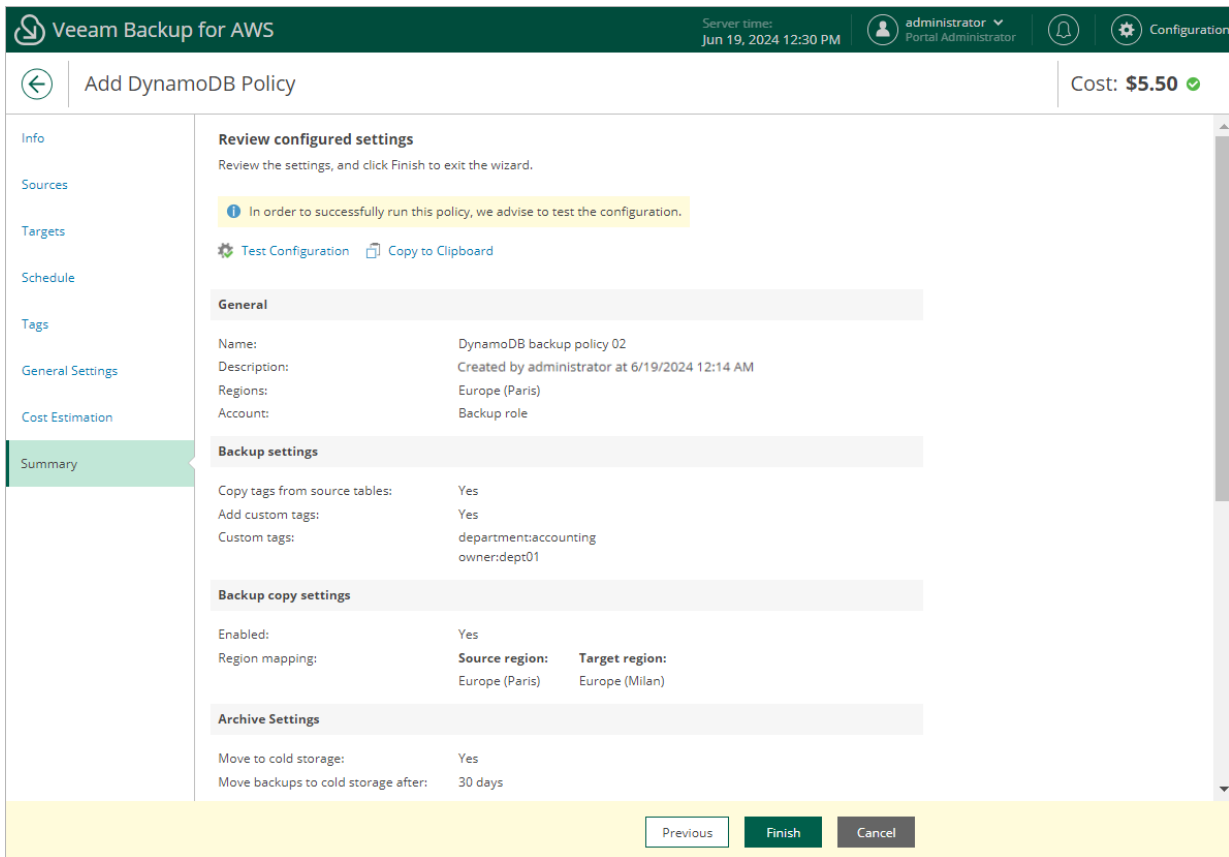
```
"iam:AttachRolePolicy",
"iam:CreatePolicy",
"iam:CreatePolicyVersion",
"iam:CreateRole",
"iam:GetAccountSummary",
"iam:GetPolicy",
"iam:GetPolicyVersion",
"iam:GetRole",
"iam:ListAttachedRolePolicies",
"iam:ListPolicyVersions",
"iam:SimulatePrincipalPolicy",
"iam:UpdateAssumeRolePolicy"
```

NOTE

Veeam Backup for AWS does not store one-time access keys in the configuration database.

3. After the required permissions are granted, close the **Test policy configuration** window, and then click **Finish** to close the **Add Policy** wizard.

Veeam Backup for AWS will save the configured backup policy.



Creating DynamoDB Backups Manually

Veeam Backup for AWS allows you to manually create backups of DynamoDB tables. You can instruct Veeam Backup for AWS to store the created backups in the same AWS Regions where the processed DynamoDB tables reside, or in a different AWS Region.

NOTE

Veeam Backup for AWS does not include backups created manually in the backup chain and does not apply the configured retention policy settings to these backups. This means that the backups are kept in your AWS environment unless you remove them manually, as described in section [Managing Backed-Up DynamoDB Data](#).

To manually create a backup of a DynamoDB table, do the following:

1. Navigate to **Resources > Databases > DynamoDB**.
2. Select the necessary table and click **Take Backup Now**.

For a DynamoDB table to be displayed in the list of available tables, an AWS Region where the table resides must be added to any of [configured DynamoDB backup policies](#), and the IAM role specified in the backup policy settings must have permissions to access the table. For more information on required permissions, see [DynamoDB Backup IAM Role Permissions](#).

3. Complete the **Take Manual Backup** wizard:

- a. At the **Account** step of the wizard, specify an IAM role whose permissions Veeam Backup for AWS will use to create the backup.

For an IAM role to be displayed in the list, it must be added to Veeam Backup for AWS as described in section [Adding IAM Roles](#).

- b. In the **Backup vault** section of the **Settings** step of the wizard, click **Edit Location Settings**.

In the **Choose region and backup vault** window, specify the following settings:

- i. From the **Target region** drop-down list, select an AWS Region where manual backups will be stored.
- ii. In the **Backup vault** section, select a backup vault that will be used to store table backups.

For a backup vault to be displayed in the list of available vaults, it must be created in the AWS Backup console as described in [AWS Documentation](#). If you have not created a backup vault for the selected AWS Region, Veeam Backup for AWS will display only the default backup vault created for the AWS Region automatically.

- iii. To save changes made to the location settings, click **Apply**.

IMPORTANT

Consider the following:

- Make sure policies assigned to the selected backup vault allow Veeam Backup for AWS to access vault resources and to perform backup, backup copy and restore operations. For more information on vault access policies, see [AWS Documentation](#).
- For Veeam Backup for AWS to be able to back up DynamoDB tables, you must configure the AWS Backup settings to enable both the Opt-in service and the advanced features for Amazon DynamoDB backups. Otherwise, Veeam Backup for AWS will automatically enable these settings for each AWS Region specified in the backup settings in your AWS account while performing backup operations. For more information on advanced DynamoDB backup, see [AWS Documentation](#).

- c. At the **Tags** section of the **Settings** step of the wizard, to assign tags to the created backup, click **Edit Tag Settings**.

In the **Tag configuration** window, specify tag settings:

- i. To assign already existing AWS tags from the processed table, select the **Copy tags from source table** check box.

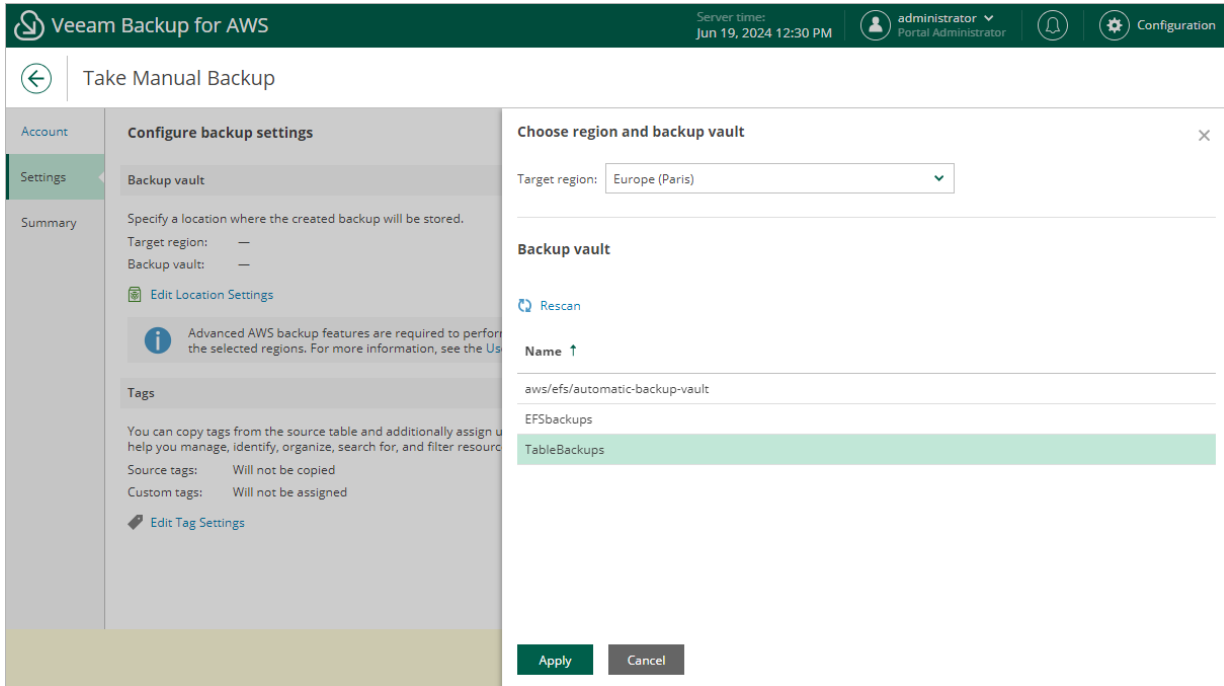
If you choose to copy tags from the source table, Veeam Backup for AWS will first create a backup of the DynamoDB table and assign to the created backup AWS tags with Veeam metadata, then Veeam Backup for AWS will copy tags from the processed table and, finally, assign the copied AWS tags to the backup.

- ii. To assign your own custom AWS tags, set the **Add custom tags to created backup** toggle to *On* and specify the tags explicitly. To do that, use the **Key** and **Value** fields to specify a key and a value for the new custom AWS tag, and then click **Add**. Note that you cannot add more than 5 custom AWS tags.

If you choose to add custom tags to the created backups, Veeam Backup for AWS will assign the specified tags right after it creates a backup.

- iii. To save changes made to the tag settings, click **Apply**.

d. At the **Summary** step of the wizard, review summary information and click **Finish**.



Performing Redshift Backup

One backup policy can be used to process one or more Redshift clusters within one AWS account. The scope of data that you can protect in an AWS account is limited by permissions of an IAM role that is specified in the backup policy settings.

NOTE

If you plan to receive email notifications on backup policy results, configure email notification settings before creating a Redshift backup policy. For more information, see [Configuring Global Notification Settings](#).

To schedule data protection tasks to run automatically, [create backup policies](#). For each protected Redshift cluster, you can also [take a backup manually](#) when needed.

IMPORTANT

Consider the following:

- Veeam Backup for AWS does not support backup of Amazon Redshift Serverless.
- Veeam Backup for AWS supports backup of Redshift clusters only to the same AWS accounts to which the source clusters belong and to the same AWS Region where the source cluster resides.

Creating Redshift Backup Policies


To create a backup policy, do the following:

1. [Launch the Add Redshift Policy wizard](#).
2. [Specify a backup policy name and description](#).
3. [Configure backup source settings](#).
4. [Configure backup target settings](#).
5. [Specify a schedule for the backup policy](#).
6. [Enable AWS tags assigning](#).
7. [Specify automatic retry settings and notification settings for the backup policy](#).
8. [Review estimated cost of the selected Redshift clusters](#).
9. [Finish working with the wizard](#).

Step 1. Launch Add Redshift Policy Wizard

To launch the **Add Redshift Policy** wizard, do the following:

1. Navigate to **Policies > Databases > Redshift**.
2. Click **Add**.



Step 2. Specify Policy Name and Description

At the **Info** step of the wizard, use the **Name** and **Description** fields to specify a name for the new backup policy and to provide a description for future reference. The name must be unique in Veeam Backup for AWS; the maximum length of the name is 127 characters, the maximum length of the description is 255 characters.

The screenshot shows the 'Add Redshift Policy' wizard in Veeam Backup for AWS. The interface is in the 'Info' step, where the user is prompted to 'Specify policy name and description'. The 'Name' field contains 'redshift_us' and the 'Description' field contains 'redshift workloads in us'. The 'Cost' is listed as 'N/A'. The 'Next' button is highlighted in green, and the 'Cancel' button is greyed out. The top navigation bar shows the user is logged in as 'administrator' (Portal Administrator) and the server time is 'Jun 12, 2024 3:17 PM'.

Veeam Backup for AWS

Server time: Jun 12, 2024 3:17 PM

administrator Portal Administrator

Configuration

← Add Redshift Policy Cost: N/A

Info

Sources

Targets

Schedule

Tags

General Settings

Cost Estimation

Summary

Specify policy name and description

Enter a name and description for the policy.

Name:

redshift_us

Description:

redshift workloads in us

Next Cancel

Step 3. Configure Backup Source Settings

At the **Sources** step of the wizard, specify backup source settings:

1. [Select an IAM role whose permissions will be used to perform Redshift backup.](#)
2. [Select AWS Regions where Redshift cluster that you plan to back up reside.](#)
3. [Select Redshift clusters to back up.](#)

Step 3.1 Specify IAM Role

In the **IAM role** section of the **Sources** step of the wizard, specify an IAM role whose permissions will be used to access AWS services and resources, and to create backups of Redshift clusters. The specified IAM role must belong to the AWS account in which the Redshift clusters that you want to protect reside, and must be assigned the permissions listed in section [Redshift Backup IAM Role Permissions](#).

For an IAM role to be displayed in the **IAM role** list, it must be added to Veeam Backup for AWS with the *Amazon Redshift Backup* operation selected for the role as described in section [Adding IAM Roles](#). If you have not added the necessary IAM role to Veeam Backup for AWS beforehand, you can do it without closing the **Add Redshift Policy** wizard. To do that, click **Add** and complete the [Add IAM Role](#) wizard.

IMPORTANT


It is recommended that you check whether the selected IAM role has all the permissions required to perform the operation. If some permissions of the IAM role are missing, the backup policy will fail to complete successfully. To run the IAM role permission check, click **Check Permissions** and follow the instructions provided in section [Checking IAM Role Permissions](#).

The screenshot shows the 'Add Redshift Policy' wizard in the Veeam Backup for AWS interface. The top navigation bar includes the Veeam logo, 'Veeam Backup for AWS', the server time 'Jun 12, 2024 3:18 PM', the user 'administrator Portal Administrator', and a 'Configuration' button. The main header shows a back arrow, 'Add Redshift Policy', and 'Cost: N/A'. A left sidebar contains navigation options: Info, Sources (highlighted), Targets, Schedule, Tags, General Settings, Cost Estimation, and Summary. The main content area is titled 'Specify source settings' and includes instructions: 'Select an IAM role to use, regions to cover and resources to process by the policy. Using tags provides dynamic selection that automatically changes the policy scope when tags are assigned to clusters.' Below this are three sections: 'IAM role' with a dropdown menu set to 'Default Backup Restore (Default Backup Restore)' and buttons for '+ Add' and 'Check Permissions'; 'Regions' with the text 'Specify one or more regions.' and a 'Choose regions...' button; and 'Resources' with the text 'Specify resources to protect or exclude.' and buttons for 'Choose resources to protect...' and 'Choose resources to exclude...'. At the bottom, there are 'Previous', 'Next', and 'Cancel' buttons.

Step 3.2 Select AWS Regions

In the **Regions** section of the **Sources** step of the wizard, choose AWS Regions where Redshift clusters that you plan to back up reside:

1. Click **Choose regions**.
2. In the **Choose regions** window, select the necessary regions, and click **Add** to include them in the backup policy.
3. To save changes made to the backup policy settings, click **Apply**.



Step 3.3 Select Redshift Clusters

In the **Resources** section of the **Sources** step of the wizard, specify the backup scope – select Redshift clusters that Veeam Backup for AWS will back up:

1. Click **Choose resources to protect**.
2. In the **Choose resources to protect** window, choose whether you want to back up all Redshift clusters from AWS Regions selected at [step 3.2](#) of the wizard, or only specific Redshift clusters.

If you select the **All resources** option, Veeam Backup for AWS will regularly check for new Redshift clusters launched in the selected regions and automatically update the backup policy settings to include these clusters into the backup scope.

If you select the **Protect only following resources** option, you must also specify the cluster explicitly:

- a. Use the **Type** drop-down list to choose whether you want to add individual Redshift clusters or AWS tags to the backup scope.

If you select the **Tag** option, Veeam Backup for AWS will back up only those Redshift clusters from the selected AWS Regions that are assigned specific tags.

- b. Use the search field to the right of the **Type** list to find the necessary resource, and then click **Protect** to add the resource to the backup scope.

For a resource to be displayed in the list of available resources, it must reside in an AWS Region that has ever been specified in any backup policy. Otherwise, the only option to discover the available resources is to click **Browse to select specific resources from the global list** and to wait for Veeam Backup for AWS to populate the resource list.

TIP

You can simultaneously add multiple resources to the backup scope. To do that, click **Browse to select specific sources from the global list**, select check boxes next to the necessary Redshift clusters or AWS tags in the list of available resources, and then click **Protect**.

If the list does not show the resources that you want to back up, click **Rescan** to launch the data collection process. As soon as the process is over, Veeam Backup for AWS will update the resource list.

If you add an AWS tag to the backup scope, Veeam Backup for AWS will regularly check for new Redshift clusters assigned the added AWS tag and automatically update the backup policy settings to include these resources in the scope. However, this applies only to Redshift clusters from the AWS Regions selected at [step 3.2](#) of the wizard. If you select an AWS tag assigned to Redshift clusters from other AWS Regions, these clusters will not be protected by the backup policy. To work around the issue, either go back to [step 3.2](#) and add the missing AWS Regions, or create a new backup policy.

3. To save changes made to the backup policy settings, click **Apply**.

TIP

As an alternative to selecting the **Protect only following resources** option and specifying the resources explicitly, you can select the **All resources** option and exclude a number of resources from the backup scope. To do that, click **Choose resources to exclude** and specify the clusters or tags that you do not want to protect – the procedure is the same as described for including resources in the backup scope.

Note that if a resource appears both in the list of included and excluded resources, Veeam Backup for AWS will still not process the resource because the list of excluded resources has a higher priority.

The screenshot shows the Veeam Backup for AWS console interface. The main window is titled "Add Redshift Policy" and includes a sidebar with navigation options like "Info", "Sources", "Targets", "Schedule", "Tags", "General Settings", "Cost Estimation", and "Summary". The "Specify source settings" section is active, showing options for IAM role, Regions, and Resources. The "Resources" section has two buttons: "Choose resources to protect..." and "Choose resources to exclude...".

The "Choose resources to protect" dialog is open, showing the following options:

- All resources
- Protect only following resources

The "Type" dropdown is set to "Cluster" and the "Name or ID" field is empty. A "Protect" button is visible. Below this, there is a search bar and a "Browse to select specific resources from the global list..." link. A "Protected resources (1)" section shows a table with one resource:

Item	ID	Value	Region
jf-check-event-cluster-...	025fc806-32e8-4c47-8e43...	—	US West (Oregon)

The "Selected: 0 of 1" text is displayed above the table. At the bottom of the dialog, there are "Apply" and "Cancel" buttons.

Step 4. Configure Backup Target Settings

At the **Targets** step of the wizard, specify backup vaults that will be used to store backups of the selected Redshift clusters:

1. In the **Backups** section, click **Choose backup vaults**.
2. In the **Choose backup vaults** window, for each AWS Region included in the policy, specify a backup vault to save and organize cluster backups. To do that:
 - a. Select an AWS Region and click **Edit**.
 - b. In the **Edit Backup Vault** window, from the **Backup vault** drop-down list, select the necessary backup vault.

For a backup vault to be displayed in the **Backup vault** list, it must be created in the AWS Backup console as described in [AWS Documentation](#). If no custom backup vaults exist in the selected AWS Region, the list will contain the default backup vault only.


IMPORTANT

Consider the following:

- Veeam Backup for AWS does not support storing backups in [logically air gapped vaults](#).
- Make sure policies assigned to the selected backup vault allow Veeam Backup for AWS to access vault resources and to perform backup and restore operations. For more information on vault access policies, see [AWS Documentation](#).
- For Veeam Backup for AWS to be able to back up Redshift clusters, you must enable the Opt-in service for the Redshift resource type in the AWS Backup settings. Otherwise, Veeam Backup for AWS will automatically enable the service for each AWS Region specified in the **Backups** section in your AWS account while performing backup operations.

c. Click **Save**.

3. To save changes made to the backup policy settings, click **Apply**.



Step 5. Specify Policy Scheduling Options

You can instruct Veeam Backup for AWS to start the backup policy automatically according to a specific backup schedule. The backup schedule defines how often data of the clusters added to the backup policy must be backed up.

To help you implement a comprehensive backup strategy, Veeam Backup for AWS allows you to create schedules of the following types:

- **Daily** – the backup policy will create restore points repeatedly throughout a day on specific days.
- **Weekly** – the backup policy will create restore points once a day on specific days.
- **Monthly** – the backup policy will create restore points once a month on a specific day.
- **Yearly** – the backup policy will create restore points once a year on a specific day.

Combining multiple schedule types together allows you to retain restore points for longer periods of time. For more information, see [Enabling Harmonized Scheduling](#).

TIP


If you do not specify the backup schedule, after you configure the backup policy, you will need to start it manually to create Redshift cluster backups. To learn how to start backup policies, see [Starting and Stopping Policies](#).

Specifying Daily Schedule

To create a daily schedule for the backup policy, at the **Schedule** step of the wizard, do the following:

1. Set the **Daily schedule** toggle to *On* and click **Edit Daily Settings**.
2. In the **Create daily schedule** window, select hours when the backup policy must create cluster backups.
If you want to protect cluster data more frequently, you can instruct the backup policy to create multiple backups per hour. To do that, click the link to the right of the **Backups** hour selection area, and specify the number of backups that the backup policy must create within an hour.
3. Use the **Run at** drop-down list to choose whether you want the backup policy to run everyday, on work days (Monday through Friday) or on specific days.
4. In the **Daily retention** section, configure retention policy settings for the daily schedule. For backups, specify the number of days (or months) for which you want to keep restore points in a backup chain.
If a restore point is older than the specified time limit, Veeam Backup for AWS removes the restore point from the chain. For more information, see [Redshift Backup Retention](#).

5. To save changes made to the backup policy settings, click **Apply**.




Specifying Weekly Schedule

To create a weekly schedule for the backup policy, at the **Schedule** step of the wizard, do the following:

1. Set the **Weekly schedule** toggle to *On* and click **Edit Weekly Settings**.
2. In the **Create weekly schedule** window, select weekdays when the backup policy must create cluster backups.
3. Use the **Create restore point at** drop-down list to schedule a specific time for the backup policy to run.
4. In the **Weekly retention** section, configure retention policy settings for the weekly schedule. For backups, specify the number of days (or months) for which you want to keep restore points in a backup chain.

If a restore point is older than the specified time limit, Veeam Backup for AWS removes the restore point from the chain. For more information, see [Redshift Backup Retention](#).

5. To save changes made to the backup policy settings, click **Apply**.



Specifying Monthly Schedule

To create a monthly schedule for the backup policy, at the **Schedule** step of the wizard, do the following:

1. Set the **Monthly schedule** toggle to *On* and click **Edit Monthly Settings**.
2. In the **Create monthly schedule** window, select months when the backup policy must create cluster backups.
3. Use the **Create restore point at** and **Run on** drop-down lists to schedule a specific time and day for the backup policy to run.

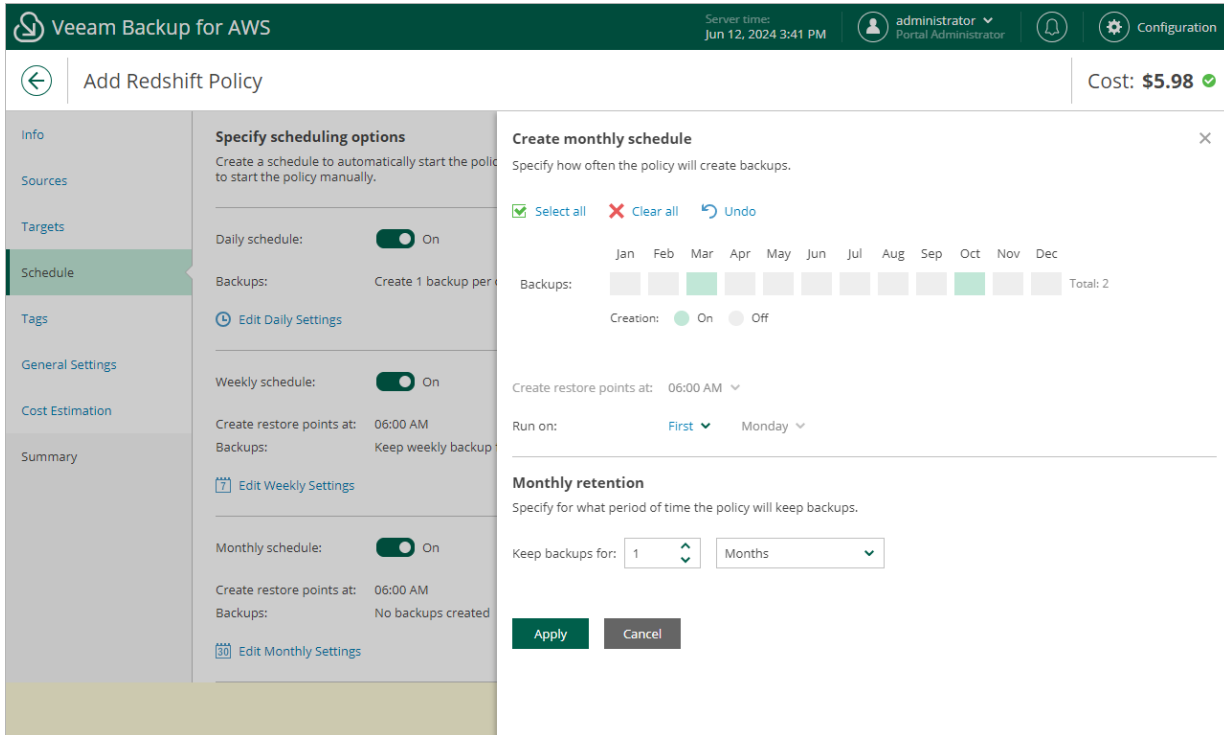
NOTE

Consider the following:

- If you have selected a specific time for the backup policy to run at the **Weekly schedule** section of the **Schedule** step of the wizard, you will not be able to change the time for the monthly schedule unless you select the *On Day* option from the **Run on** drop-down list.
 - If you select the **On day** option, **harmonized scheduling** cannot be guaranteed.
4. In the **Monthly retention** section, configure retention policy settings for the monthly schedule. For backups, specify the number of days (or months) for which you want to keep restore points in a backup chain.

If a restore point is older than the specified time limit, Veeam Backup for AWS removes the restore point from the chain. For more information, see [Redshift Backup Retention](#).

5. To save changes made to the backup policy settings, click **Apply**.



Specifying Yearly Schedule

To create a yearly schedule for the backup policy, at the **Schedule** step of the wizard, do the following:

1. Set the **Yearly schedule** toggle to *On* and click **Edit Yearly Settings**.
2. In the **Create yearly schedule** window, specify a day, month and time when the backup policy must create cluster backups.

For example, if you select *First, Friday, January* and *06:00 PM*, the backup policy will run every first Friday of January at 06:00 PM.

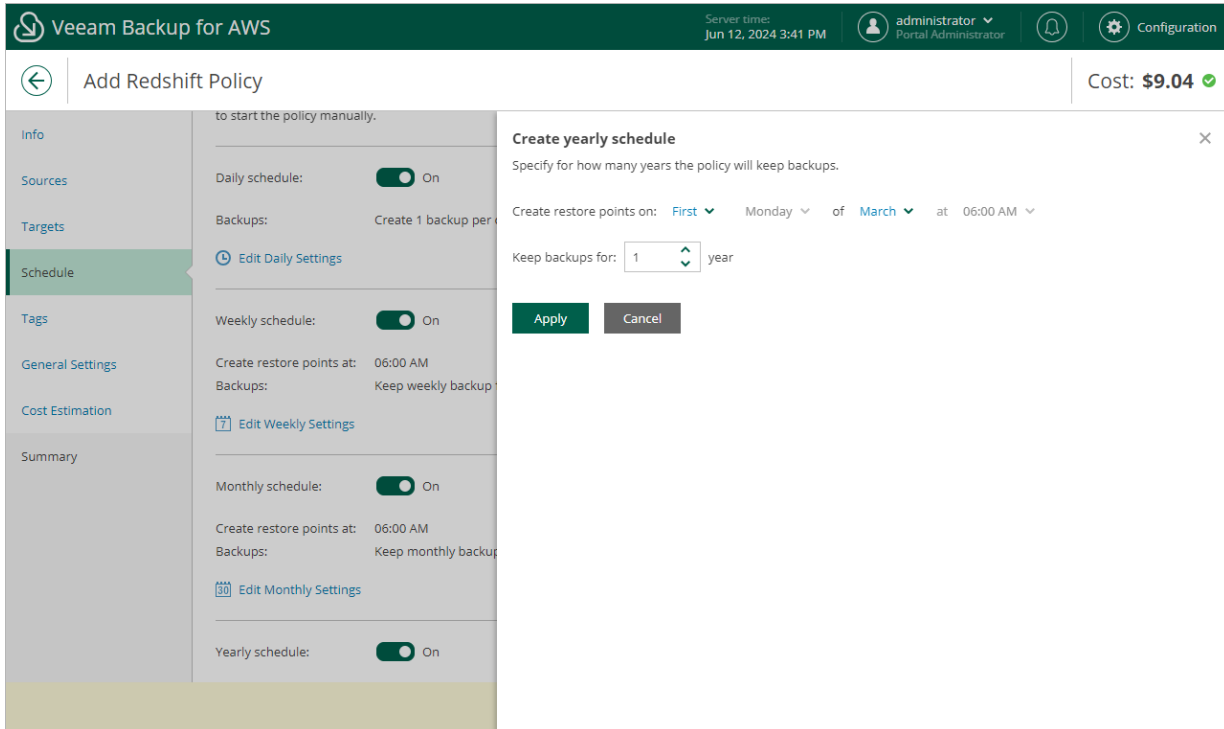
NOTE

Consider the following:

- If you have selected a specific time and day for the backup policy to run at the **Weekly schedule** or **Monthly schedule** sections of the **Schedule** step of the wizard, you will not be able to change the time and day for the yearly schedule unless you select the *On Day* option from the **Create restore point on** drop-down list.
 - If you select the *On day* option, [harmonized scheduling](#) cannot be guaranteed.
3. In the **Keep backups for** field, specify the number of years for which you want to keep restore points in a backup chain.

If a restore point is older than the specified time limit, Veeam Backup for AWS removes the restore from the chain. For more information, see [Redshift Backup Retention](#).

4. To save changes made to the backup policy settings, click **Apply**.



Enabling Harmonized Scheduling

When you combine multiple types of schedules, Veeam Backup for AWS applies the harmonization mechanism that allows you to leverage restore points for long-term retentions instead of taking a new restore point every time. The mechanism simplifies the backup schedule, optimizes the backup performance and reduces the cost of retaining restore points.


With harmonized scheduling, Veeam Backup for AWS can keep restore points created according to a daily, weekly or monthly schedule for longer periods of time: Redshift backups can be kept for weeks, months and years.

For Veeam Backup for AWS to use the harmonization mechanism, there must be specified at least 2 different schedules: one schedule will control the regular creation of restore points, while another schedule will control the process of storing restore points. In terms of harmonized scheduling, Veeam Backup for AWS re-uses restore points created according to a more-frequent schedule (daily, weekly or monthly) to achieve the desired retention for less-frequent schedules (weekly, monthly and yearly). Each restore point is marked with a flag of the related schedule type: the (D) flag is used to mark restore points created daily, (W) – weekly, (M) – monthly, and (Y) – yearly. Veeam Backup for AWS uses these flags to control the retention period for the created restore points. Once a flag of a less-frequent schedule is assigned to a restore point, this restore point can no longer be removed – it is kept for the period defined in the retention settings. When the specified retention period is over, the flag is unassigned from the restore point. If the restore point does not have any other flags assigned, it is removed according to the retention settings of a more-frequent schedule.

Consider the following example. You want a backup policy to create backups of your Redshift clusters once a day, to keep 3 daily backups in the backup chain, and also to keep one of the created backups for 2 weeks. In this case, you create 2 schedules when configuring the backup policy settings – daily and weekly:

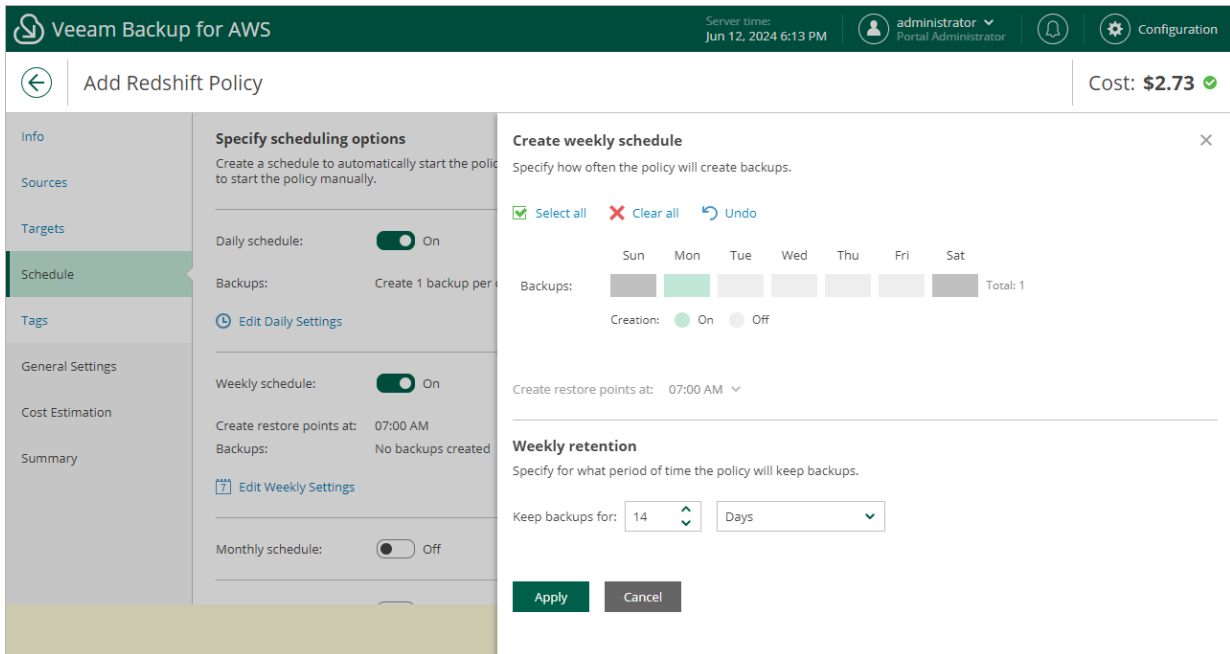
- In the daily scheduling settings, you select hours and days when backups will be created (for example, *7:00 AM; Weekdays*), and specify a number of days for which you want to keep daily restore points in a backup chain (for example, 3).

Veeam Backup for AWS will propagate these settings to the schedule of a lower frequency (which is the weekly schedule in our example).



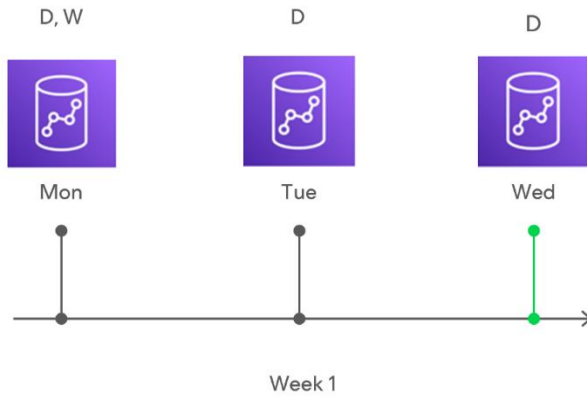
- In the weekly scheduling settings, you specify which one of the backups created by the daily schedule will be retained for a longer period, and choose for how long you want to keep the selected backup.

For example, if you want to keep the daily restore point created on Monday for 2 weeks, you select *7:00 AM, Monday* and specify 14 days to keep in the weekly schedule settings.




According to the specified scheduling settings, Veeam Backup for AWS will create Redshift backups in the following way:

1. On the first work day (Monday), a backup session will start at 7:00 AM to create the first restore point. The restore point will be marked with the (D) flag as it was created according to the daily schedule.
 Since *7:00 AM, Monday* is specified in weekly schedule settings, Veeam Backup for AWS will also assign the (W) flag to this restore point. As a result, 2 flags (D,W) will be assigned to the restore point.
2. On the same week, after starting the next backup sessions, the created restore points will be marked with the (D) flag.



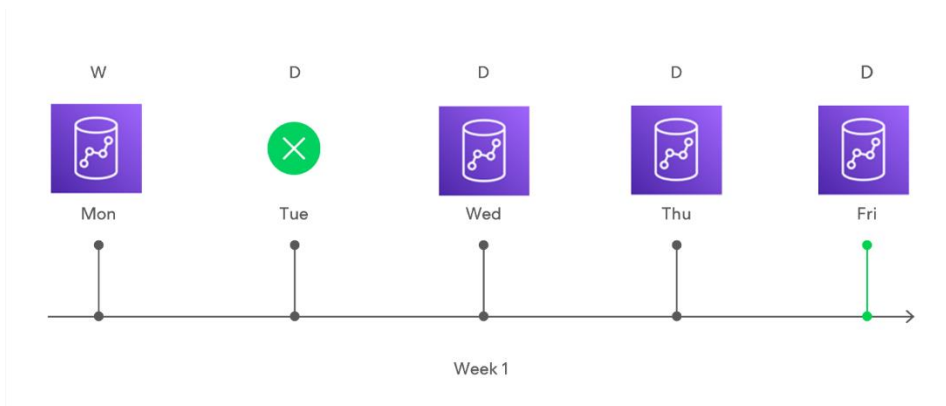
3. On the fourth work day (Thursday), after a backup session runs at 7:00 AM, the created restore point will be marked with the (D) flag.

By this moment, the earliest restore point in the backup chain will get older than the specified retention limit. However, Veeam Backup for AWS will not remove the earliest restore point (*7:00 AM, Monday*) with the (D) flag from the backup chain as this restore point is also marked with a flag of a less-frequent schedule. Instead, Veeam Backup for AWS will unassign the (D) flag from the restore point. This restore point will be kept for the retention period specified in the weekly scheduling settings (that is, for 2 weeks).

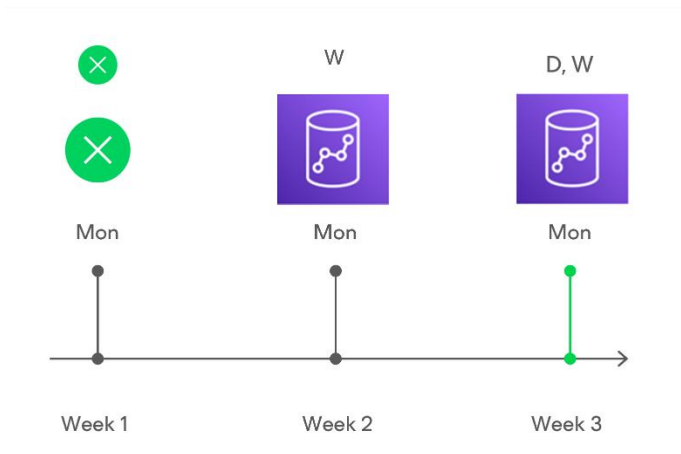


- On the fifth working day (Friday), after a backup session runs at 7:00 AM, the created restore point will be marked with the (D) flag.

By this moment, the restore point created on Tuesday with the (D) flag will get older than the specified retention limit. Veeam Backup for AWS will remove from the backup chain the restore point created at 7:00 AM on Tuesday as no flags of a less-frequent schedule are assigned to this restore point.



- Veeam Backup for AWS will continue creating restore points for the next week in the same way as described in steps 1-4.
- On week 3, after a backup session runs at 7:00 AM on Monday, the earliest weekly restore point in the backup chain will get older than the specified retention limit. Veeam Backup for AWS will unassign the (W) flag from the earliest weekly restore point. Since no other flags are assigned to this restore point, Veeam Backup for AWS will remove this restore point from the backup chain.



Step 6. Enable AWS Tags Assigning

At the **Tags** step of the wizard, choose whether you want to assign AWS tags to backups.

- To assign already existing AWS tags from the processed Redshift clusters, select the **Copy tags from source clusters** check box.

If you choose to copy tags from the source clusters, Veeam Backup for AWS will first create a backup of the Redshift cluster and assign to the created backup AWS tags with Veeam metadata, then Veeam Backup for AWS will copy tags from the processed cluster and, finally, assign the copied AWS tags to the backup.

- To assign your own custom AWS tags, set the **Add custom tags to created backups** toggle to *On* and specify the AWS tags explicitly. To do that, use the **Key** and **Value** fields to specify a key and value for the new custom AWS tag, and then click **Add**. Note that you cannot add more than 5 custom AWS tags.

If you choose to add custom tags to the created backups, Veeam Backup for AWS will assign the specified tags right after it creates a backup.

The screenshot shows the 'Add Redshift Policy' wizard in the Veeam Backup for AWS interface. The 'Tags' step is active, showing options to copy tags from source clusters and add custom tags. The 'Add custom tags to created backups' toggle is set to 'On'. Two custom tags are listed: 'location' with value 'us' and 'owners: dept01'. The cost is estimated at \$9.04.

Specify tag settings
You can copy tags from source clusters and additionally assign up to 5 custom tags to backups created by the policy. Tags can help you manage, identify, organize, search for, and filter resources.

Copy tags from source clusters

Add custom tags to created backups: On

Key: location Value: us + Add

owners: dept01 x

A maximum of 5 custom tags is allowed.

Cost: \$9.04 ✓

Buttons: Previous, Next, Cancel

Step 7. Specify General Settings

At the **General Settings** step of the wizard, you can enable automatic retries and specify notification settings for the backup policy.

Automatic Retry Settings

To instruct Veeam Backup for AWS to run the backup policy again if it fails on the first try, do the following:

1. In the **Schedule** section of the step, select the **Automatically retry failed policy** check box.
2. In the field to the right of the check box, specify the maximum number of attempts to run the backup policy. The time interval between retries is 60 seconds.

When retrying backup policies, Veeam Backup for AWS processes only those clusters that failed to be backed up during the previous attempt.

Email Notification Settings

NOTE

To be able to specify email notification settings for the Redshift Backup policy, you must configure [global notification settings](#) first.

To instruct Veeam Backup for AWS to send email notifications for the backup policy, do the following:

1. In the **Notifications** section of the step, set the **Enabled** toggle to *On*.
If you set the toggle to *Off*, Veeam Backup for AWS will send notifications according to the configured global notification settings.
2. In the **Email** field, specify an email address of a recipient.
Use a semicolon to separate multiple recipient addresses. Do not use spaces after semicolons between the specified email addresses.
3. Use the **Notify on** list to choose whether you want Veeam Backup for AWS to send email notifications in case the backup policy completes successfully, completes with warnings or completes with errors.
4. Select the **Suppress notifications until the last retry** check box to receive a notification about the final backup policy result.
If you do not select the check box, Veeam Backup for AWS will send a notification for every backup policy retry.

NOTE

If you specify the same email recipient in both backup policy notification and [global notification settings](#), Veeam Backup for AWS will override the configured global notification settings and will send each notification to this recipient only once to avoid notification duplicates.

The screenshot shows the 'Add Redshift Policy' configuration page in Veeam Backup for AWS. The page is titled 'Add Redshift Policy' and has a cost of '\$9.04'. The left sidebar contains navigation options: Info, Sources, Targets, Schedule, Tags, General Settings (selected), Cost Estimation, and Summary. The main content area is titled 'Configure retry and notification settings' and includes the following sections:

- Schedule:** A checkbox for 'Automatically retry failed policy' is checked, with a dropdown menu set to '3' times. A yellow information box below states: 'Automatic retry settings are only applicable on a scheduled run of the policy'.
- Notifications:** A toggle switch for 'Enabled' is turned 'On'. The 'Email' field contains 'donna_ortiz@company.com'. Below this, there are four checkboxes for notification types: 'Failure', 'Warning', 'Success', and 'Suppress notifications until the last retry', all of which are checked.

At the bottom of the page, there are three buttons: 'Previous', 'Next', and 'Cancel'.

Step 8. Review Estimated Cost

[This step applies only if you have created a schedule for the backup policy at the **Schedule** step of the wizard]

At the **Cost Estimation** step of the wizard, review the estimated monthly cost of AWS services and resources that will be consumed to protect the clusters added to the backup policy.

The total estimated cost includes the cost of creating backups of the Redshift clusters. For each cluster included in the backup policy, Veeam Backup for AWS takes into account the number of restore points to be kept in the backup chain and the configured scheduling settings.

NOTE

To calculate the estimated cost, Veeam Backup for AWS uses capabilities of the [AWS Pricing Calculator](#) that estimates the cost of services in USD only. This calculator is intended for informational and estimation purposes only.

The estimated cost may occur to be significantly higher due to the backup frequency and AWS backup charges. To reduce the cost, you can try the following workarounds:

- To reduce high AWS backup charges, adjust the backup retention settings to keep less restore points in the backup chain.
- To optimize the cost of storing backups, configure the scheduling settings to run the backup policy less frequently.

TIP

You can save the cost estimation as a .CSV or .XML file. To do that, click **Export to** and select the necessary format.

Veeam Backup for AWS

Server time: Jun 12, 2024 3:43 PM

administrator Portal Administrator

Configuration

← Add Redshift Policy Cost: \$9.04 ✓

Info

Sources

Targets

Schedule

Tags

General Settings

Cost Estimation

Summary

Review cost estimation

The estimated cost takes into account the configured target settings, the specified scheduling options and the number of resources to protect.

Note that Veeam Backup for AWS makes predefined assumptions to calculate the cost, which means that the results should be used only as an approximation. For more information on cost calculation, see [this Veeam KB article](#).

Estimated monthly cost: \$9.04

Cluster

Cluster ↑	Backup	Total
jf-check-event-cluster-oregon	\$9.04	\$9.04

Related Resources

[How AWS Pricing Works](#)

Step 9. Finish Working with Wizard

At the **Summary** step of the wizard, it is recommended that you run the backup policy configuration check before you click **Finish**.

The configuration check will verify whether specified IAM roles have all the required permissions. To run the check, click **Test Configuration**. Veeam Backup for AWS will display the **Test policy configuration** window where you can track the progress and view the results of the check. If the IAM role permissions are insufficient or policy settings are not configured properly, the check will complete with errors, and the list of permissions that must be granted to the IAM role and policy configuration issues will be displayed in the **Test policy configuration** window.

You can grant the missing permissions to the IAM role [using the AWS Management Console](#) or instruct Veeam Backup for AWS to do it.

To let Veeam Backup for AWS grant the missing permissions:

1. In the **Test policy configuration** window, click the **Grant** link.
2. In the **Grant Permissions** window, provide one-time access keys of an IAM user that is authorized to update permissions of IAM roles, and then click **Apply**.

The IAM user must have the following permissions:

```
"iam:AttachRolePolicy",  
"iam:CreatePolicy",  
"iam:CreatePolicyVersion",  
"iam:CreateRole",  
"iam:GetAccountSummary",  
"iam:GetPolicy",  
"iam:GetPolicyVersion",  
"iam:GetRole",  
"iam:ListAttachedRolePolicies",  
"iam:ListPolicyVersions",  
"iam:SimulatePrincipalPolicy",  
"iam:UpdateAssumeRolePolicy"
```

NOTE

Veeam Backup for AWS does not store one-time access keys in the configuration database.

3. After the required permissions are granted, close the **Test policy configuration** window, and then click **Finish** to close the **Add Policy** wizard.

Veeam Backup for AWS will save the configured backup policy.

Server time: Jun 12, 2024 3:43 PM administrator Portal Administrator Configuration

← Add Redshift Policy Cost: \$9.04 ✓

Review configured settings
Review the settings, and click Finish to exit the wizard.

Info
Sources
Targets
Schedule
Tags
General Settings
Cost Estimation
Summary

General
Name: reshift_us
Description: protection of redshift workloads in us
Regions: US West (Oregon)
Account: Default Backup Restore

Backup settings
Copy tags from source clusters: Yes
Add custom tags: Yes
Custom tags: owner:dept01

Backup schedule
Daily retention: Create 1 restore points and keep for 14 Days
Weekly retention: Create 1 restore points and keep for 1 Months
Monthly retention: Create 2 restore points and keep for 1 Months
Yearly retention: Create restore point on First Monday of March at 06:00 AM
Keep backups for 1 years

General settings
Automatic retry enabled: Yes
Notifications enabled: Yes

Resources
Added resources: jf-check-event-cluster-oregon
Excluded resources: —

Previous Finish Cancel

Creating Redshift Backups Manually

Veeam Backup for AWS allows you to manually create backups of Redshift clusters. You can instruct Veeam Backup for AWS to store the created backups only in the same AWS Regions where the processed Redshift clusters reside.

NOTE

Veeam Backup for AWS does not include backups created manually in the backup chain and does not apply the configured retention policy settings to these backups. This means that the backups are kept in your AWS environment unless you remove them manually, as described in section [Managing Backed-Up Redshift Data](#).

To manually create a backup of a Redshift cluster, do the following:

1. Navigate to **Resources > Databases > Redshift**.

2. Select the necessary cluster and click **Take Backup Now**.

For a Redshift cluster to be displayed in the list of available clusters, an AWS Region where the cluster resides must be added to any of [configured Redshift backup policies](#), and the IAM role specified in the backup policy settings must have permissions to access the cluster. For more information on required permissions, see [Redshift Backup IAM Role Permissions](#).

3. Complete the **Take Manual Backup** wizard:

- a. At the **Account** step of the wizard, specify an IAM role whose permissions Veeam Backup for AWS will use to create the backup.

For an IAM role to be displayed in the list, it must be added to Veeam Backup for AWS as described in section [Adding IAM Roles](#).

- b. In the **Backup vault** section of the **Settings** step of the wizard, click **Edit Location Settings**.

In the **Choose a backup vault** window, select a backup vault that will be used to store cluster backups. For a backup vault to be displayed in the list of available vaults, it must be created in the AWS Backup console as described in [AWS Documentation](#).

Note that since [cross-region copying of Redshift backups](#) is not supported by the AWS Backup service, the list of available vaults will include only those backup vaults that exist in the AWS Region where the selected source cluster resides.

IMPORTANT

Consider the following:

- Make sure policies assigned to the selected backup vault allow Veeam Backup for AWS to access vault resources and to perform backup and restore operations. For more information on vault access policies, see [AWS Documentation](#).
- For Veeam Backup for AWS to be able to back up Redshift clusters, you must enable the Opt-in service for the Redshift resource type in the AWS Backup settings. Otherwise, Veeam Backup for AWS will automatically enable the service for the AWS Region specified in the backup settings in your AWS account while performing backup operations.

- c. At the **Tags** section of the **Settings** step of the wizard, to assign tags to the created backup, click **Edit Tag Settings**.

In the **Tag configuration** window, specify tag settings:

- i. To assign already existing AWS tags from the processed cluster, select the **Copy tags from source cluster** check box.

If you choose to copy tags from the source cluster, Veeam Backup for AWS will first create a backup of the Redshift cluster and assign to the created backup AWS tags with Veeam metadata, then Veeam Backup for AWS will copy tags from the processed cluster and, finally, assign the copied AWS tags to the backup.

- ii. To assign your own custom AWS tags, set the **Add custom tags to created backup** toggle to *On* and specify the tags explicitly. To do that, use the **Key** and **Value** fields to specify a key and a value for the new custom AWS tag, and then click **Add**. Note that you cannot add more than 5 custom AWS tags.

If you choose to add custom tags to the created backups, Veeam Backup for AWS will assign the specified tags right after it creates a backup.

- iii. To save changes made to the tag settings, click **Apply**.

d. At the **Summary** step of the wizard, review summary information and click **Finish**.

The screenshot shows the Veeam Backup for AWS interface. At the top, the header includes the Veeam logo, the text 'Veeam Backup for AWS', the server time 'Jun 12, 2024 3:51 PM', and the user 'administrator Portal Administrator'. The main navigation bar shows 'Take Manual Backup' with a back arrow. On the left, a sidebar contains 'Account', 'Settings', and 'Summary'. The 'Settings' section is active, showing 'Configure backup settings' with options for 'Backup vault', 'Target region: US West (Oregon)', and 'Tags'. A modal dialog titled 'Choose a backup vault' is open, displaying a list of vaults: 'aws/efs/automatic-backup-vault', 'Default', 'jf_vault', 'jf-long-backup-Oregon', 'jf-vault-to-delete', and 'veeamon-backupvault-oregon'. An information message at the top of the dialog states: 'The AWS backup service opt-in setting must be enabled for the selected region. If required, it will be automatically enabled. For more information, see the User Guide.' At the bottom of the dialog are 'Apply' and 'Cancel' buttons.

Performing EFS Backup

One backup policy can be used to process one or more EFS file systems within one AWS account. The scope of data that you can protect in an AWS account is limited by permissions of an IAM role that is specified in the backup policy settings.

NOTE

If you plan to receive email notifications on backup policy results, configure email notification settings before creating an EFS backup policy. For more information, see [Configuring Global Notification Settings](#).

To schedule data protection tasks to run automatically, [create backup policies](#). For each protected EFS systems, you can also [take a backup manually](#) when needed.

IMPORTANT

You can back up EFS file systems only to the same AWS accounts where the source file systems belong.

Creating EFS Backup Policies

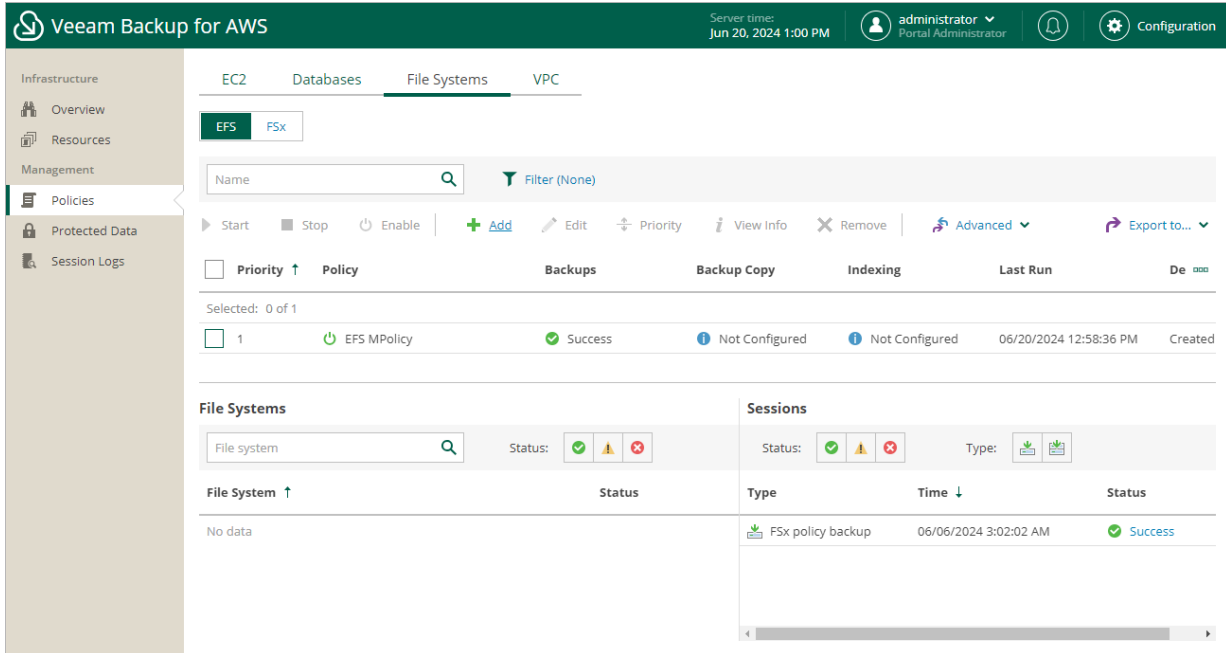
To create a backup policy, do the following:

1. [Launch the Add EFS Policy wizard](#).
2. [Specify a backup policy name and description](#).
3. [Configure backup source settings](#).
4. [Enable indexing for the processed file systems](#).
5. [Configure backup target settings](#).
6. [Specify a schedule for the backup policy](#).
7. [Enable AWS tags assigning](#).
8. [Specify automatic retry settings and notification settings for the backup policy](#).
9. [Review estimated cost of the selected EFS file systems](#).
10. [Finish working with the wizard](#).

Step 1. Launch Add EFS Policy Wizard

To launch the **Add EFS Policy** wizard, do the following:

1. Navigate to **Policies > File Systems > EFS**.
2. Click **Add**.



Step 2. Specify Policy Name and Description

At the **Info** step of the wizard, use the **Name** and **Description** fields to specify a name for the new backup policy and to provide a description for future reference. The name must be unique in Veeam Backup for AWS; the maximum length of the name is 127 characters, the maximum length of the description is 255 characters.

The screenshot shows the 'Add EFS Policy' wizard in Veeam Backup for AWS. The interface includes a top navigation bar with the Veeam logo, 'Veeam Backup for AWS', server time 'Jun 20, 2024 1:00 PM', and user 'administrator Portal Administrator'. A left sidebar lists steps: Info (selected), Sources, Indexing, Targets, Schedule, Tags, General Settings, Cost Estimation, and Summary. The main area is titled 'Specify policy name and description' with the instruction 'Enter a name and description for the policy.' It contains two input fields: 'Name' with the value 'EFS Backup Policy' and 'Description' with the value 'Backup of file system for D01'. A 'Cost: N/A' indicator is in the top right. At the bottom, there are 'Next' and 'Cancel' buttons.

Step 3. Configure Backup Source Settings

At the **Sources** step of the wizard, specify backup source settings:

1. [Select an IAM role whose permissions will be used to perform EFS file system backup.](#)
2. [Select AWS Regions where EFS file systems that you plan to back up reside.](#)
3. [Select EFS file systems to back up.](#)

Step 3.1 Specify IAM Role

In the **IAM role** section of the **Sources** step of the wizard, specify an IAM role whose permissions will be used to access AWS services and resources, and to create backups of Amazon EFS file systems. The specified IAM role must belong to the AWS account in which the file systems that you want to protect reside, and must be assigned permissions listed in section [EFS Backup IAM Role Permissions](#).

For an IAM role to be displayed in the **IAM role** list, it must be added to Veeam Backup for AWS with the *Amazon EFS Backup* operation selected as described in section [Adding IAM Roles](#). If you have not added the necessary IAM role to Veeam Backup for AWS beforehand, you can do it without closing the **Add EFS Policy** wizard. To do that, click **Add** and complete the [Add IAM Role](#) wizard.

IMPORTANT


It is recommended that you check whether the selected IAM role has all the permissions required to perform the operation. If some permissions of the IAM role are missing, the backup policy will fail to complete successfully. To run the IAM role permission check, click **Check Permissions** and follow the instructions provided in section [Checking IAM Role Permissions](#).

The screenshot shows the 'Add EFS Policy' wizard in the Veeam Backup for AWS interface. The top navigation bar includes the Veeam logo, 'Veeam Backup for AWS', server time 'Jun 20, 2024 1:01 PM', user 'administrator Portal Administrator', and a 'Configuration' link. The main header shows a back arrow, 'Add EFS Policy', and 'Cost: N/A'. A left sidebar lists steps: Info, Sources (selected), Indexing, Targets, Schedule, Tags, General Settings, Cost Estimation, and Summary. The main content area is titled 'Specify source settings' and includes instructions: 'Select an IAM role to use, regions to cover and resources to process by the policy. Using tags provides dynamic selection that automatically changes the policy scope when tags are assigned to file systems.' Below this are three sections: 'IAM role' with a dropdown menu showing 'Default Backup Restore (Default Backup Restore)' and buttons for '+ Add' and 'Check Permissions'; 'Regions' with the text 'Specify one or more regions.' and a 'Choose regions...' button; and 'Resources' with the text 'Specify resources to protect or exclude.' and buttons for 'Choose resources to protect...' and 'Choose resources to exclude...'. At the bottom, there are 'Previous', 'Next', and 'Cancel' buttons.

Step 3.2 Select AWS Regions

In the **Regions** section of the **Sources** step of the wizard, select AWS Regions where EFS file systems that you plan to back up reside:

1. Click **Choose regions**.
2. In the **Choose regions** window, select the necessary AWS Regions from the **Available Regions** list, and click **Add**.
3. To save changes made to the backup policy settings, click **Apply**.



Step 3.3 Select EFS File Systems

In the **Resources** section of the **Sources** step of the wizard, specify the backup scope – select EFS file systems that Veeam Backup for AWS will back up:

1. Click **Choose resources to protect**.
2. In the **Choose resources to protect** window, choose whether you want to back up all EFS file systems from AWS Regions selected at [step 3.2](#) of the wizard, or only specific file systems.

If you select the **All resources** option, Veeam Backup for AWS will regularly check for new EFS file systems reside in the selected regions and automatically update the backup policy settings to include these file systems into the backup scope.

If you select the **Protect only following resources** option, you must specify the EFS file systems explicitly:

- a. Use the **Type** drop-down list to choose whether you want to add individual file systems or AWS tags to the backup scope.

If you select the *Tag* option, Veeam Backup for AWS will back up only those file systems that reside in the selected AWS Regions under specific AWS tags.

- b. Use the search field of the **Name or ID** drop-down list to find the necessary resource, and then click **Protect** to add the resource to the backup scope.

For a resource to be displayed in the list of available resources, it must reside in an AWS Region that has ever been specified in any backup policy. Otherwise, the only option to discover the available resources is to click **Browse to select specific resources from the global list** and to wait for Veeam Backup for AWS to populate the resource list.

NOTE

By default, Veeam Backup for AWS uses AWS CloudTrail to track changes in your EFS resources. If no trails are configured in the source AWS account, Veeam Backup for AWS will automatically access AWS resources and populate the list of available file systems or AWS tags only once in 24 hours. To manually force the data collection process, click **Rescan**.

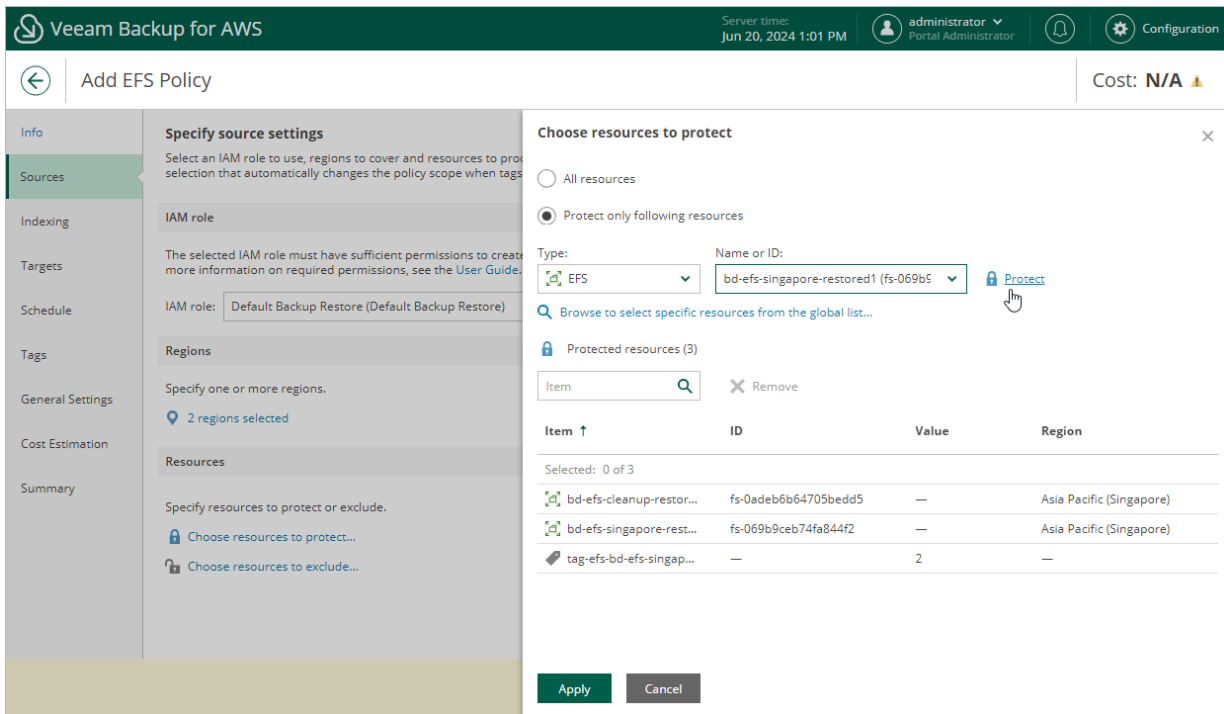
If you add an AWS tag to the backup scope, Veeam Backup for AWS will regularly check for new Amazon EFS file systems assigned the added AWS tag and automatically update the backup policy settings to include these file systems in the scope. However, this applies only to file systems from the AWS Regions selected at [step 3.2](#) of the wizard. If you select a tag assigned to file systems from other regions, these file systems will not be protected by the backup policy. To work around the issue, either go back to [step 3.2](#) and add the missing regions, or create a new backup policy.

- c. To save changes made to the backup policy settings, click **Apply**.

TIP

As an alternative to selecting the **Protect only following resources** option and specifying the resources explicitly, you can select the **All resources** option and exclude a number of resources from the backup scope. To do that, click **Choose resources to exclude** and specify the file system or tags that you do not want to protect – the procedure is the same as described for including resources in the backup scope.

Note that if a resource appears both in the list of included and excluded resources, Veeam Backup for AWS will still not process the resource because the list of excluded resources has a higher priority.



The screenshot shows the Veeam Backup for AWS console interface. The main window is titled "Add EFS Policy" and has a "Cost: N/A" indicator. On the left, there is a navigation menu with options like "Info", "Sources", "Indexing", "Targets", "Schedule", "Tags", "General Settings", "Cost Estimation", and "Summary". The "Specify source settings" section is active, showing "IAM role" set to "Default Backup Restore (Default Backup Restore)" and "Regions" with "2 regions selected".

The "Choose resources to protect" dialog is open, showing two radio button options: "All resources" (unselected) and "Protect only following resources" (selected). Below these, there are input fields for "Type" (set to "EFS") and "Name or ID" (set to "bd-efs-singapore-restored1 (fs-069b5...)"). A "Protect" button is visible next to the "Name or ID" field. Below the input fields, there is a link: "Browse to select specific resources from the global list...".

Below the link, there is a section for "Protected resources (3)". It includes a search bar and a "Remove" button. A table lists the resources:

Item	ID	Value	Region
Selected: 0 of 3			
bd-efs-cleanup-restor...	fs-0adeb6b64705bedd5	—	Asia Pacific (Singapore)
bd-efs-singapore-rest...	fs-069b9ceb74fa844f2	—	Asia Pacific (Singapore)
tag-efs-bd-efs-singap...	—	2	—

At the bottom of the dialog, there are "Apply" and "Cancel" buttons.

Step 4. Enable EFS Indexing

At the **Indexing** step of the wizard, you can instruct Veeam Backup for AWS to perform indexing of the processed EFS file systems. EFS indexing allows you to perform EFS file-level recovery operations without specifying the exact paths to the necessary files folders and to restore them using different restore points during one restore session. While performing EFS indexing of a file system, Veeam Backup for AWS creates a catalog of all files and directories (an index) and saves the index to a backup repository. This index is further used to reproduce the file system structure and to enable browsing and searching for specific files within an EFS backup.

To learn how indexing works, see [EFS Backup](#).

NOTE

To perform indexing of the EFS file systems, Veeam Backup for AWS deploys a worker instance per each processed file system in the same AWS account where the file system resides – production account. By default, the most appropriate network settings of AWS Regions are used to deploy these worker instances. However, you can add [specific worker configurations](#) that will be used to deploy worker instances used for EFS indexing operations.

Limitations and Requirements

Before you enable EFS indexing, consider the following:

- EFS indexing is not supported in the *Free* edition of Veeam Backup for AWS. For more information, see [Licensing](#).
- Each processed EFS file system for which you want to perform indexing must meet the following requirements:
 - A file system must have at least one mount target created.
 - A mount target that will be used by worker instances to connect to the file system must be associated with a security group that allows inbound access on port **2049**.
- If no specific [worker configurations](#) are added to Veeam Backup for AWS, the most appropriate network settings of AWS Regions are used to deploy worker instances for EFS indexing operations. For Veeam Backup for AWS to be able to deploy a worker instance used to create an index of a file system:
 - A VPC in which the file system has the mount target must have at least one security group that allows outbound access on ports **2049** and **443**. These ports are used by worker instances to mount the file system and to communicate with [AWS services](#).
 - The DNS resolution option must be enabled for the VPC. For more information, see [AWS Documentation](#).
 - As Veeam Backup for AWS uses public access to communicate with worker instances, the [public IPv4 addressing](#) attribute must be enabled at least for one subnet in the Availability Zone in which the file system has a mount target and the VPC to which the subnet belongs must have an [internet gateway attached](#). VPC and subnet route tables must have routes that direct internet-bound traffic to this internet gateway.

If you want worker instances to operate in a private network, enable the [private network deployment](#) functionality and configure [specific VPC endpoints](#) for the subnet to let Veeam Backup for AWS use private IPv4 addresses. Alternatively, configure VPC interface endpoints as described in section [Appendix C. Configuring Endpoints in AWS](#).

Enabling EFS Indexing

To enable indexing of the processed file systems, do the following:

1. Set the **Enable indexing** toggle to *On*.
2. In the **Repositories** window, select a repository where the created EFS indexes will be stored, and click **Apply**.

For a backup repository to be displayed in the **Repositories** list, it must be added to Veeam Backup for AWS as described in section [Adding Backup Repositories Using Web UI](#). The list shows only backup repositories of the *S3 Standard* storage class that have encryption enabled and immutability disabled.

3. In the **IAM role** section, choose an IAM role that will be attached to the worker instances and used by Veeam Backup for AWS to communicate with these instances. The role must be assigned permissions listed in section [Worker Deployment Role Permissions in Production Accounts](#).

For an IAM role to be displayed in the list, it must be added to Veeam Backup for AWS with the *Production worker role* selected as described in section [Adding IAM Roles](#). The list shows only IAM roles that belong to the production account – account to which the file systems belong. Note that the specified IAM role must be included in one or more instance profiles. For more information on instance profiles, see [AWS Documentation](#).

IMPORTANT

It is recommended that you check whether both the IAM role specified at [step 3.1](#) of the wizard and the IAM role specified in the **IAM role** section have the required permissions. If some permissions of the IAM role are missing, the backup policy will fail to complete successfully. To run the IAM role permission check, click **Check Permissions** and follow the instructions provided in section [Checking IAM Role Permissions](#).

The screenshot shows the 'Add EFS Policy' wizard in Veeam Backup for AWS. The 'Indexing' section is active, with 'Enable indexing' set to 'On' and 'Indexes will be stored in' set to 'Choose repository...'. The 'IAM role' section shows 'IAM role: bd-fsx-autogen-new-permission-list (Created by...)'. A 'Repositories' dialog box is open, displaying a table of repositories.

Repository	Region	Folder	Description
Repo config back...	Europe (Paris)	Backups	Created by bd at 4/1...
SQS Repo	Europe (Paris)	SQS	Created by bd at 6/3...
SQS Repo acc22	Europe (Paris)	SQS	Created by bd at 6/3...

Step 5. Configure Backup Target Settings

By default, backup policies create only backups of processed EFS file systems. At the **Targets** step of the wizard, you can specify the following backup target settings:

- [Specify backup vaults where Veeam Backup for AWS will store EFS file system backups.](#)
- [Instruct Veeam Backup for AWS to copy EFS file system backups to other AWS Regions.](#)

Configuring Backup Settings

To specify backup vaults used to store backups of the selected EFS file systems, do the following:

1. In the **Backups** section of the **Targets** step of the wizard, click **Choose backup vaults**.
2. In the **Choose backup vaults** window, for each AWS Region included in the policy, specify a backup vault to save and organize file system backups. To do that:
 - a. Select an AWS Region and click **Edit**.
 - b. In the **Edit Backup Vault** window, from the **Backup vault** drop-down list, select the necessary backup vault.

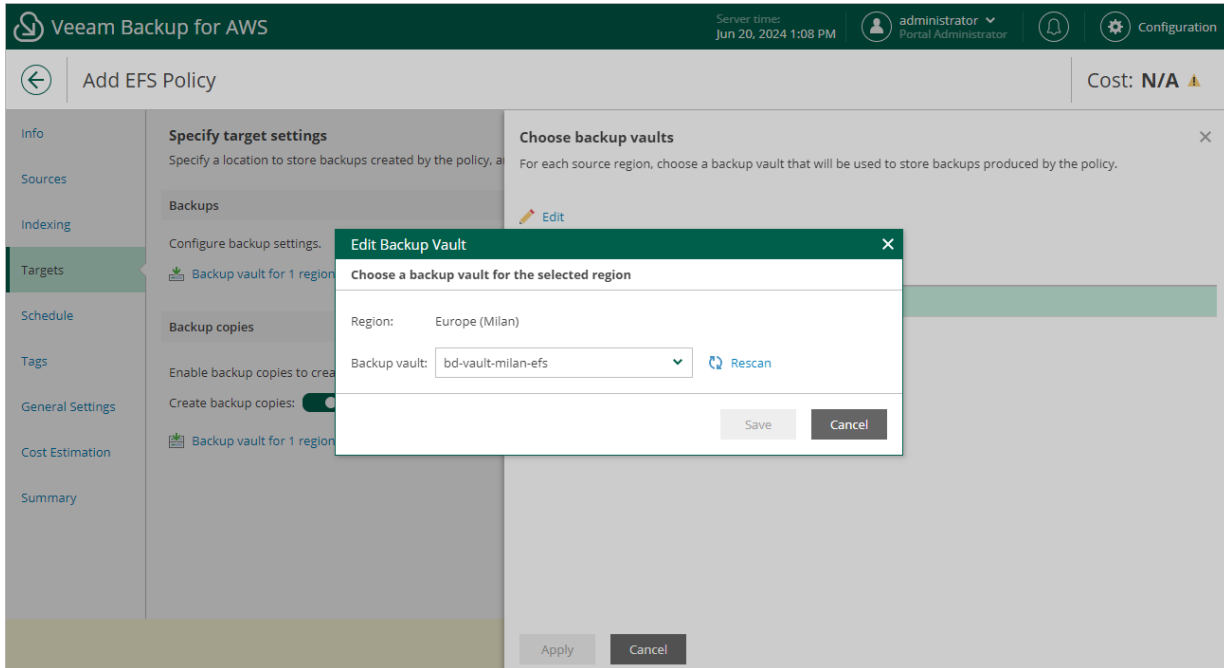
For a backup vault to be displayed in the **Backup vault** list, it must be created in the AWS Backup console as described in [AWS Documentation](#). If you have not created a backup vault for the selected AWS Region, Veeam Backup for AWS will display only the default backup vault existing in this region.

IMPORTANT

Consider the following:

- Veeam Backup for AWS does not support storing backups in [logically air gapped vaults](#).
 - Make sure policies assigned to the selected backup vault allow Veeam Backup for AWS to access vault resources and to perform backup, backup copy and restore operations. For more information on vault access policies, see [AWS Documentation](#).
- c. Click **Save**.

3. To save changes made to the backup policy settings, click **Apply**.



Enabling Additional Backup Copy

If you want to copy EFS file system backups to other AWS Regions, do the following:

1. In the **Backup copies** section of the **Targets** step of the wizard, set the **Create backup copies** toggle to *On*.
2. In the **Choose backup vaults** window, configure the following mapping settings for each AWS Region where original file systems reside:
 - a. Select a source AWS Region in the list and click **Edit Region Mapping**.
 - b. In the **Edit Region Mapping** window, specify the following settings:
 - i. From the **Target region** drop-down list, select the target AWS Region to which Veeam Backup for AWS must copy created backups of the selected file systems.
 - ii. From the **Backup vault** drop-down list, select a backup vault that will be used to store the copied backups.

For a backup vault to be displayed in the **Backup vault** list, it must be created in the AWS Backup console as described in [AWS Documentation](#). If you have not created a backup vault for the selected AWS Region, Veeam Backup for AWS will display only the default backup vault existing in this region.

IMPORTANT


Consider the following:

- Veeam Backup for AWS does not support storing backups in [logically air gapped vaults](#).
- Make sure policies assigned to the selected backup vault allow Veeam Backup for AWS to access vault resources and to perform backup, backup copy and restore operations. For more information on vault access policies, see [AWS Documentation](#).

iii. Click **Save**.

To configure mapping for all source AWS Regions at once, click **Set Mapping for All Regions** and specify settings as described in [step 2.b](#).

c. To save changes made to the backup policy settings, click **Apply**.



Step 6. Specify Policy Scheduling Options

You can instruct Veeam Backup for AWS to start the backup policy automatically according to a specific backup schedule. The backup schedule defines how often data stored in file systems added to the backup policy must be backed up.

To help you implement a comprehensive backup strategy, Veeam Backup for AWS allows you to create schedules of the following types:

- **Daily** – the backup policy will create restore points repeatedly throughout a day on specific days.
- **Weekly** – the backup policy will create restore points once a day on specific days.
- **Monthly** – the backup policy will create restore points once a month on a specific day.
- **Yearly** – the backup policy will create restore points once a year on a specific day.

Combining multiple schedule types together allows you to retain restore points for longer periods of time. For more information, see [Enabling Harmonized Scheduling](#).

NOTE

If you do not specify the backup schedule, after you configure the backup policy, you will need to start it manually to create EFS file system backups. For information on how to start backup policies, see [Starting and Stopping Policies](#).

Specifying Daily Schedule

To create a daily schedule for the backup policy, at the **Schedule** step of the wizard, do the following:

1. Set the **Daily schedule** toggle to *On* and click **Edit Daily Settings**.
2. In the **Create daily schedule** window, select hours when the backup policy must create file system backups and backup copies.

If you want to protect file system data more frequently, you can instruct the backup policy to create multiple backups per hour. To do that, click the link to the right of the **Backups** hour selection area, and specify the number of backups that the backup policy must create within an hour.

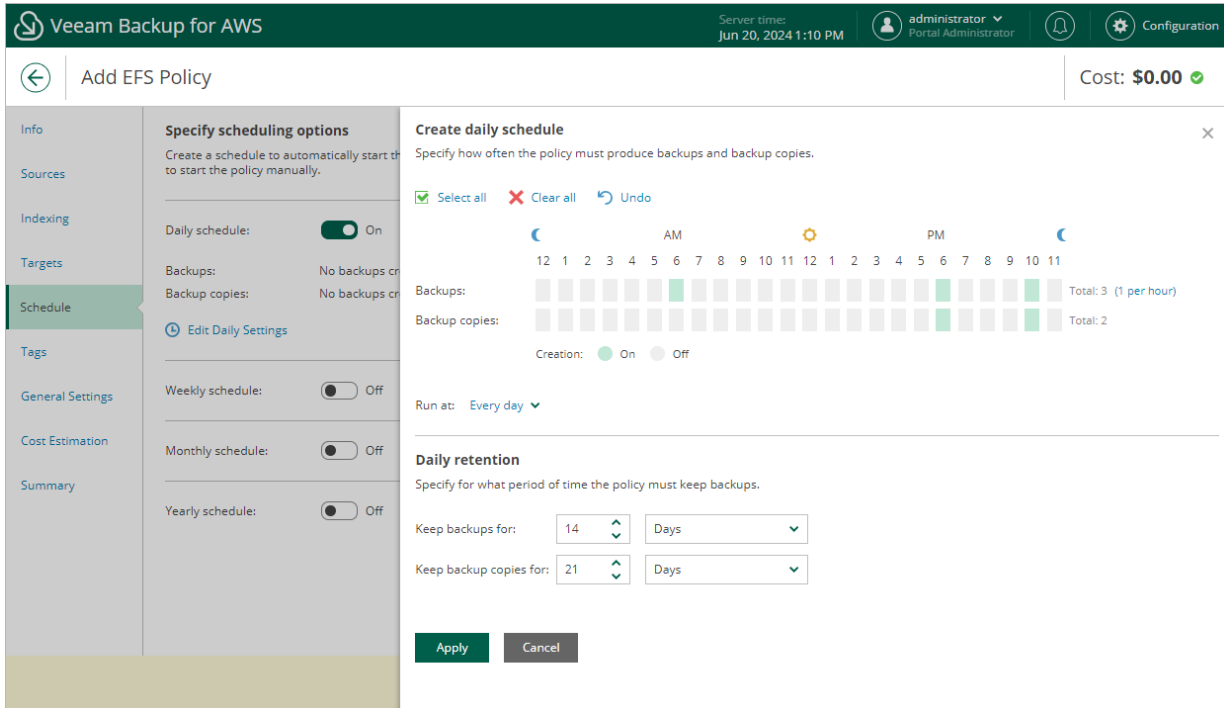
NOTE

Veeam Backup for AWS does not create backup copies independently from file system backups. That is why when you select hours for backup copies, the same hours are automatically selected for backups. To learn how Veeam Backup for AWS performs backup, see [EFS Backup](#).

3. Use the **Run at** drop-down list to choose whether you want the backup policy to run everyday, on work days (Monday through Friday) or on specific days.
4. In the **Daily retention** section, configure retention policy settings for the daily schedule. For backups and backup copies, specify the number of days (or months) for which you want to keep restore points in a backup chain.

If a restore point is older than the specified time limit, Veeam Backup for AWS removes the restore point from the chain. For more information, see [EFS Backup Retention](#).

5. To save changes made to the backup policy settings, click **Apply**.



Specifying Weekly Schedule

To create a weekly schedule for the backup policy, at the **Schedule** step of the wizard, do the following:

1. Set the **Weekly schedule** toggle to *On* and click **Edit Weekly Settings**.
2. In the **Create weekly schedule** window, select weekdays when the backup policy must create file system backups and backup copies.

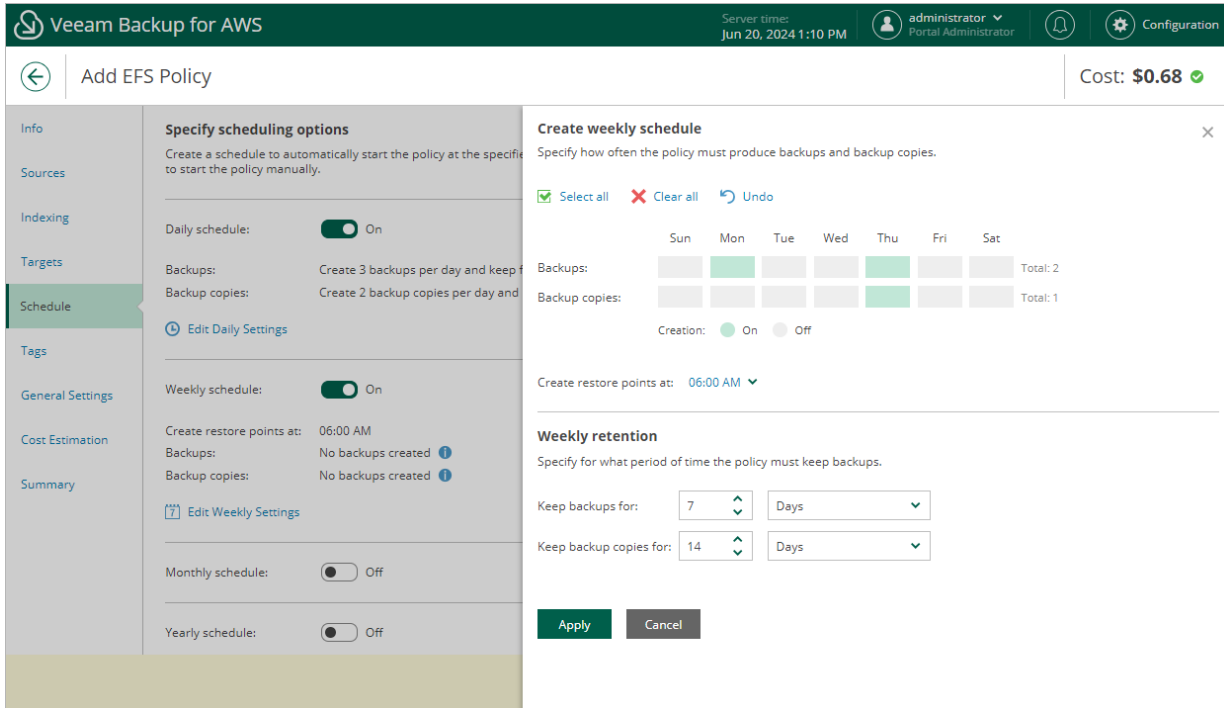
NOTE

Veeam Backup for AWS does not create backup copies independently from file system backups. That is why when you select days to create backup copies, the same days are automatically selected for backups. To learn how Veeam Backup for AWS performs backup, see [EFS Backup](#).

3. Use the **Create restore point at** drop-down list to schedule a specific time for the backup policy to run.
4. In the **Weekly retention** section, configure retention policy settings for the weekly schedule. For backups and backup copies, specify the number of days (or months) for which you want to keep restore points in a backup chain.

If a restore point is older than the specified time limit, Veeam Backup for AWS removes the restore point from the chain. For more information, see [EFS Backup Retention](#).

5. To save changes made to the backup policy settings, click **Apply**.



Specifying Monthly Schedule

To create a monthly schedule for the backup policy, at the **Schedule** step of the wizard, do the following:

1. Set the **Monthly schedule** toggle to *On* and click **Edit Monthly Settings**.
2. In the **Create monthly schedule** window, select months when the backup policy must create file system backups and backup copies.

NOTE

Veeam Backup for AWS does not create backup copies independently from EFS backups. That is why when you select hours for backup copies, the same hours are automatically selected for backups. To learn how Veeam Backup for AWS performs backup, see [EFS Backup](#).

3. Use the **Create restore point at** and **Run on** drop-down lists to schedule a specific time and day for the backup policy to run.

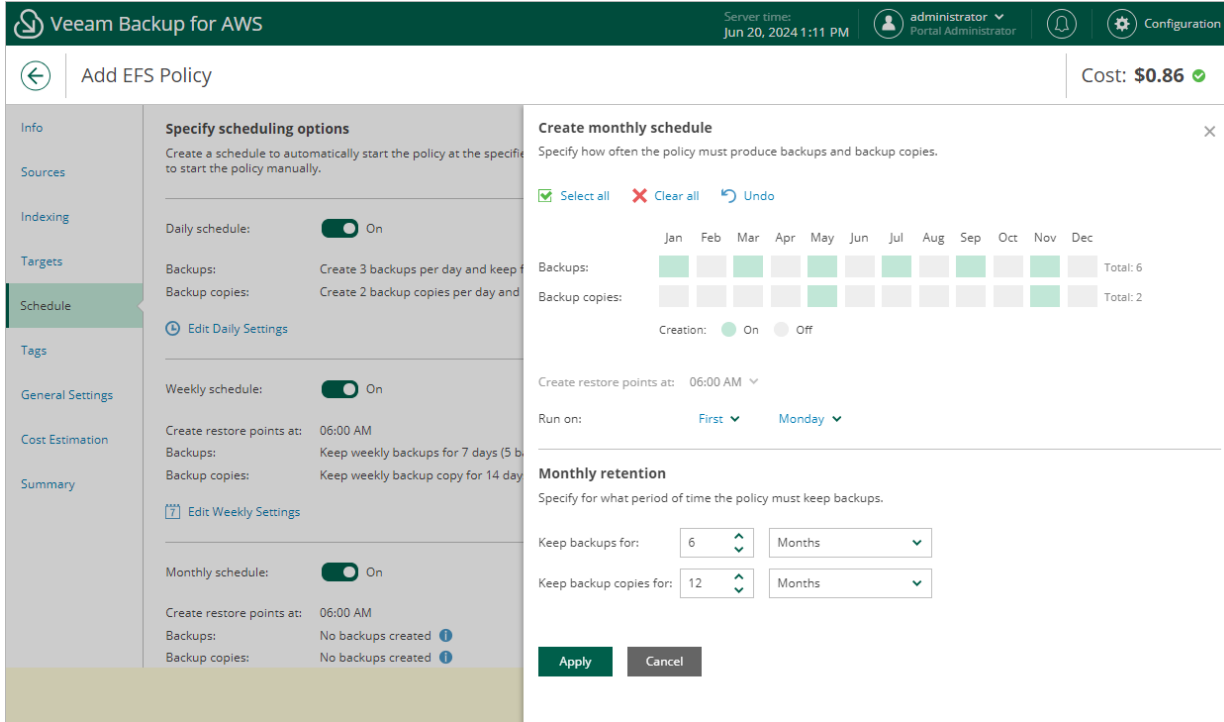
NOTE

Consider the following:

- If you have selected a specific time for the backup policy to run at the **Weekly schedule** section of the **Schedule** step of the wizard, you will not be able to change the time for the monthly schedule unless you select the *On Day* option from the **Run on** drop-down list.
 - If you select the **On day** option, [harmonized scheduling](#) cannot be guaranteed.
4. In the **Monthly retention** section, configure retention policy settings for the monthly schedule. For backups and backup copies, specify the number of days (or months) for which you want to keep restore points in a backup chain.

If a restore point is older than the specified time limit, Veeam Backup for AWS removes the restore point from the chain. For more information, see [EFS Backup Retention](#).

5. To save changes made to the backup policy settings, click **Apply**.



Specifying Yearly Schedule

The yearly schedule is applied only to EFS file system backups, no backup copies are created according to this schedule.

To create a yearly schedule for the backup policy, at the **Schedule** step of the wizard, do the following:

1. Set the **Yearly schedule** toggle to *On* and click **Edit Yearly Settings**.
2. In the **Create yearly schedule** window, specify a day, month and time when the backup policy must create file system backups.

For example, if you select *First, Friday, January* and *06:00 PM*, the backup policy will run every first Friday of January at 06:00 PM.

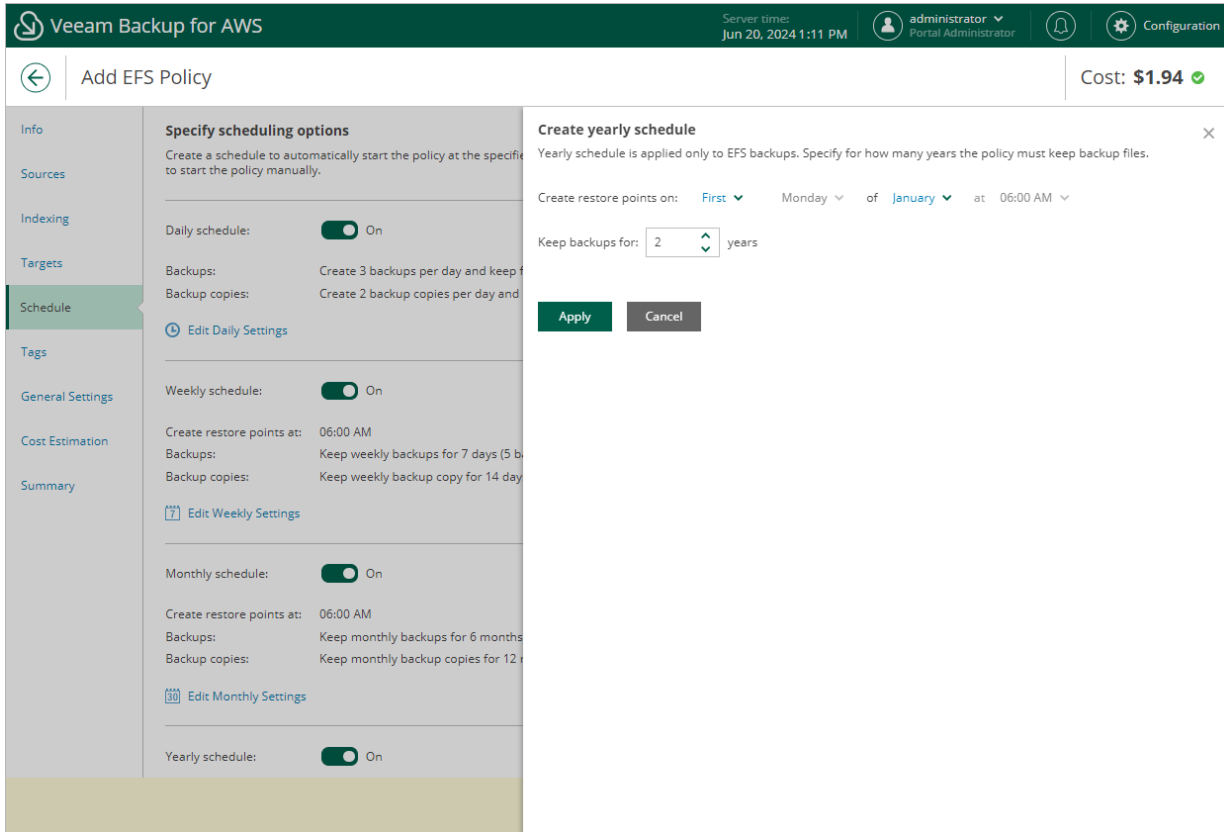
NOTE2

Consider the following:

- If you have selected a specific time and day for the backup policy to run at the **Weekly schedule** or **Monthly schedule** sections of the **Schedule** step of the wizard, you will not be able to change the time and day for the yearly schedule unless you select the *On Day* option from the **Create restore point on** drop-down list.
 - If you select the *On day* option, **harmonized scheduling** cannot be guaranteed.
3. In the **Keep backups for** field, specify the number of years for which you want to keep restore points in a backup chain.

If a restore point is older than the specified time limit, Veeam Backup for AWS removes the restore from the chain. For more information, see [EFS Backup Retention](#).

4. To save changes made to the backup policy settings, click **Apply**.



Enabling Harmonized Scheduling

When you combine multiple types of schedules, Veeam Backup for AWS applies the harmonization mechanism that allows you to leverage restore points for long-term retentions instead of taking a new restore point every time. The mechanism simplifies the backup schedule, optimizes the backup performance and reduces the cost of retaining restore points.

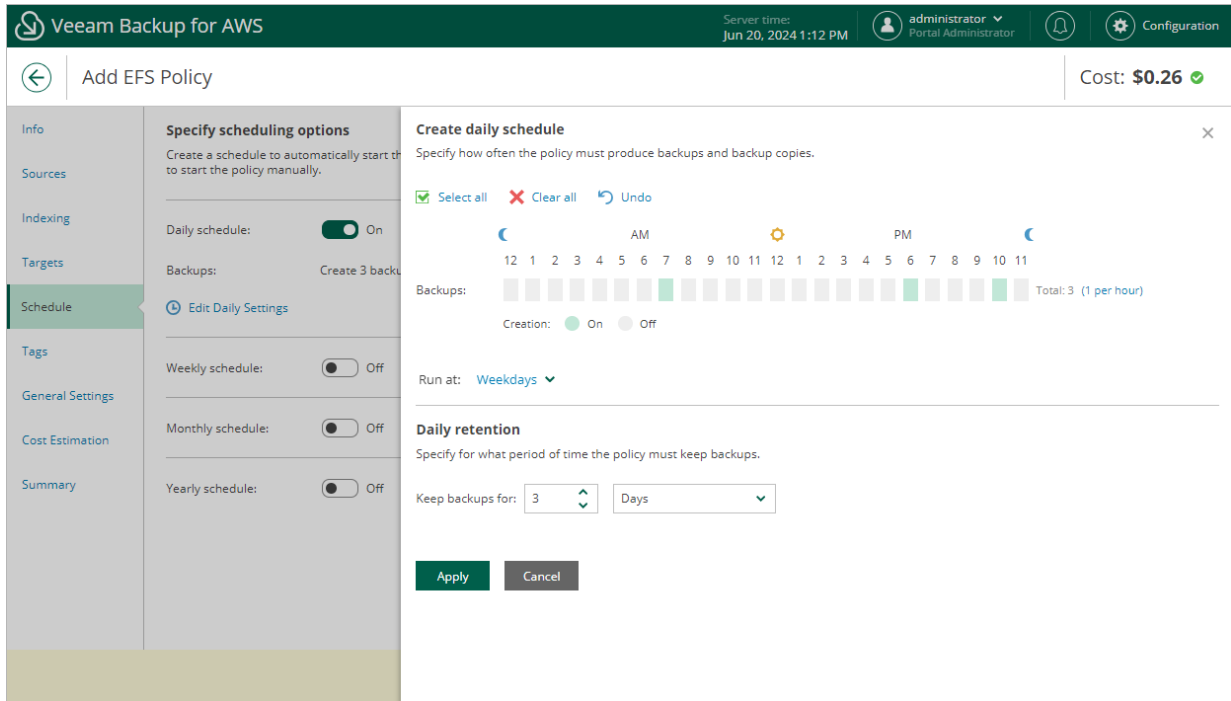
With harmonized scheduling, Veeam Backup for AWS can keep restore points created according to a daily, weekly or monthly schedule for longer periods of time: EFS backups and backup copies can be kept for weeks, months and years.

For Veeam Backup for AWS to use the harmonization mechanism, there must be specified at least 2 different schedules: one schedule will control the regular creation of restore points, while another schedule will control the process of storing restore points. In terms of harmonized scheduling, Veeam Backup for AWS re-uses restore points created according to a more-frequent schedule (daily, weekly or monthly) to achieve the desired retention for less-frequent schedules (weekly, monthly and yearly). Each restore point is marked with a flag of the related schedule type: the (D) flag is used to mark restore points created daily, (W) – weekly, (M) – monthly, and (Y) – yearly. Veeam Backup for AWS uses these flags to control the retention period for the created restore points. Once a flag of a less-frequent schedule is assigned to a restore point, this restore point can no longer be removed – it is kept for the period defined in the retention settings. When the specified retention period is over, the flag is unassigned from the restore point. If the restore point does not have any other flags assigned, it is removed according to the retention settings of a more-frequent schedule.

Consider the following example. You want a backup policy to create backups of your file systems once a day, to keep 3 daily backups in the backup chain, and also to keep one of the created backups for 2 weeks. In this case, you create 2 schedules when configuring the backup policy settings – daily and weekly:


- In the daily scheduling settings, you select hours and days when backups will be created (for example, *7:00 AM; Weekdays*), and specify a number of days for which you want to keep daily restore points in a backup chain (for example, 3).

Veeam Backup for AWS will propagate these settings to the schedule of a lower frequency (which is the weekly schedule in our example).




- In the weekly scheduling settings, you specify which one of the backups created by the daily schedule will be retained for a longer period, and choose for how long you want to keep the selected backup.

For example, if you want to keep the daily restore point created on Monday for 2 weeks, you select *7:00 AM, Monday* and specify 14 days to keep in the weekly schedule settings.




According to the specified scheduling settings, Veeam Backup for AWS will create EFS backups in the following way:

- On the first work day (Monday), a backup session will start at 7:00 AM to create the first restore point. The restore point will be marked with the (D) flag as it was created according to the daily schedule. Since *7:00 AM, Monday* is specified in weekly schedule settings, Veeam Backup for AWS will also assign the (W) flag to this restore point. As a result, 2 flags (D,W) will be assigned to the restore point.
- On the same week, after backup sessions run on Tuesday and Wednesday, the created restore points will be marked with the (D) flag.




- On the fourth work day (Thursday), after a backup session runs at 7:00 AM, the created restore point will be marked with the (D) flag.

By this moment, the earliest restore point in the backup chain will get older than the specified retention limit. However, Veeam Backup for AWS will not remove the earliest restore point (7:00 AM, Monday) with the (D) flag from the backup chain as this restore point is also marked with a flag of a less-frequent schedule. Instead, Veeam Backup for AWS will unassign the (D) flag from the restore point. This restore point will be kept for the retention period specified in the weekly scheduling settings (that is, for 2 weeks).



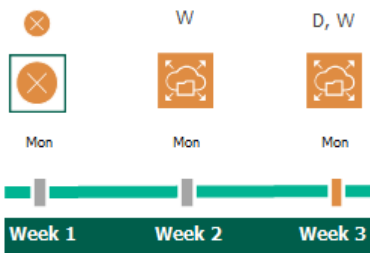
- On the fifth working day (Friday), after a backup session runs at 7:00 AM, the created restore point will be marked with the (D) flag.

By this moment, the restore point created on Tuesday with the (D) flag will get older than the specified retention limit. Veeam Backup for AWS will remove from the backup chain the restore point created at 7:00 AM on Tuesday as no flags of a less-frequent schedule are assigned to this restore point.



- Veeam Backup for AWS will continue creating restore points for the next week in the same way as described in steps 1-4.

- On week 3, after a backup session runs at 7:00 AM on Monday, the earliest weekly restore point in the backup chain will get older than the specified retention limit. Veeam Backup for AWS will unassign the (W) flag from the earliest weekly restore point. Since no other flags are assigned to this restore point, Veeam Backup for AWS will remove this restore point from the backup chain.



Step 7. Enable AWS Tags Assigning

At the **Tags** step of the wizard, choose whether you want to assign AWS tags to backups and backup copies.

- To assign already existing AWS tags from the processed EFS file systems, select the **Copy tags from source file systems** check box.

If you choose to copy tags from the source file systems, Veeam Backup for AWS will first create a backup or backup copy of the EFS file system and assign to the created backup AWS tags with Veeam metadata, then Veeam Backup for AWS will copy tags from the processed file system and, finally, assign the copied AWS tags to the backup.

- To assign your own custom AWS tags, set the **Add custom tags to created backups** toggle to *On* and specify the AWS tags explicitly. To do that, use the **Key** and **Value** fields to specify a key and value for the new custom AWS tag, and then click **Add**. Note that you cannot add more than 5 custom AWS tags.

If you choose to add custom tags to the created snapshots, Veeam Backup for AWS will assign the specified tags right after it creates a backup or backup copy.

The screenshot shows the 'Add EFS Policy' wizard in Veeam Backup for AWS, specifically the 'Tags' step. The interface includes a top navigation bar with the Veeam logo, server time (Jun 20, 2024 1:12 PM), user information (administrator, Portal Administrator), and a Configuration icon. The main header shows 'Add EFS Policy' and a cost of '\$2.10'. A left sidebar lists navigation options: Info, Sources, Indexing, Targets, Schedule, Tags (highlighted), General Settings, Cost Estimation, and Summary. The 'Specify tag settings' section contains a checked checkbox for 'Copy tags from source file systems' and a toggle for 'Add custom tags to created backups' set to 'On'. Below this, there are input fields for 'Key' (containing 'user') and 'Value' (containing 'donna_ortiz'), with an '+ Add' button. A list of existing tags shows 'owner: dept01' with a close icon. A note states 'A maximum of 5 custom tags is allowed.' At the bottom, there are 'Previous', 'Next', and 'Cancel' buttons.

Step 8. Specify General Settings

At the **General Settings** step of the wizard, you can enable automatic retries and specify notification settings for the backup policy.

Automatic Retry Settings

To instruct Veeam Backup for AWS to run the backup policy again if it fails on the first try, do the following:

1. In the **Schedule** section of the step, select the **Automatically retry failed policy** check box.
2. In the field to the right of the check box, specify the maximum number of attempts to run the backup policy. The time interval between retries is 60 seconds.

When retrying backup policies, Veeam Backup for AWS processes only those file systems that failed to be backed up during the previous attempt.

Email Notification Settings

NOTE

To be able to specify email notification settings for the EFS Backup policy, you must configure [global notification settings](#) first.

To instruct Veeam Backup for AWS to send email notifications for the backup policy, do the following:

1. In the **Notifications** section of the step, set the **Enabled** toggle to *On*.
If you set the toggle to *Off*, Veeam Backup for AWS will send notifications according to the configured global notification settings.
2. In the **Email** field, specify an email address of a recipient.
Use a semicolon to separate multiple recipient addresses. Do not use spaces after semicolons between the specified email addresses.
3. Use the **Notify on** list to choose whether you want Veeam Backup for AWS to send email notifications in case the backup policy completes successfully, completes with warnings or completes with errors.
4. Select the **Suppress notifications until the last retry** check box to receive a notification about the final backup policy result.
If you do not select the check box, Veeam Backup for AWS send a notification for every backup policy retry.

NOTE

If you specify the same email recipient in both backup policy notification and [global notification settings](#), Veeam Backup for AWS will override the configured global notification settings and will send each notification to this recipient only once to avoid notification duplicates.

The screenshot shows the 'Add EFS Policy' configuration page in the Veeam Backup for AWS console. The page is titled 'Add EFS Policy' and shows a cost of '\$2.10'. The left sidebar contains navigation options: Info, Sources, Indexing, Targets, Schedule, Tags, General Settings (selected), Cost Estimation, and Summary. The main content area is titled 'Configure retry and notification settings' and includes the following sections:

- Schedule:** A checkbox labeled 'Automatically retry failed policy:' is checked, with a value of '3' in a spinner box and the text 'times'.
- Notifications:** A toggle switch labeled 'Enabled:' is set to 'On'. Below it, an 'Email:' field contains the text 'donna_ortiz@company.mail'. Under the heading 'Notify on', there are three checked checkboxes: 'Failure', 'Warning', and 'Success'. At the bottom of this section, there is a checked checkbox labeled 'Suppress notifications until the last retry'.

At the bottom of the page, there are three buttons: 'Previous', 'Next', and 'Cancel'.

Step 9. Review Estimated Cost

[This step applies only if you have created a schedule for the backup policy at the **Schedule** step of the wizard]

At the **Cost Estimation** step of the wizard, review the estimated monthly cost of AWS services and resources that will be consumed to protect the file systems added to the backup policy. The total estimated cost includes the following:

- The cost of creating backups of the EFS file systems.
For each file system included in the backup policy, Veeam Backup for AWS takes into account the number of restore points to be kept in the backup chain and the configured scheduling settings.
- The cost of creating backup copies and maintaining them in the target AWS Region.
For each file system included in the backup policy, Veeam Backup for AWS takes into account the number of restore points to be kept in the backup chain and the configured scheduling settings.
- The cost of sending API requests to Veeam Backup for AWS during data protection operations.

NOTE

To calculate the estimated cost, Veeam Backup for AWS uses capabilities of the [AWS Pricing Calculator](#) that estimates the cost of services in USD only. This calculator is intended for informational and estimation purposes only.

The estimated cost may occur to be significantly higher due to the backup frequency, cross-region data transfer and AWS backup charges. To reduce the cost, you can try the following workarounds:

- To reduce high AWS backup charges, adjust the backup retention settings to keep less restore points in the backup chain.
- To optimize the cost of storing backups, configure the scheduling settings to run the backup policy less frequently.

TIP

You can save the cost estimation as a .CSV or .XML file. To do that, click **Export to** and select the necessary format.

Server time: Jun 20, 2024 1:13 PM administrator Portal Administrator Configuration

← Add EFS Policy Cost: **\$2.10** ✓

Review cost estimation

The estimated cost takes into account the configured target settings, the specified scheduling options and the number of resources to protect.

Note that Veeam Backup for AWS makes predefined [assumptions](#) to calculate the cost, which means that the results should be used only as an approximation. For more information on cost calculation, see [this Veeam KB article](#).

\$1.18 Backups

\$0.71 Backup Copies

\$0.21 Traffic

Estimated monthly cost: \$2.10

File system

File System ↑	Backup	Backup Copy	Traffic
bd-efs-cleanup-resto...	\$0.24	\$0.14	\$0.04
bd-efs-singapore	\$0.24	\$0.14	\$0.04
bd-efs-singapore	\$0.24	\$0.14	\$0.04
bd-efs-singapore-res...	\$0.24	\$0.14	\$0.04
bd-efs-singapore-res...	\$0.24	\$0.14	\$0.04

Previous Next Cancel

Related Resources

[How AWS Pricing Works](#)

Step 10. Finish Working with Wizard

At the **Summary** step of the wizard, it is recommended that you run the backup policy configuration check before you click **Finish**.

The configuration check will verify whether specified IAM roles have all the required permissions, and networks settings are configured properly to launch worker instances. To run the check, click **Test Configuration**. Veeam Backup for AWS will display the **Test policy configuration** window where you can track the progress and view the results of the check. If the IAM role permissions are insufficient or policy settings are not configured properly, the check will complete with errors, and the list of permissions that must be granted to the IAM role and policy configuration issues will be displayed in the **Test policy configuration** window.

You can grant the missing permissions to the IAM role [using the AWS Management Console](#) or instruct Veeam Backup for AWS to do it.

To let Veeam Backup for AWS grant the missing permissions:

1. In the **Test policy configuration** window, click the **Grant** link.
2. In the **Grant Permissions** window, provide one-time access keys of an IAM user that is authorized to update permissions of IAM roles, and then click **Apply**.

The IAM user must have the following permissions:

```
"iam:AttachRolePolicy",  
"iam:CreatePolicy",  
"iam:CreatePolicyVersion",  
"iam:CreateRole",  
"iam:GetAccountSummary",  
"iam:GetPolicy",  
"iam:GetPolicyVersion",  
"iam:GetRole",  
"iam:ListAttachedRolePolicies",  
"iam:ListPolicyVersions",  
"iam:SimulatePrincipalPolicy",  
"iam:UpdateAssumeRolePolicy"
```

NOTE

Veeam Backup for AWS does not store one-time access keys in the configuration database.

3. After the required permissions are granted, close the **Test policy configuration** window, and then click **Finish** to close the **Add Policy** wizard.

Veeam Backup for AWS will save the configured backup policy.

Server time: Jun 20, 2024 1:13 PM administrator Portal Administrator Configuration

← Add EFS Policy Cost: \$2.10 ✓

Review configured settings
Review the settings, and click Finish to exit the wizard.

Info
Sources
Indexing
Targets
Schedule
Tags
General Settings
Cost Estimation
Summary

General

Name: EFS Backup Policy
Description: Backup of file system for D01
Regions: Asia Pacific (Singapore)
Europe (Frankfurt)
Account: Default Backup Restore

Backup settings

Copy tags from source file systems: Yes
Add custom tags: Yes
Custom tags: owner:dept01

Backup schedule

Daily retention: Create 2 restore points and keep for 14 Days
Weekly retention: Create 2 restore points and keep for 7 Days
Monthly retention: Create 6 restore points and keep for 6 Months
Yearly retention: Create restore point on First Monday of January at 06:00 AM
Keep backups for 2 years

Backup copy settings

Enabled: Yes
Region mapping: Source region: Asia Pacific (Singapore) Target region: Asia Pacific (Sydney)

Previous Finish Cancel

Creating EFS Backups Manually

Veeam Backup for AWS allows you to manually create backups of Amazon EFS file systems. You can instruct Veeam Backup for AWS to store the created backups in the same AWS Regions where the processed file systems reside, or in a different AWS Region.

NOTE

Veeam Backup for AWS does not include EFS backups created manually in the EFS backup chain and does not apply the configured retention policy settings to these backups. This means that the backups are kept in your AWS environment unless you remove them manually, as described in section [Managing Backed-Up EFS Data](#).

To manually create a backup of an EFS file system, do the following:

1. Navigate to **Resources > File Systems > EFS**.

NOTE

By default, Veeam Backup for AWS uses an AWS CloudTrail trail to track changes in your EFS resources. If no trails are configured in the source AWS account, Veeam Backup for AWS will access AWS resources and populate the list of available file systems or AWS tags only once in 24 hours. To force the data collection process manually, click **Rescan**.

2. Select the necessary file system and click **Take Backup Now**.

For an EFS file system to be displayed in the list of available file systems, an AWS Region where the file system resides must be added to any of [configured EFS backup policies](#), and the IAM role specified in the backup policy settings must have permissions to access the file system. For more information on required permissions, see [EFS Backup IAM Role Permissions](#).

3. Complete the **Take Manual Backup** wizard:

- a. At the **Account** step of the wizard, specify an IAM role whose permissions Veeam Backup for AWS will use to create the backup.

For an IAM role to be displayed in the list, it must be added to Veeam Backup for AWS as described in section [Adding IAM Roles](#).

- b. In the **Backup vault** section of the **Settings** step of the wizard, click **Edit Location Settings**.

In the **Choose region and backup vault** window, specify the following settings:

- i. From the **Target region** drop-down list, select an AWS Region where manual backups will be stored.
- ii. In the **Backup vault** section, select a backup vault that will be used to store file system backups.

For a backup vault to be displayed in the list of available vaults, it must be created in the AWS Backup console as described in [AWS Documentation](#). If you have not created a backup vault for the selected AWS Region, Veeam Backup for AWS will display only the default backup vault created for the AWS Region automatically.

- iii. To save changes made to the location settings, click **Apply**.

- c. At the **Tags** section of the **Settings** step of the wizard, if you want to assign tags to the created backup, click **Edit Tag Settings**.

In the **Tag configuration** window, specify tag settings:

- i. To assign already existing AWS tags from the processed file system, select the **Copy tags from source file system** check box.

If you choose to copy tags from source file system, Veeam Backup for AWS will first create a backup of the EFS file system and assign to the created backup AWS tags with Veeam metadata, then Veeam Backup for AWS will copy tags from the processed file system and, finally, assign the copied AWS tags to the backup.

- ii. To assign your own custom AWS tags, set the **Add custom tags to created backup** toggle to *On* and specify the tags explicitly. To do that, use the **Key** and **Value** fields to specify a key and a value for the new custom AWS tag, and then click **Add**. Note that you cannot add more than 5 custom AWS tags.

If you choose to add custom tags to created backups, Veeam Backup for AWS will assign the specified tags right after it creates a backup.

- iii. To save changes made to the tag settings, click **Apply**.

d. At the **Summary** step of the wizard, review summary information and click **Finish**.

The screenshot shows the Veeam Backup for AWS interface. At the top, the header includes the Veeam logo, the text "Veeam Backup for AWS", the server time "Jun 20, 2024 2:52 PM", and the user "administrator Portal Administrator". The main navigation bar shows "Take Manual Backup" with a back arrow. On the left, a sidebar contains "Account", "Settings", and "Summary". The "Configure backup settings" panel is active, showing "Backup vault" (Target region: —, Backup vault: —) and "Tags" (Source tags: Will not be copied, Custom tags: Will not be assigned). A modal window titled "Choose region and backup vault" is open, showing "Target region" set to "Europe (Spain)" and a list of backup vaults: "aws/efs/automatic-backup-vault", "bd-vault-spain" (highlighted), "Default", "jf_redshift_vault", "nm-explicit-deny", "nm-v7-regress", and "nm-vault-check". "Apply" and "Cancel" buttons are at the bottom of the modal.

Performing FSx Backup

One backup policy can be used to process one or more FSx file systems within one AWS account. The scope of data that you can protect in an AWS account is limited by permissions of an IAM role that is specified in the backup policy settings.

NOTE

If you plan to receive email notifications on backup policy results, configure email notification settings before creating an FSx backup policy. For more information, see [Configuring Global Notification Settings](#).

To schedule data protection tasks to run automatically, [create backup policies](#). For each protected FSx file systems, you can also [take a backup manually](#) when needed.

IMPORTANT

Consider the following:

- Veeam Backup for AWS supports backup of FSx file systems only to the same AWS accounts to which the source file systems belong.
- Veeam Backup for AWS does not support backup of Amazon FSx for NetApp ONTAP file systems.
- Veeam Backup for AWS does not support backup of Amazon FSx for Lustre file systems with the [Scratch deployment type](#).
- Veeam Backup for AWS does not support backup of Amazon FSx for Lustre file systems with the [data repository association](#).

Creating FSx Backup Policies

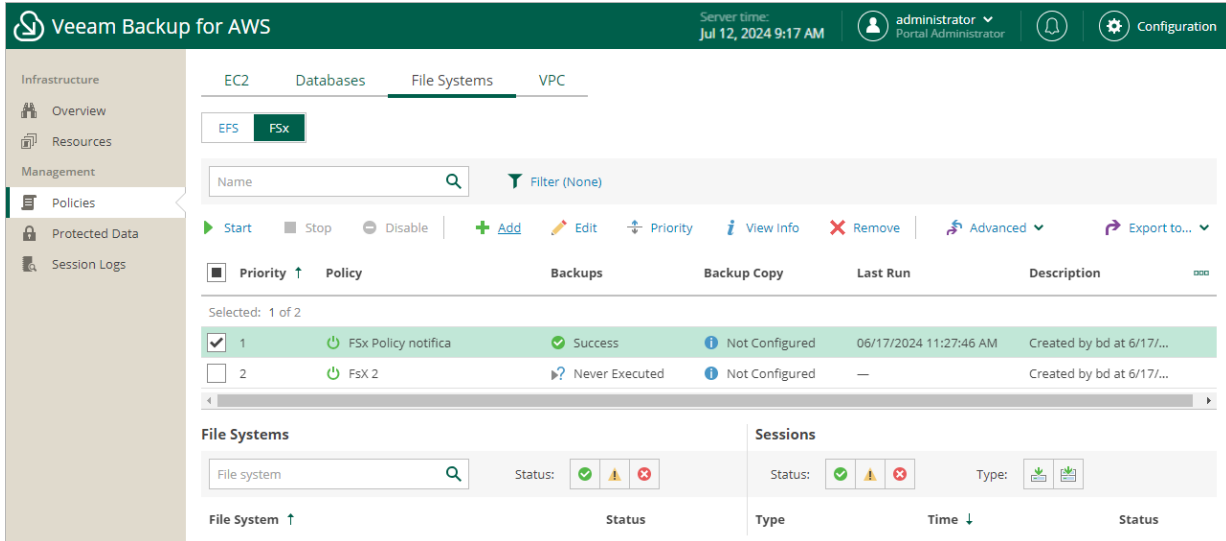
To create a backup policy, do the following:

1. [Launch the Add FSx Policy wizard](#).
2. [Specify a backup policy name and description](#).
3. [Configure backup source settings](#).
4. [Configure backup target settings](#).
5. [Specify a schedule for the backup policy](#).
6. [Enable AWS tags assigning](#).
7. [Specify automatic retry settings and notification settings for the backup policy](#).
8. [Review estimated cost of the selected FSx file systems](#).
9. [Finish working with the wizard](#).

Step 1. Launch Add FSx Policy Wizard

To launch the **Add FSx Policy** wizard, do the following:

1. Navigate to **Policies > File Systems > FSx**.
2. Click **Add**.



Step 2. Specify Policy Name and Description

At the **Info** step of the wizard, use the **Name** and **Description** fields to specify a name for the new backup policy and to provide a description for future reference. The name must be unique in Veeam Backup for AWS; the maximum length of the name is 127 characters, the maximum length of the description is 255 characters.

The screenshot shows the 'Add FSx Policy' wizard in Veeam Backup for AWS. The interface includes a top navigation bar with the Veeam logo, 'Veeam Backup for AWS', server time 'Jul 12, 2024 9:19 AM', user 'administrator Portal Administrator', and a 'Configuration' icon. Below the navigation bar, the wizard title 'Add FSx Policy' is displayed with a back arrow and a 'Cost: N/A' indicator. A left-hand sidebar contains a list of steps: Info (selected), Sources, Targets, Schedule, Tags, General Settings, Cost Estimation, and Summary. The main content area is titled 'Specify policy name and description' and contains the instruction 'Enter a name and description for the policy.' There are two input fields: 'Name:' with the value 'fsx backup policy 01' and 'Description:' with the value 'backup policy for dept-01'. At the bottom of the wizard, there are 'Next' and 'Cancel' buttons.

Step 3. Configure Backup Source Settings

At the **Sources** step of the wizard, specify backup source settings:

1. [Select an IAM role whose permissions will be used to perform FSx file system backup.](#)
2. [Select AWS Regions where FSx file systems that you plan to back up reside.](#)
3. [Select FSx file systems to back up.](#)

Step 3.1 Specify IAM Role

In the **IAM role** section of the **Sources** step of the wizard, specify an IAM role whose permissions will be used to access AWS services and resources, and to create backups of Amazon FSx file systems. The specified IAM role must belong to the AWS account in which the file systems that you want to protect reside, and must be assigned permissions listed in section [FSx Backup IAM Role Permissions](#).

For an IAM role to be displayed in the **IAM role** list, it must be added to Veeam Backup for AWS with the *Amazon FSx Backup* operation selected as described in section [Adding IAM Roles](#). If you have not added the necessary IAM role to Veeam Backup for AWS beforehand, you can do it without closing the **Add FSx Policy** wizard. To do that, click **Add** and complete the [Add IAM Role](#) wizard.

IMPORTANT


It is recommended that you check whether the selected IAM role has all the permissions required to perform the operation. If some permissions of the IAM role are missing, the backup policy will fail to complete successfully. To run the IAM role permission check, click **Check Permissions** and follow the instructions provided in section [Checking IAM Role Permissions](#).

The screenshot shows the 'Add FSx Policy' wizard in the Veeam Backup for AWS interface. The top navigation bar includes the Veeam logo, 'Veeam Backup for AWS', server time 'Jul 12, 2024 9:21 AM', user 'administrator Portal Administrator', and 'Configuration' settings. The main header shows a back arrow, 'Add FSx Policy', and 'Cost: N/A'. A left sidebar contains navigation options: Info, Sources (selected), Targets, Schedule, Tags, General Settings, Cost Estimation, and Summary. The main content area is titled 'Specify source settings' and includes instructions: 'Select an IAM role to use, regions to cover and resources to process by the policy. Using tags provides dynamic selection that automatically changes the policy scope when tags are assigned to file systems.' Below this, there are sections for 'IAM role' (with a dropdown menu showing 'Default Backup Restore (Default Backup Restore)' and buttons for '+ Add' and 'Check Permissions'), 'Regions' (with a 'Choose regions...' button), and 'Resources' (with buttons for 'Choose resources to protect...' and 'Choose resources to exclude...'). At the bottom, there are 'Previous', 'Next', and 'Cancel' buttons.

Step 3.2 Select AWS Regions

In the **Regions** section of the **Sources** step of the wizard, select AWS Regions where FSx file systems that you plan to back up reside:

1. Click **Choose regions**.
2. In the **Choose regions** window, select the necessary AWS Regions from the **Available Regions** list, and click **Add**.
3. To save changes made to the backup policy settings, click **Apply**.



Step 3.3 Select FSx File Systems

In the **Resources** section of the **Sources** step of the wizard, specify the backup scope – select FSx file systems that Veeam Backup for AWS will back up:

1. Click **Choose resources to protect**.
2. In the **Choose resources to protect** window, choose whether you want to back up all FSx file systems from AWS Regions selected at [step 3.2](#) of the wizard, or only specific file systems.

If you select the **All resources** option, Veeam Backup for AWS will regularly check for new FSx file systems reside in the selected regions and automatically update the backup policy settings to include these file systems into the backup scope.

If you select the **Protect only following resources** option, you must specify the FSx file systems explicitly:

- a. Use the **Type** drop-down list to choose whether you want to add individual file systems or AWS tags to the backup scope.

If you select the *Tag* option, Veeam Backup for AWS will back up only those file systems that reside in the selected AWS Regions under specific AWS tags.

- b. Use the search field of the **Name or ID** drop-down list to find the necessary resource, and then click **Protect** to add the resource to the backup scope.

For a resource to be displayed in the list of available resources, it must reside in an AWS Region that has ever been specified in any backup policy. Otherwise, the only option to discover the available resources is to click **Browse to select specific resources from the global list** and to wait for Veeam Backup for AWS to populate the resource list.

TIP

You can simultaneously add multiple resources to the backup scope. To do that, click **Browse to select specific sources from the global list**, select check boxes next to the necessary FSx file systems or AWS tags in the list of available resources, and then click **Protect**.

If the list does not show the resources that you want to back up, click **Rescan** to launch the data collection process. As soon as the process is over, Veeam Backup for AWS will update the resource list.

If you add an AWS tag to the backup scope, Veeam Backup for AWS will regularly check for new Amazon FSx file systems assigned the added AWS tag and automatically update the backup policy settings to include these file systems in the scope. However, this applies only to file systems from the AWS Regions selected at [step 3.2](#) of the wizard. If you select a tag assigned to file systems from other regions, these file systems will not be protected by the backup policy. To work around the issue, either go back to [step 3.2](#) and add the missing regions, or create a new backup policy.

4. To save changes made to the backup policy settings, click **Apply**.

TIP

As an alternative to selecting the **Protect only following resources** option and specifying the resources explicitly, you can select the **All resources** option and exclude a number of resources from the backup scope. To do that, click **Choose resources to exclude** and specify the file system or tags that you do not want to protect – the procedure is the same as described for including resources in the backup scope.

Note that if a resource appears both in the list of included and excluded resources, Veeam Backup for AWS will still not process the resource because the list of excluded resources has a higher priority.

The screenshot shows the 'Add FSx Policy' configuration window in Veeam Backup for AWS. The 'Choose resources to protect' dialog is open, with the 'Protect only following resources' radio button selected. The 'Type' is set to 'FSx' and the 'Name or ID' is 'jf-test (fs-0e7cc4d400b62a08d)'. A 'Protect' button is visible next to the resource name. Below, a table lists the protected resources:

Item	ID	Value	Region
Selected: 0 of 1			
jf-test	fs-0e7cc4d400b62a08d	—	Europe (Stockholm)

At the bottom of the dialog, there are 'Apply' and 'Cancel' buttons.

Step 4. Configure Backup Target Settings

By default, backup policies create only backups of processed FSx file systems. At the **Targets** step of the wizard, you can specify the following backup target settings:

- [Specify backup vaults where Veeam Backup for AWS will store FSx file system backups.](#)
- [Instruct Veeam Backup for AWS to copy FSx file system backups to other AWS Regions.](#)

Configuring Backup Settings

To specify backup vaults used to store backups of the selected FSx file systems, do the following:

1. In the **Backups** section of the **Targets** step of the wizard, click **Choose backup vaults**.
2. In the **Choose backup vaults** window, for each AWS Region included in the policy, specify a backup vault to save and organize file system backups. To do that:
 - a. Select an AWS Region and click **Edit**.
 - b. In the **Edit Backup Vault** window, from the **Backup vault** drop-down list, select the necessary backup vault.


For a backup vault to be displayed in the **Backup vault** list, it must be created in the AWS Backup console as described in [AWS Documentation](#). If you have not created a backup vault for the selected AWS Region, Veeam Backup for AWS will display only the default backup vault existing in this region.

IMPORTANT

Consider the following:

- Veeam Backup for AWS does not support storing backups in [logically air gapped vaults](#).
 - Make sure that the type of file system added to the backup scope is supported by the AWS Backup service in the source AWS Region. Otherwise, the backup operation will fail to complete successfully. For the list of supported AWS Regions, see [AWS Documentation](#).
 - Make sure policies assigned to the selected backup vault allow Veeam Backup for AWS to access vault resources and to perform backup, backup copy and restore operations. For more information on vault access policies, see [AWS Documentation](#).
 - For Veeam Backup for AWS to be able to back up FSx file systems, you must enable the Opt-in service for the FSx resource type in the AWS Backup settings. Otherwise, Veeam Backup for AWS will automatically enable the service for each AWS Region specified in the **Backups** section in your AWS account while performing backup operations.
- c. Click **Save**.

3. To save changes made to the backup policy settings, click **Apply**.



Enabling Additional Backup Copy

[This step applies only if you have selected file systems resided in default AWS Regions at the **Sources** step of the wizard. There is a limitation on the AWS Backup service side – cross-region copying of FSx backups is not supported for [opt-in Regions](#)]

If you want to copy FSx file system backups to other AWS Regions, do the following:

1. In the **Backup copies** section of the **Targets** step of the wizard, set the **Create backup copies** toggle to *On*.
2. In the **Choose backup vaults** window, configure the following mapping settings for each AWS Region where original file systems reside:
 - a. Select a source AWS Region in the list and click **Edit Region Mapping**.
 - b. In the **Edit Region Mapping** window, specify the following settings:
 - i. From the **Target region** drop-down list, select the target AWS Region to which Veeam Backup for AWS must copy created backups of the selected file systems. The region list shows only the default AWS Regions (that is, the AWS Regions activated for the AWS account by default).
 - ii. From the **Backup vault** drop-down list, select a backup vault that will be used to store the copied backups.

For a backup vault to be displayed in the **Backup vault** list, it must be created in the AWS Backup console as described in [AWS Documentation](#). If you have not created a backup vault for the selected AWS Region, Veeam Backup for AWS will display only the default backup vault existing in this region.

IMPORTANT

Consider the following:

- Veeam Backup for AWS does not support storing backups in [logically air gapped vaults](#).
- Since cross-region copying of FSx backups is not supported for opt-in Regions, the list of available regions will depend on the original location of the selected file system. If the original location is a default AWS Region, the **Target region** list will contain all default AWS Regions; if the original location is an opt-in Region, there is no possibility to specify any AWS Regions.
- Make sure policies assigned to the selected backup vault allow Veeam Backup for AWS to access vault resources and to perform backup, backup copy and restore operations. For more information on vault access policies, see [AWS Documentation](#).
- For Veeam Backup for AWS to be able to back up FSx file systems, you must enable the Opt-in service for the FSx resource type in the AWS Backup settings. Otherwise, Veeam Backup for AWS will automatically enable the service for each AWS Region specified in the **Backup copies** section in your AWS account while performing backup operations.

iii. Click **Save**.

To configure mapping for all source AWS Regions at once, click **Set Mapping for All Regions** and specify settings as described in [step 2.b](#).

c. To save changes made to the backup policy settings, click **Apply**.

The screenshot displays the Veeam Backup for AWS console interface. The main window is titled 'Add FSx Policy' and shows various configuration sections. A modal window titled 'Edit Region Mapping' is open, prompting the user to 'Choose a target region and backup vault for the selected source region'. The modal shows the source region as 'Europe (Stockholm)', the target region as 'Europe (Paris)', and the backup vault as 'pi-parisvault'. A 'Rescan' button is visible next to the backup vault dropdown. A message at the bottom of the modal states: 'The AWS backup service opt-in setting must be enabled for the selected region. If required, it will be automatically enabled. For more information, see the User Guide.' The 'Save' button is highlighted with a mouse cursor. In the background, a table shows the mapping for other source regions: Europe (Milan) is unsupported, and Europe (Stockholm) is mapped to Europe (Paris) with the backup vault 'bd-paris-vault'.

Source Region	Target Region	Backup Vault
Europe (Milan)	Unsupported	Unsupported
Europe (Stockholm)	Europe (Paris)	bd-paris-vault

Step 5. Specify Policy Scheduling Options

You can instruct Veeam Backup for AWS to start the backup policy automatically according to a specific backup schedule. The backup schedule defines how often data stored in file systems added to the backup policy must be backed up.

NOTE

If some of the file systems that you plan to protect have daily automatic backup enabled or a weekly maintenance window configured, you must take into account that automatic backup and maintenance window have a higher priority. If you schedule a backup policy to run approximately at the same time when automatic backup or maintenance window starts, the backup policy will be queued.

To help you implement a comprehensive backup strategy, Veeam Backup for AWS allows you to create schedules of the following types:

- **Daily** – the backup policy will create restore points repeatedly throughout a day on specific days.
- **Weekly** – the backup policy will create restore points once a day on specific days.
- **Monthly** – the backup policy will create restore points once a month on a specific day.
- **Yearly** – the backup policy will create restore points once a year on a specific day.

Combining multiple schedule types together allows you to retain restore points for longer periods of time. For more information, see [Enabling Harmonized Scheduling](#).

TIP

If you do not specify the backup schedule, after you configure the backup policy, you will need to start it manually to create FSx file system backups. For information on how to start backup policies, see [Starting and Stopping Policies](#).

Specifying Daily Schedule

To create a daily schedule for the backup policy, at the **Schedule** step of the wizard, do the following:

1. Set the **Daily schedule** toggle to *On* and click **Edit Daily Settings**.
2. In the **Create daily schedule** window, select hours when the backup policy must create file system backups and backup copies.

If you want to protect file system data more frequently, you can instruct the backup policy to create multiple backups per hour. To do that, click the link to the right of the **Backups** hour selection area, and specify the number of backups that the backup policy must create within an hour.

NOTE

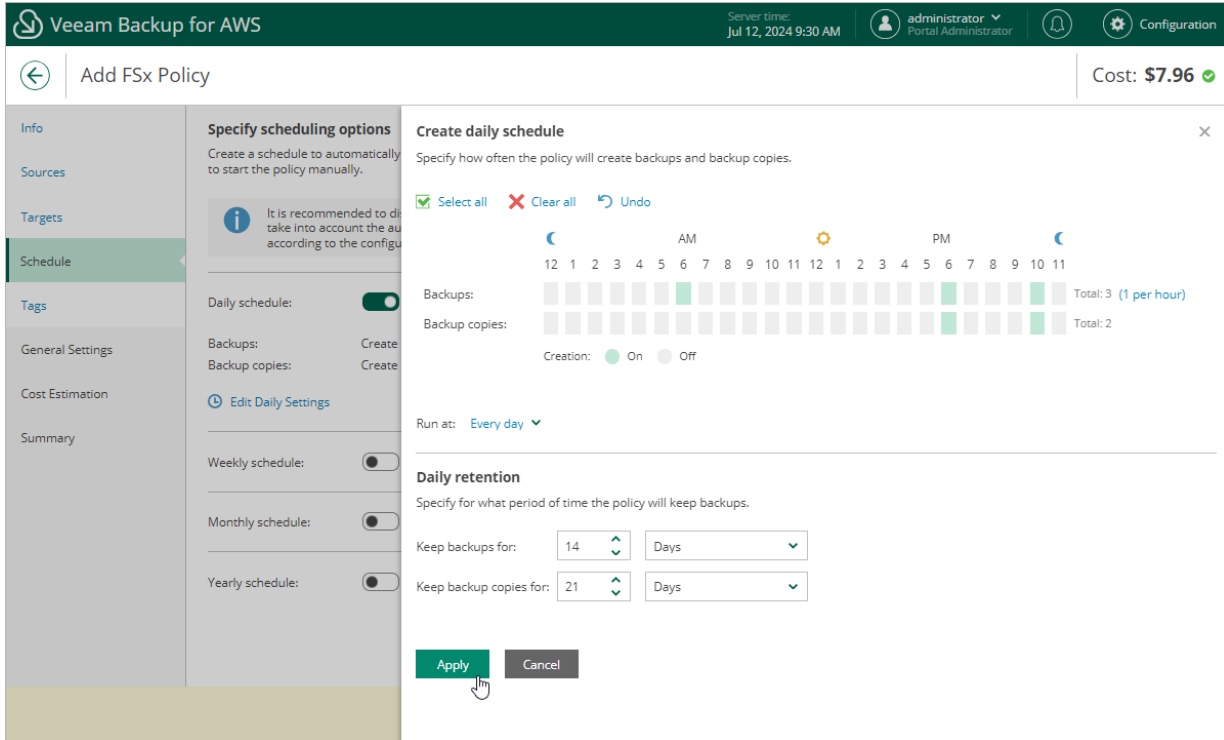
Veeam Backup for AWS does not create backup copies independently from file system backups. That is why when you select hours for backup copies, the same hours are automatically selected for backups. To learn how Veeam Backup for AWS performs backup, see [FSx Backup](#).

3. Use the **Run at** drop-down list to choose whether you want the backup policy to run everyday, on work days (Monday through Friday) or on specific days.

- In the **Daily retention** section, configure retention policy settings for the daily schedule. For backups and backup copies, specify the number of days (or months) for which you want to keep restore points in a backup chain.

If a restore point is older than the specified time limit, Veeam Backup for AWS removes the restore point from the chain. For more information, see [FSx Backup Retention](#).

- To save changes made to the backup policy settings, click **Apply**.



Specifying Weekly Schedule

To create a weekly schedule for the backup policy, at the **Schedule** step of the wizard, do the following:

- Set the **Weekly schedule** toggle to *On* and click **Edit Weekly Settings**.
- In the **Create weekly schedule** window, select weekdays when the backup policy must create file system backups and backup copies.


NOTE

Veeam Backup for AWS does not create backup copies independently from file system backups. That is why when you select days to create backup copies, the same days are automatically selected for backups. To learn how Veeam Backup for AWS performs backup, see [FSx Backup](#).

- Use the **Create restore point at** drop-down list to schedule a specific time for the backup policy to run.
- In the **Weekly retention** section, configure retention policy settings for the weekly schedule. For backups and backup copies, specify the number of days (or months) for which you want to keep restore points in a backup chain.

If a restore point is older than the specified time limit, Veeam Backup for AWS removes the restore point from the chain. For more information, see [FSx Backup Retention](#).

5. To save changes made to the backup policy settings, click **Apply**.



Specifying Monthly Schedule

To create a monthly schedule for the backup policy, at the **Schedule** step of the wizard, do the following:

1. Set the **Monthly schedule** toggle to *On* and click **Edit Monthly Settings**.
2. In the **Create monthly schedule** window, select months when the backup policy must create file system backups and backup copies.

NOTE

Veeam Backup for AWS does not create backup copies independently from FSx backups. That is why when you select hours for backup copies, the same hours are automatically selected for backups. To learn how Veeam Backup for AWS performs backup, see [FSx Backup](#).

3. Use the **Create restore point at** and **Run on** drop-down lists to schedule a specific time and day for the backup policy to run.


NOTE

Consider the following:

- If you have selected a specific time for the backup policy to run at the **Weekly schedule** section of the **Schedule** step of the wizard, you will not be able to change the time for the monthly schedule unless you select the *On Day* option from the **Run on** drop-down list.
 - If you select the **On day** option, [harmonized scheduling](#) cannot be guaranteed.
4. In the **Monthly retention** section, configure retention policy settings for the monthly schedule. For backups and backup copies, specify the number of days (or months) for which you want to keep restore points in a backup chain.

If a restore point is older than the specified time limit, Veeam Backup for AWS removes the restore point from the chain. For more information, see [FSx Backup Retention](#).

5. To save changes made to the backup policy settings, click **Apply**.



Specifying Yearly Schedule

The yearly schedule is applied only to FSx file system backups, no backup copies are created according to this schedule.

To create a yearly schedule for the backup policy, at the **Schedule** step of the wizard, do the following:

1. Set the **Yearly schedule** toggle to *On* and click **Edit Yearly Settings**.
2. In the **Create yearly schedule** window, specify a day, month and time when the backup policy must create file system backups.

For example, if you select *First, Friday, January* and *06:00 PM*, the backup policy will run every first Friday of January at 06:00 PM.


NOTE2

Consider the following:

- If you have selected a specific time and day for the backup policy to run at the **Weekly schedule** or **Monthly schedule** sections of the **Schedule** step of the wizard, you will not be able to change the time and day for the yearly schedule unless you select the *On Day* option from the **Create restore point on** drop-down list.
 - If you select the *On day* option, **harmonized scheduling** cannot be guaranteed.
3. In the **Keep backups for** field, specify the number of years for which you want to keep restore points in a backup chain.

If a restore point is older than the specified time limit, Veeam Backup for AWS removes the restore from the chain. For more information, see [FSx Backup Retention](#).

4. To save changes made to the backup policy settings, click **Apply**.



Enabling Harmonized Scheduling

When you combine multiple types of schedules, Veeam Backup for AWS applies the harmonization mechanism that allows you to leverage restore points for long-term retentions instead of taking a new restore point every time. The mechanism simplifies the backup schedule, optimizes the backup performance and reduces the cost of retaining restore points.

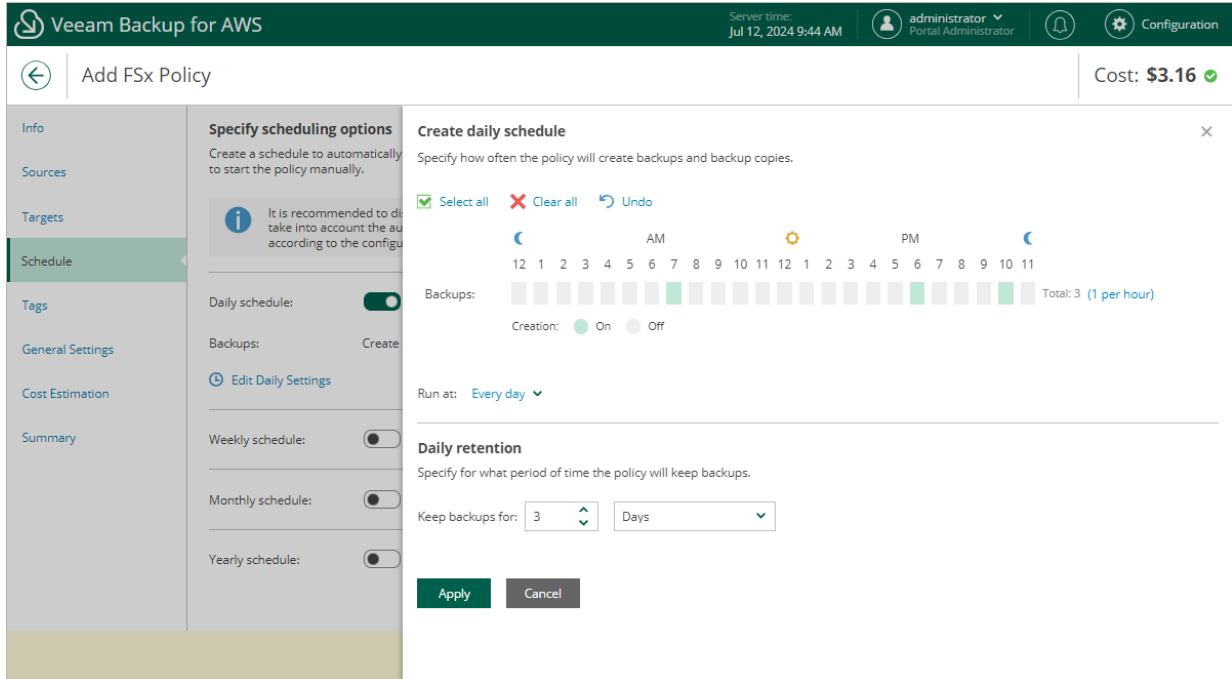
With harmonized scheduling, Veeam Backup for AWS can keep restore points created according to a daily, weekly or monthly schedule for longer periods of time: FSx backups and backup copies can be kept for weeks, months and years.

For Veeam Backup for AWS to use the harmonization mechanism, there must be specified at least 2 different schedules: one schedule will control the regular creation of restore points, while another schedule will control the process of storing restore points. In terms of harmonized scheduling, Veeam Backup for AWS re-uses restore points created according to a more-frequent schedule (daily, weekly or monthly) to achieve the desired retention for less-frequent schedules (weekly, monthly and yearly). Each restore point is marked with a flag of the related schedule type: the (D) flag is used to mark restore points created daily, (W) – weekly, (M) – monthly, and (Y) – yearly. Veeam Backup for AWS uses these flags to control the retention period for the created restore points. Once a flag of a less-frequent schedule is assigned to a restore point, this restore point can no longer be removed – it is kept for the period defined in the retention settings. When the specified retention period is over, the flag is unassigned from the restore point. If the restore point does not have any other flags assigned, it is removed according to the retention settings of a more-frequent schedule.

Consider the following example. You want a backup policy to create backups of your file systems once a day, to keep 3 daily backups in the backup chain, and also to keep one of the created backups for 2 weeks. In this case, you create 2 schedules when configuring the backup policy settings – daily and weekly:

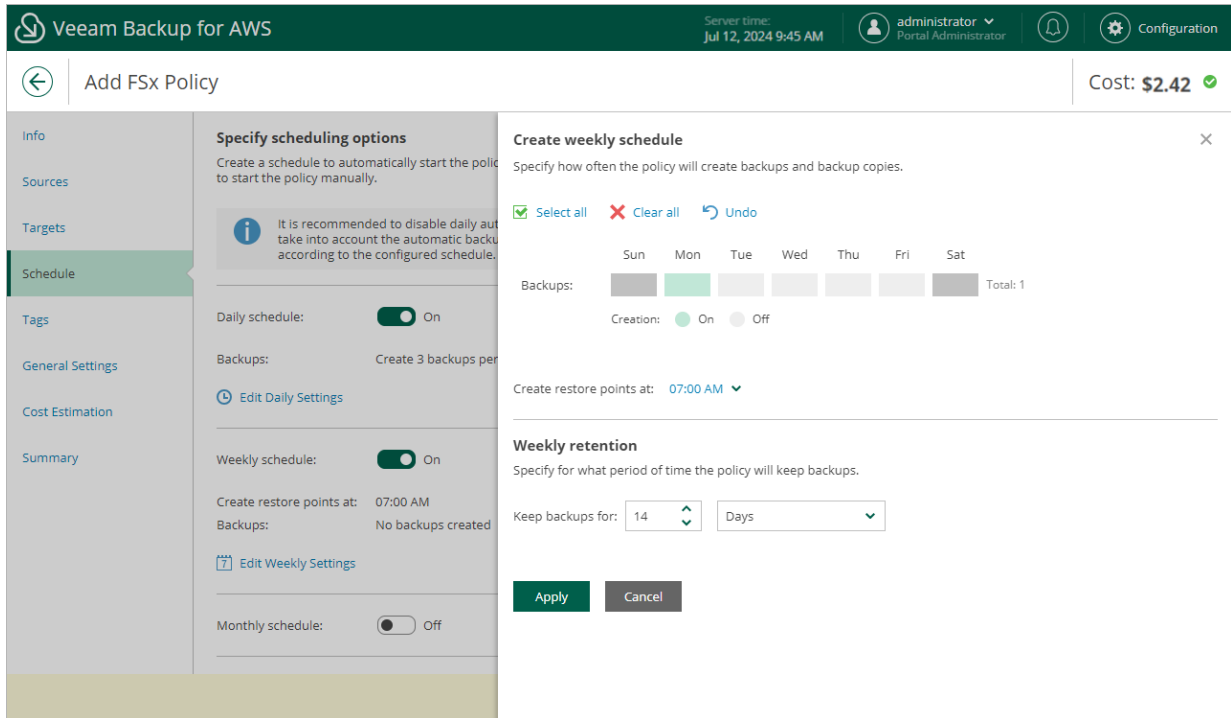
- In the daily scheduling settings, you select hours and days when backups will be created (for example, *7:00 AM; Weekdays*), and specify a number of days for which you want to keep daily restore points in a backup chain (for example, 3).

Veeam Backup for AWS will propagate these settings to the schedule of a lower frequency (which is the weekly schedule in our example).



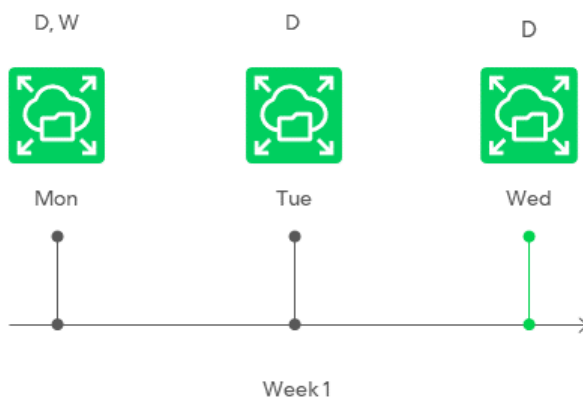
- In the weekly scheduling settings, you specify which one of the backups created by the daily schedule will be retained for a longer period, and choose for how long you want to keep the selected backup.

For example, if you want to keep the daily restore point created on Monday for 2 weeks, you select *7:00 AM, Monday* and specify 14 days to keep in the weekly schedule settings.




According to the specified scheduling settings, Veeam Backup for AWS will create FSx backups in the following way:

- On the first work day (Monday), a backup session will start at 7:00 AM to create the first restore point. The restore point will be marked with the (D) flag as it was created according to the daily schedule. Since *7:00 AM, Monday* is specified in weekly schedule settings, Veeam Backup for AWS will also assign the (W) flag to this restore point. As a result, 2 flags (D,W) will be assigned to the restore point.
- On the same week, after backup sessions run on Tuesday and Wednesday, the created restore points will be marked with the (D) flag.




- On the fourth work day (Thursday), after a backup session runs at 7:00 AM, the created restore point will be marked with the (D) flag.

By this moment, the earliest restore point in the backup chain will get older than the specified retention limit. However, Veeam Backup for AWS will not remove the earliest restore point (7:00 AM, Monday) with the (D) flag from the backup chain as this restore point is also marked with a flag of a less-frequent schedule. Instead, Veeam Backup for AWS will unassign the (D) flag from the restore point. This restore point will be kept for the retention period specified in the weekly scheduling settings (that is, for 2 weeks).




- On the fifth working day (Friday), after a backup session runs at 7:00 AM, the created restore point will be marked with the (D) flag.

By this moment, the restore point created on Tuesday with the (D) flag will get older than the specified retention limit. Veeam Backup for AWS will remove from the backup chain the restore point created at 7:00 AM on Tuesday as no flags of a less-frequent schedule are assigned to this restore point.



- Veeam Backup for AWS will continue creating restore points for the next week in the same way as described in steps 1-4.

6. On week 3, after a backup session runs at 7:00 AM on Monday, the earliest weekly restore point in the backup chain will get older than the specified retention limit. Veeam Backup for AWS will unassign the (W) flag from the earliest weekly restore point. Since no other flags are assigned to this restore point, Veeam Backup for AWS will remove this restore point from the backup chain.



Step 6. Enable AWS Tags Assigning

At the **Tags** step of the wizard, choose whether you want to assign AWS tags to backups and backup copies.

- To assign already existing AWS tags from the processed FSx file systems, select the **Copy tags from source file systems** check box.

If you choose to copy tags from the source file systems, Veeam Backup for AWS will first create a backup or backup copy of the FSx file system and assign to the created backup AWS tags with Veeam metadata, then Veeam Backup for AWS will copy tags from the processed file system and, finally, assign the copied AWS tags to the backup.

- To assign your own custom AWS tags, set the **Add custom tags to created backups** toggle to *On* and specify the AWS tags explicitly. To do that, use the **Key** and **Value** fields to specify a key and value for the new custom AWS tag, and then click **Add**. Note that you cannot add more than 5 custom AWS tags.

If you choose to add custom tags to the created snapshots, Veeam Backup for AWS will assign the specified tags right after it creates a backup or backup copy.

The screenshot shows the 'Add FSx Policy' wizard in Veeam Backup for AWS. The interface is in a dark theme. At the top, the Veeam logo and 'Veeam Backup for AWS' are on the left, and 'Server time: Jul 12, 2024 9:34 AM', 'administrator Portal Administrator', and 'Configuration' are on the right. The main header shows a back arrow, 'Add FSx Policy', and 'Cost: \$25.21'. A left sidebar contains navigation links: Info, Sources, Targets, Schedule, Tags (highlighted), General Settings, Cost Estimation, and Summary. The main content area is titled 'Specify tag settings' and includes a sub-header: 'You can copy tags from source file systems and additionally assign up to 5 custom tags to backups and backup copies created by the policy. Tags can help you manage, identify, organize, search for, and filter resources.' Below this, there is a checked checkbox for 'Copy tags from source file systems'. A toggle for 'Add custom tags to created backups:' is set to 'On'. There are two input fields: 'Key:' with 'user' and 'Value:' with 'donna_ortiz'. To the right of the 'Value' field is a '+ Add' button with a mouse cursor over it. Below these fields is a list of existing tags: 'owner: dept01'. A note at the bottom states 'A maximum of 5 custom tags is allowed.' At the bottom of the wizard, there are three buttons: 'Previous', 'Next', and 'Cancel'.

Step 7. Specify General Settings

At the **General Settings** step of the wizard, you can enable automatic retries and specify notification settings for the backup policy.

Automatic Retry Settings

To instruct Veeam Backup for AWS to run the backup policy again if it fails on the first try, do the following:

1. In the **Schedule** section of the step, select the **Automatically retry failed policy** check box.
2. In the field to the right of the check box, specify the maximum number of attempts to run the backup policy. The time interval between retries is 60 seconds.

When retrying backup policies, Veeam Backup for AWS processes only those file systems that failed to be backed up during the previous attempt.

Email Notification Settings

NOTE

To be able to specify email notification settings for the FSx Backup policy, you must configure [global notification settings](#) first.

To instruct Veeam Backup for AWS to send email notifications for the backup policy, do the following:

1. In the **Notifications** section of the step, set the **Enabled** toggle to *On*.
If you set the toggle to *Off*, Veeam Backup for AWS will send notifications according to the configured global notification settings.
2. In the **Email** field, specify an email address of a recipient.
Use a semicolon to separate multiple recipient addresses. Do not use spaces after semicolons between the specified email addresses.
3. Use the **Notify on** list to choose whether you want Veeam Backup for AWS to send email notifications in case the backup policy completes successfully, completes with warnings or completes with errors.
4. Select the **Suppress notifications until the last retry** check box to receive a notification about the final backup policy result.
If you do not select the check box, Veeam Backup for AWS send a notification for every backup policy retry.

NOTE

If you specify the same email recipient in both backup policy notification and [global notification settings](#), Veeam Backup for AWS will override the configured global notification settings and will send each notification to this recipient only once to avoid notification duplicates.

The screenshot shows the 'Add FSx Policy' configuration page in the Veeam Backup for AWS console. The page is titled 'Add FSx Policy' and shows a cost of '\$25.61'. The left sidebar contains navigation options: Info, Sources, Targets, Schedule, Tags, General Settings (selected), Cost Estimation, and Summary. The main content area is titled 'Configure retry and notification settings' and includes the following sections:

- Schedule:** A section for configuring the policy schedule.
- Automatic Retries:** A checkbox labeled 'Automatically retry failed policy:' is checked, with a dropdown menu set to '3' times.
- Notification Note:** A yellow callout box states: 'Automatic retry settings are only applicable on a scheduled run of the policy'.
- Notifications:** A section for configuring email notifications.
 - Enabled:** A toggle switch is turned 'On'.
 - Email:** A text input field contains the email address 'donna-ortiz@gmail.com'.
 - Notification Types:** Three checkboxes are checked: 'Failure', 'Warning', and 'Success'.
 - Suppress Notifications:** A checkbox labeled 'Suppress notifications until the last retry' is checked.

At the bottom of the page, there are three buttons: 'Previous', 'Next', and 'Cancel'.

Step 8. Review Estimated Cost

[This step applies only if you have created a schedule for the backup policy at the **Schedule** step of the wizard]

At the **Cost Estimation** step of the wizard, review the estimated monthly cost of AWS services and resources that will be consumed to protect the file systems added to the backup policy. The total estimated cost includes the following:

- The cost of creating backups of the FSx file systems.
For each file system included in the backup policy, Veeam Backup for AWS takes into account the number of restore points to be kept in the backup chain and the configured scheduling settings.
- The cost of creating backup copies and maintaining them in the target AWS Region.
For each file system included in the backup policy, Veeam Backup for AWS takes into account the number of restore points to be kept in the backup chain and the configured scheduling settings.
- The cost of sending API requests to Veeam Backup for AWS during data protection operations.

NOTE

To calculate the estimated cost, Veeam Backup for AWS uses capabilities of the [AWS Pricing Calculator](#) that estimates the cost of services in USD only. This calculator is intended for informational and estimation purposes only.

The estimated cost may occur to be significantly higher due to the backup frequency, cross-region data transfer and AWS backup charges. To reduce the cost, you can try the following workarounds:

- To reduce high AWS backup charges, adjust the backup retention settings to keep less restore points in the backup chain.
- To optimize the cost of storing backups, configure the scheduling settings to run the backup policy less frequently.

TIP

You can save the cost estimation as a .CSV or .XML file. To do that, click **Export to** and select the necessary format.

The screenshot shows the 'Add FSx Policy' configuration page in Veeam Backup for AWS. The top navigation bar includes the Veeam logo, 'Server time: Jul 12, 2024 9:41 AM', user 'administrator', and 'Configuration' settings. The page title is 'Add FSx Policy' with a back arrow and a cost indicator 'Cost: \$25.61'. A left sidebar lists navigation options: Info, Sources, Targets, Schedule, Tags, General Settings, Cost Estimation (selected), and Summary. The main content area is titled 'Review cost estimation' and includes explanatory text and three cost breakdown cards: Backups (\$13.84), Backup copies (\$9.11), and Traffic (\$2.66). A dark green box highlights the 'Estimated monthly cost: \$25.61'. Below this is a search bar for 'File System' and an 'Export to...' button. A table summarizes the costs for the 'jf-test' file system.

File System ↑	Backup	Backup Copy	Traffic	Total
jf-test	\$13.84	\$9.11	\$2.66	\$25.61

Navigation buttons at the bottom: Previous, Next, Cancel.

Related Resources

[How AWS Pricing Works](#)

Step 9. Finish Working with Wizard

At the **Summary** step of the wizard, it is recommended that you run the backup policy configuration check before you click **Finish**.

The configuration check will verify whether specified IAM roles have all the required permissions. To run the check, click **Test Configuration**. Veeam Backup for AWS will display the **Test policy configuration** window where you can track the progress and view the results of the check. If the IAM role permissions are insufficient or policy settings are not configured properly, the check will complete with errors, and the list of permissions that must be granted to the IAM role and policy configuration issues will be displayed in the **Test policy configuration** window.

You can grant the missing permissions to the IAM role [using the AWS Management Console](#) or instruct Veeam Backup for AWS to do it.

To let Veeam Backup for AWS grant the missing permissions:

1. In the **Test policy configuration** window, click the **Grant** link.
2. In the **Grant Permissions** window, provide one-time access keys of an IAM user that is authorized to update permissions of IAM roles, and then click **Apply**.

The IAM user must have the following permissions:

```
"iam:AttachRolePolicy",
"iam:CreatePolicy",
"iam:CreatePolicyVersion",
"iam:CreateRole",
"iam:GetAccountSummary",
"iam:GetPolicy",
"iam:GetPolicyVersion",
"iam:GetRole",
"iam>ListAttachedRolePolicies",
"iam>ListPolicyVersions",
"iam:SimulatePrincipalPolicy",
"iam:UpdateAssumeRolePolicy"
```

NOTE

Veeam Backup for AWS does not store one-time access keys in the configuration database.

3. After the required permissions are granted, close the **Test policy configuration** window, and then click **Finish** to close the **Add Policy** wizard.

Veeam Backup for AWS will save the configured backup policy.

The screenshot shows the 'Add FSx Policy' wizard in Veeam Backup for AWS. The interface is in a dark theme. At the top, the header shows 'Veeam Backup for AWS', 'Server time: Jul 12, 2024 9:42 AM', and the user 'administrator Portal Administrator'. The main title is 'Add FSx Policy' with a cost of '\$25.61'. The left sidebar has a 'Summary' tab selected. The main content area is titled 'Review configured settings' and includes a warning message: 'In order to successfully run this policy, we advise to test the configuration.' Below this are buttons for 'Test Configuration' and 'Copy to Clipboard'. The settings are organized into sections: 'General' (Name: fsx backup policy 01, Description: backup policy for dept-01, Regions: Europe (Milan), Europe (Stockholm), Account: Default: Backup Restore), 'Backup settings' (Copy tags from source file systems: Yes, Add custom tags: Yes, Custom tags: user :donna_ortiz, owner:dept01), 'Backup schedule' (Daily retention: Create 3 restore points and keep for 14 Days, Weekly retention: Create 2 restore points and keep for 7 Days, Monthly retention: Create 5 restore points and keep for 6 Months, Yearly retention: Create restore point on First Monday of January at 06:00 AM, Keep backups for 2 years), and 'Backup copy settings' (Enabled: Yes, Region mapping: Source region: Europe (Milan), Europe (Stockholm); Target region: —, Europe (Paris)). At the bottom, there are 'Previous', 'Finish', and 'Cancel' buttons.

Creating FSx Backups Manually

Veeam Backup for AWS allows you to manually create backups of Amazon FSx file systems. You can instruct Veeam Backup for AWS to store the created backups in the same AWS Regions where the processed file systems reside, or in a different AWS Region.

NOTE

Veeam Backup for AWS does not include FSx backups created manually in the FSx backup chain and does not apply the configured retention policy settings to these backups. This means that the backups are kept in your AWS environment unless you remove them manually, as described in section [Managing Backed-Up FSx Data](#).

To manually create a backup of an FSx file system, do the following:

1. Navigate to **Resources > File Systems > FSx**.
2. Select the necessary file system and click **Take Backup Now**.

For an FSx file system to be displayed in the list of available file systems, an AWS Region where the file system resides must be added to any of [configured FSx backup policies](#), and the IAM role specified in the backup policy settings must have permissions to access the file system. For more information on required permissions, see [FSx Backup IAM Role Permissions](#).

IMPORTANT

Veeam Backup for AWS does not support taking manual backup of multiple file systems that belong to different AWS accounts or reside in different AWS Regions.

3. Complete the **Take Manual Backup** wizard:

- a. At the **Account** step of the wizard, specify an IAM role whose permissions Veeam Backup for AWS will use to create the backup.

For an IAM role to be displayed in the list, it must be added to Veeam Backup for AWS as described in section [Adding IAM Roles](#).

- b. In the **Backup vault** section of the **Settings** step of the wizard, click **Edit Location Settings**.

In the **Choose region and backup vault** window, specify the following settings:

- i. From the **Target region** drop-down list, choose an AWS Region where manual backups will be stored.

Since cross-region copying of FSx backups is not supported for [opt-in Regions](#), the list of available regions will depend on the original location of the selected file system. If the original location is an opt-in Region, the **Target region** list will contain only this AWS Region; if the original location is a default AWS Region (that is, one of the AWS Regions activated for the AWS account by default), the **Target region** list will contain all default AWS Regions.

- ii. In the **Backup vault** section, select a backup vault that will be used to store file system backups.

For a backup vault to be displayed in the list of available vaults, it must be created in the AWS Backup console as described in [AWS Documentation](#). If no backup vaults were created in the selected AWS Region, Veeam Backup for AWS will display only the default backup vault created for the AWS Region automatically.

- iii. To save changes made to the location settings, click **Apply**.

IMPORTANT

Consider the following:

- Make sure that the type of file system added to the backup scope is supported by the AWS Backup service in the source AWS Region. Otherwise, the backup operation will fail to complete successfully. For the list of supported AWS Regions, see [AWS Documentation](#).
- Make sure policies assigned to the selected backup vault allow Veeam Backup for AWS to access vault resources and to perform backup, backup copy and restore operations. For more information on vault access policies, see [AWS Documentation](#).
- For Veeam Backup for AWS to be able to back up FSx file systems, you must enable the Opt-in service for the FSx resource type in the AWS Backup settings. Otherwise, Veeam Backup for AWS will automatically enable the service for each AWS Region specified in the backup settings in your AWS account while performing backup operations.

- c. At the **Tags** section of the **Settings** step of the wizard, if you want to assign tags to the created backup, click **Edit Tag Settings**.

In the **Tag configuration** window, specify tag settings:

- i. To assign already existing AWS tags from the processed file system, select the **Copy tags from source file system** check box.

If you choose to copy tags from source file system, Veeam Backup for AWS will first create a backup of the FSx file system and assign to the created backup AWS tags with Veeam metadata, then Veeam Backup for AWS will copy tags from the processed file system and, finally, assign the copied AWS tags to the backup.

- ii. To assign your own custom AWS tags, set the **Add custom tags to created backup** toggle to *On* and specify the tags explicitly. To do that, use the **Key** and **Value** fields to specify a key and a value for the new custom AWS tag, and then click **Add**. Note that you cannot add more than 5 custom AWS tags.

If you choose to add custom tags to created backups, Veeam Backup for AWS will assign the specified tags right after it creates a backup.

- iii. To save changes made to the tag settings, click **Apply**.

- d. At the **Summary** step of the wizard, review summary information and click **Finish**.

The screenshot shows the 'Take Manual Backup' wizard in Veeam Backup for AWS, specifically the 'Summary' step. The interface includes a top navigation bar with the Veeam logo, server time (Jun 17, 2024 12:33 PM), user information (administrator, Portal Administrator), and a Configuration icon. The main content area is titled 'Review configured settings' and contains several sections: 'IAM role' (IAM role name: Default Backup Restore), 'Location' (Backup vault: bd-vault-frankfurt, Region: Europe (Frankfurt)), 'Tags' (Copy tags from source file system: Yes, Add custom tags: owner:dept01), and 'Resources' (Added resources: bd-frankfurt-openzfs-regress-single-az1-custom-key). An information message states: 'After you click Finish, the backup will be created. To view the session progress, switch to the Session Logs page.' There is a checked checkbox for 'Go to Sessions'. At the bottom, there are three buttons: 'Previous', 'Finish', and 'Cancel'.

Performing VPC Configuration Backup

To protect the Amazon VPC configuration and settings, Veeam Backup for AWS comes with a preconfigured VPC Configuration Backup policy. With this policy, you can protect VPC configurations of AWS Regions in your AWS accounts.

The VPC Configuration Backup policy is disabled by default. To start protecting your Amazon VPC configuration, [edit backup policy settings](#) and [enable the policy](#).

IMPORTANT

Veeam Backup for AWS does not support backup of the following VPC configuration components: VPC Traffic Mirroring, AWS Network Firewall, Route 53 Resolver DNS Firewall, AWS Verified Access, VPC Flow Logs, carrier gateways, customer IP pools, transit gateway policy tables, and core networks in route tables.

Editing VPC Configuration Backup Policy

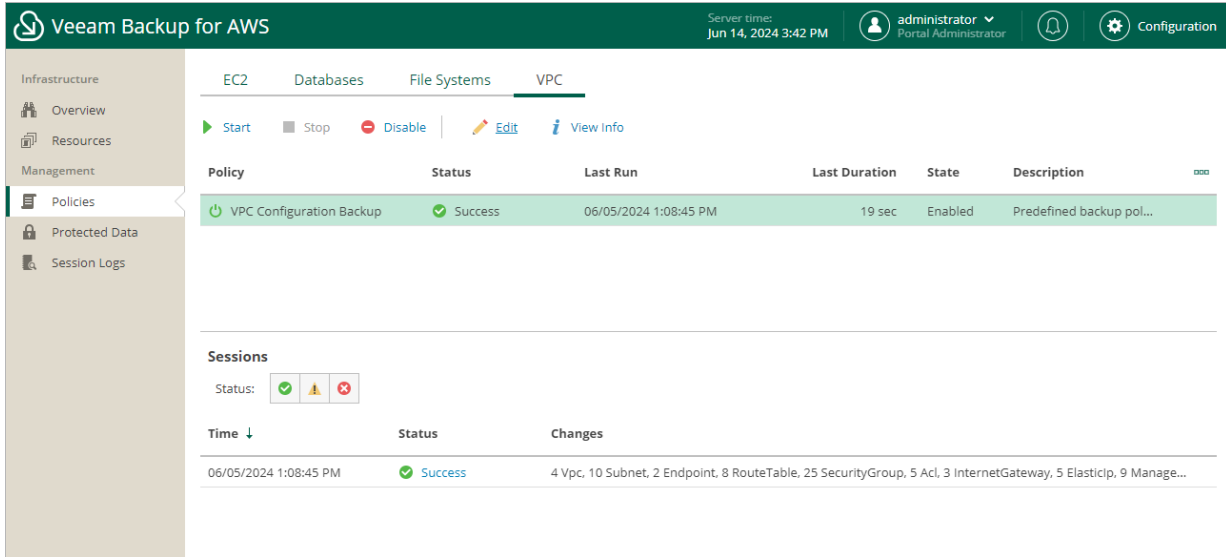
To configure the VPC Configuration Backup policy settings, do the following:

1. [Launch the VPC Configuration Backup wizard](#).
2. [Select AWS Regions to protect](#).
3. [Specify a backup repository to store an additional backup copy](#).
4. [Configure retentions settings for VPC configuration backups](#).
5. [Specify automatic retry settings and notification settings for the backup policy](#).
6. [Finish working with the wizard](#).

Step 1. Launch VPC Configuration Backup Wizard

To launch the **VPC Configuration Backup** wizard, do the following:

1. Navigate to **Policies > VPC**.
2. Click **Edit**.



Step 2. Select AWS Regions

At the **Regions** step of the wizard, select AWS Regions whose VPC configuration you want to back up.

Veeam Backup for AWS allows you to automatically collect and back up VPC configuration data for all AWS Regions selected for EC2, RDS, DynamoDB, Redshift, EFS and FSx backup policies. To do that, [enable automatic protection](#) for AWS Regions. To retrieve VPC configurations of all automatically protected AWS Regions, Veeam Backup for AWS will use permissions of IAM roles specified in the settings of backup policies that protect instances residing in these AWS Regions.

You can also configure the VPC Configuration Backup policy to protect configuration data for AWS Regions that are not specified in the settings of any backup policy, or choose another IAM role whose permissions Veeam Backup for AWS will use to collect the VPC configuration data of the automatically protected AWS Regions. To do that, [manually add AWS Regions](#) to the VPC Backup policy and configure backup settings for them.

Enabling Automatic Protection

To instruct Veeam Backup for AWS to protect VPC configuration of all AWS Regions specified in EC2, RDS, DynamoDB, Redshift, EFS and FSx backup policy settings, in the **Automatically protected regions** section, set the **Automatically collect VPC settings** toggle to *On*.

To retrieve VPC configurations of all automatically protected AWS Regions, Veeam Backup for AWS will use permissions of IAM roles specified in the settings of backup policies that protect instances residing in these AWS Regions. It is recommended that you check whether IAM roles whose permissions EC2, RDS, DynamoDB, Redshift, EFS and FSx backup policies use to perform data protection operations have all the permissions required to perform Amazon VPC configuration backup. If some permissions of the IAM role are missing, the backup policy will fail.

To run the IAM role permission check:

1. In the **Automatically Protected Regions** section, click the **Discovered regions** link.
2. In the **Discovered regions** window, select the IAM role whose permissions you want to check.
3. Click **Check Permissions**.

Veeam Backup for AWS will display the **AWS Permission Check** window where you can view the progress and results of the performed check. If some permissions of the IAM role are missing, the check will complete with errors. You can view the list of permissions that must be granted to IAM roles in the **Missing Permissions** column. For more information on required permissions, see [VPC Configuration Backup IAM Role Permissions](#).

You can grant the missing permissions to IAM roles in the AWS Management Console or instruct Veeam Backup for AWS to do it. To learn how to grant permissions to IAM roles using the AWS Management Console, see [AWS Documentation](#). To let Veeam Backup for AWS grant the missing permissions:

- a. In the **AWS Permission Check** window, click **Grant**.
- b. In the **Grant Permissions Window**, provide one-time access keys of an IAM user that is authorized to update permissions of the IAM role, and then click **Apply**.

The IAM user whose access keys are used to update the IAM role must have the following permissions:

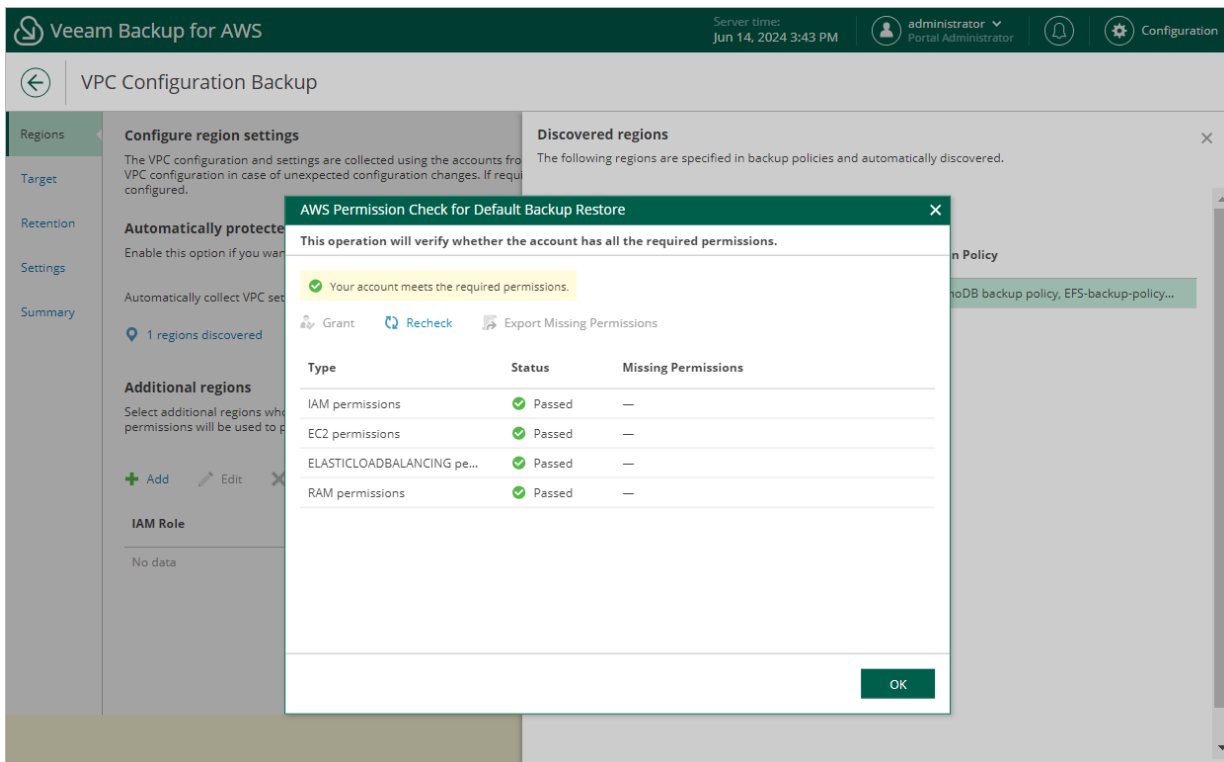
```

"iam:AttachRolePolicy",
"iam:CreatePolicy",
"iam:CreatePolicyVersion",
"iam:CreateRole",
"iam:GetAccountSummary",
"iam:GetPolicy",
"iam:GetPolicyVersion",
"iam:GetRole",
"iam:ListAttachedRolePolicies",
"iam:ListPolicyVersions",
"iam:SimulatePrincipalPolicy",
"iam:UpdateAssumeRolePolicy"

```

NOTE

Veeam Backup for AWS does not store one-time access keys in the configuration database.



Adding AWS Regions Manually

To add an AWS Region to the VPC Backup policy, or to choose another IAM role for collecting VPC configuration data, do the following:

1. In the **Additional regions** section, click **Add**.

2. In the **Configure account settings** window, from the **IAM role** drop-down list, select an IAM role whose permissions Veeam Backup for AWS will use to perform Amazon VPC configuration backup. In the **Account** field, the ID of the AWS account in which the IAM role was created will be displayed. The specified IAM role must be assigned the permissions listed in section [VPC Configuration Backup IAM Role Permissions](#).

For an IAM role to be displayed in the **IAM role** list, it must be added to Veeam Backup for AWS with the *Amazon VPC Backup operation* selected as described in section [Adding IAM Roles](#). If you have not added the necessary IAM role to Veeam Backup for AWS beforehand, you can do it without closing the **VPC Configuration Backup** wizard. To do that, click **Add** and complete the **Add IAM Role** wizard.

3. In the **Regions** section, select the necessary AWS Regions from the **Available Regions** list on the left, and then click **Add**.
4. To save changes made to the backup policy settings, click **Apply**.
5. To check whether IAM role specified for the selected AWS Regions has all the permissions required to perform Amazon VPC configuration backup, in the **Additional regions** section, click **Check Permissions**.

Veeam Backup for AWS will display the **AWS Permission Check window** where you can view the progress and results of the performed check. If some permissions of the IAM role are missing, the check will complete with errors. You can view the list of permissions that must be granted to IAM roles in the **Missing Permissions** column. For more information on required permissions, see [VPC Configuration Backup IAM Role Permissions](#).

You can grant the missing permissions to IAM roles in the AWS Management Console or instruct Veeam Backup for AWS to do it. To learn how to grant permissions to IAM roles using the AWS Management Console, see [AWS Documentation](#). To let Veeam Backup for AWS grant the missing permissions:

- a. In the **AWS Permission Check** window, click **Grant**.
- b. In the **Grant Permissions Window**, provide one-time access keys of an IAM user that is authorized to update permissions of the IAM role, and then click **Apply**.

The IAM user whose access keys are used to update the IAM role must have the following permissions:

```
"iam:AttachRolePolicy",  
"iam:CreatePolicy",  
"iam:CreatePolicyVersion",  
"iam:CreateRole",  
"iam:GetAccountSummary",  
"iam:GetPolicy",  
"iam:GetPolicyVersion",  
"iam:GetRole",  
"iam:ListAttachedRolePolicies",  
"iam:ListPolicyVersions",  
"iam:SimulatePrincipalPolicy",  
"iam:UpdateAssumeRolePolicy"
```

NOTE

Veeam Backup for AWS does not store one-time access keys in the configuration database.

You can add, edit or remove additional AWS Regions from the VPC Backup policy.

Veeam Backup for AWS Server time: Jun 14, 2024 3:45 PM administrator Portal Administrator Configuration

VPC Configuration Backup

Regions

Configure region settings

The VPC configuration and settings are collected using the accounts from VPC configuration in case of unexpected configuration changes. If required, you can configure the accounts.

Automatically protected regions

Enable this option if you want to automatically collect VPC settings for a region.

Automatically collect VPC settings: On

1 regions discovered

Additional regions

Select additional regions whose VPC configuration you want to protect. Permissions will be used to protect the selected region.

+ Add Edit Remove Check Permissions

IAM Role	Account	Region
No data		

Configure account settings

Choose an IAM role to use and specify regions that will be protected. The IAM role requires sufficient permissions to read the VPC configurations of the specified regions.

IAM role: Default Backup Restore (Default Backup Restore) + Add

Account: 611610175276

Regions

Specify regions from which the VPC configuration will be collected.

Available Regions (16)

- Asia Pacific (Sydney)
- Asia Pacific (Tokyo)
- Canada (Central)
- Europe (Frankfurt)
- Europe (Ireland)
- Europe (London)
- Europe (Paris)
- South America (Sao Paulo)
- US East (N. Virginia)
- US East (Ohio)

Selected Regions (2)

- Europe (Milan)
- Europe (Stockholm)

Apply Cancel

Step 3. Enable Additional Backup Copy

By default, Veeam Backup for AWS stores VPC configuration backups in the Veeam Backup for AWS database. You can instruct Veeam Backup for AWS to save additional VPC configuration backup copies to a backup repository. To do that:

1. At the **Target** step of the wizard, set the **Enable additional copy** toggle to *On*.
2. In the **Repository** window, select a backup repository that will be used to store the additional configuration backup copies.

For a backup repository to be displayed in the **Repository** list, it must be added to Veeam Backup for AWS as described in section [Adding Backup Repositories Using Web UI](#). The list shows only backup repositories of the *S3 Standard* storage class.

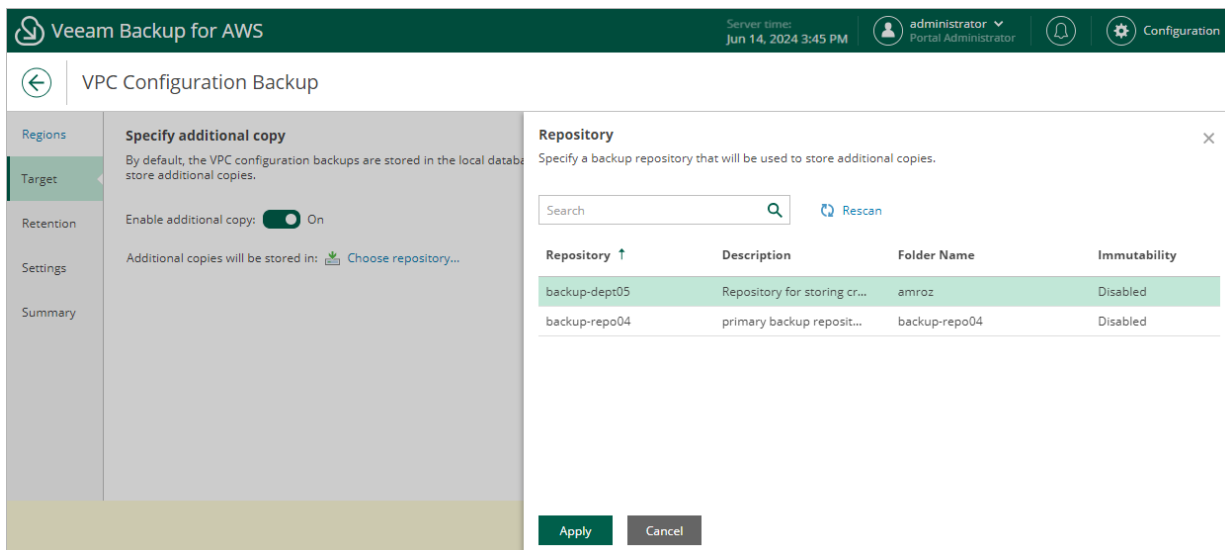
3. To save changes made to the backup policy settings, click **Apply**.

NOTE

When choosing a backup repository, consider the following:

- If you want to encrypt the backed-up VPC configuration data, select a repository with encryption enabled.
- If you want to make the backed-up VPC configuration data immutable for the period specified in [retention settings](#) of the backup policy, select a repository with immutability enabled. Note that Veeam Backup for AWS does not apply generations to VPC backups.

For more information on encryption and immutability, see [Adding Backup Repositories Using Web UI](#).



The screenshot shows the Veeam Backup for AWS web interface. The main window is titled 'VPC Configuration Backup' and is in the 'Target' step. The 'Enable additional copy' toggle is turned on. A 'Repository' dialog box is open, showing a search bar and a table of backup repositories.

Repository ↑	Description	Folder Name	Immutability
backup-dept05	Repository for storing cr...	amroz	Disabled
backup-repo04	primary backup reposit...	backup-repo04	Disabled

Step 4. Configure Retention Settings

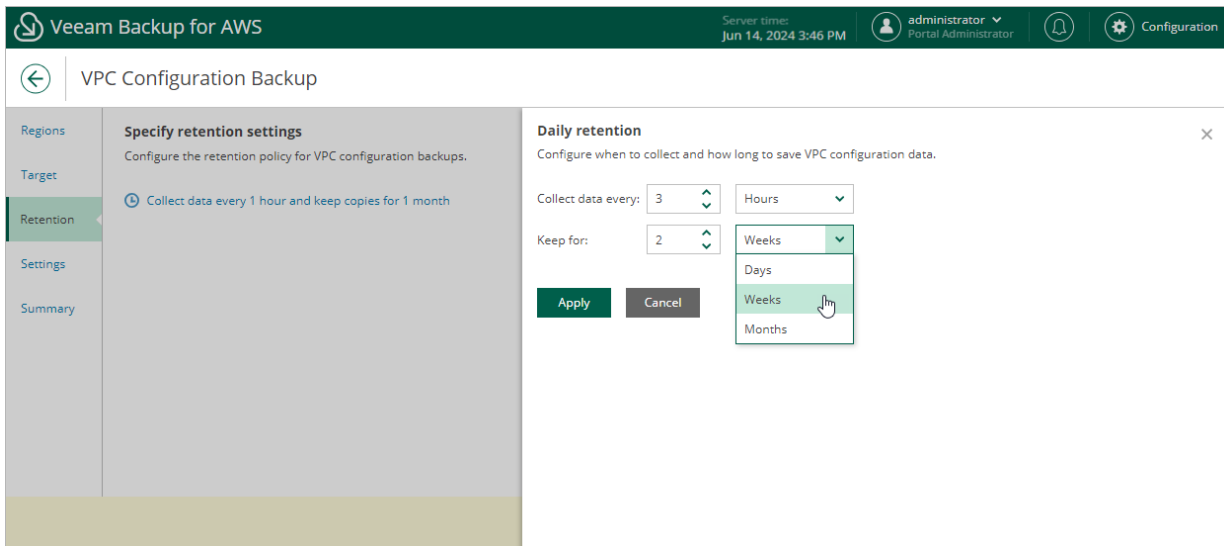
At the **Retention** step of the wizard, specify retention settings for VPC configuration backups:

1. Click the **Collect data** link.
2. In the **Daily retention** window, specify how often the data will be backed up and for how long the backups will be stored.

If a restore point is older than the specified time limit, Veeam Backup for AWS removes the restore point from the backup chain. For more information, see [VPC Configuration Backup Retention](#).

NOTE

Veeam Backup for AWS applies the retention settings configured for the VPC Configuration Backup policy both to VPC configuration backups stored in the Veeam Backup for AWS database and to VPC configuration backups stored in the backup repository selected for the policy. For VPC configuration backups stored in backup repositories that are not specified in the VPC Configuration Backup policy settings, Veeam Backup for AWS applies retention settings saved in the backup metadata.



Step 5. Specify Email Notification Settings

At the **Settings** step of the wizard, you can specify email notification settings for the VPC Backup policy.

NOTE

If you want to receive daily reports and email notifications on the VPC Configuration Backup policy results, you must configure [global notification settings](#) first.

To instruct Veeam Backup for AWS to send email notifications for the backup policy, do the following:

1. In the **Notifications** section, set the **Receive daily report** toggle to *On*.
If you set the toggle to *Off*, Veeam Backup for AWS will send notifications according to the configured global notification settings.
2. In the **Email** field, specify an email address of a recipient.
Use a semicolon to separate multiple recipient addresses. Do not use spaces after semicolons between the specified email addresses.
3. Use the **Notify on** list to choose whether you want Veeam Backup for AWS to send email notifications in case the backup policy completes successfully, completes with warnings or completes with errors.

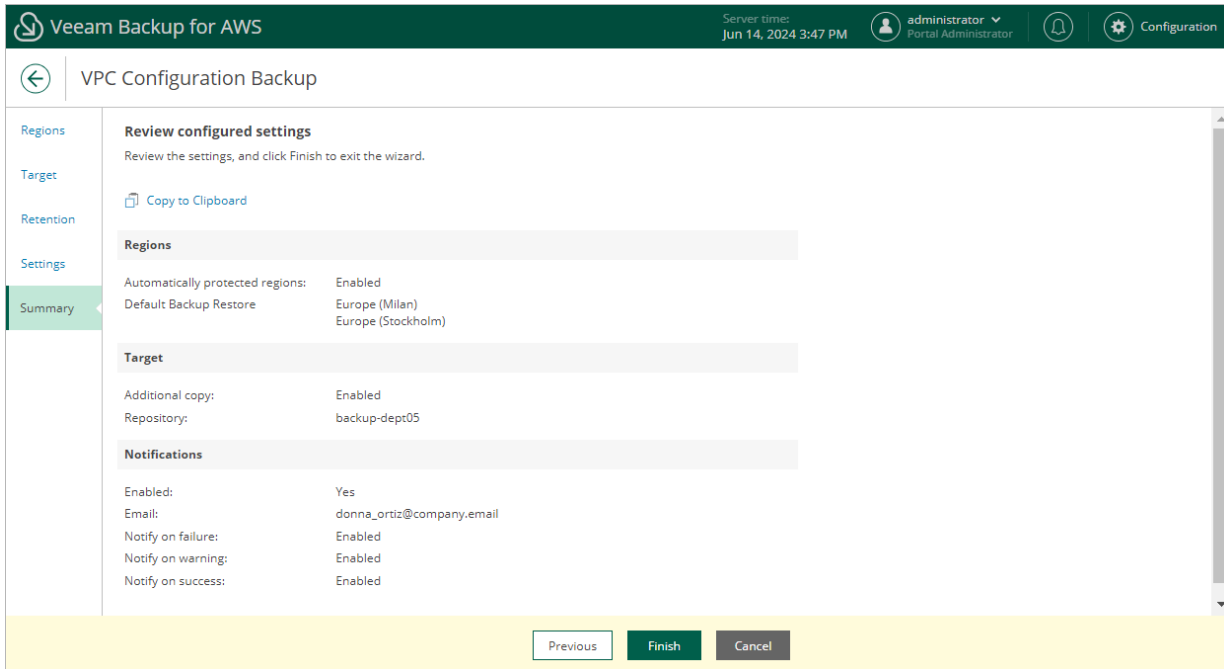
NOTE

If you specify the same email recipient in both backup policy notification and [global notification settings](#), Veeam Backup for AWS will override the configured global notification settings and will send each notification to this recipient only once to avoid notification duplicates.

The screenshot shows the 'Configure notification settings' step in the Veeam Backup for AWS wizard. The interface includes a top navigation bar with the Veeam logo, 'Veeam Backup for AWS', server time (Jun 14, 2024 3:46 PM), user information (administrator, Portal Administrator), and a 'Configuration' icon. The main content area is titled 'VPC Configuration Backup' and has a left sidebar with navigation options: 'Regions', 'Target', 'Retention', 'Settings' (selected), and 'Summary'. The 'Settings' section is titled 'Configure notification settings' and contains the following elements: a sub-header 'Configure daily email notifications.', a 'Notifications' section with a 'Receive daily report' toggle set to 'On', an 'Email' input field containing 'donna_ortiz@company.email', and a 'Notify on' section with three checked checkboxes: 'Failure', 'Warning', and 'Success'. At the bottom of the wizard, there are three buttons: 'Previous', 'Next' (highlighted in green), and 'Cancel'.

Step 6. Finish Working with Wizard

At the **Summary** step of the wizard, review configuration information and click **Finish**.




Enabling and Disabling VPC Configuration Backup Policy

By default, Veeam Backup for AWS comes with the disabled VPC Configuration Backup Policy. You can [manually start](#) or enable the disabled backup policy at any time you need.

To enable or disable the VPC Configuration Backup policy, do the following:

1. Navigate to **Policies > VPC**.
2. Click **Enable** or **Disable**.



Starting and Stopping VPC Configuration Backup Policy

You can start the VPC Configuration Backup policy manually, for example, if you want to create an additional restore point in the backup chain and do not want to modify the configured backup policy schedule. You can also stop a backup policy if the backup process is about to take long, and you do not want the policy to have an impact on the production environment during business hours.

To start or stop a backup policy, do the following:

1. Navigate to **Policies > VPC**.
2. Click **Start** or **Stop**.

The screenshot displays the Veeam Backup for AWS console interface. The top navigation bar includes the Veeam logo, the text 'Veeam Backup for AWS', the server time 'Jun 14, 2024 3:49 PM', and the user 'administrator Portal Administrator'. The left sidebar shows a navigation menu with categories: Infrastructure (Overview, Resources), Management (Policies, Protected Data, Session Logs), and Configuration. The main content area is titled 'VPC' and contains a table of backup policies. The 'VPC Configuration Backup' policy is highlighted in green, showing a 'Success' status and a last run time of '06/05/2024 1:08:45 PM'. Below the policy table, there is a 'Sessions' section with a status indicator (Success) and a table of session details. The session table shows a successful backup on '06/05/2024 1:08:45 PM' with a status of 'Success' and a list of changes including '4 Vpc, 10 Subnet, 2 Endpoint, 8 RouteTable, 25 SecurityGroup, 5 Acl, 3 InternetGateway, 5 Elasticip, 9 Manage...'.

Policy	Status	Last Run	Last Duration	State	Description
VPC Configuration Backup	Success	06/05/2024 1:08:45 PM	19 sec	Enabled	Predefined backup pol...

Time ↓	Status	Changes
06/05/2024 1:08:45 PM	Success	4 Vpc, 10 Subnet, 2 Endpoint, 8 RouteTable, 25 SecurityGroup, 5 Acl, 3 InternetGateway, 5 Elasticip, 9 Manage...

Managing Backup Policies

You can manage and edit created EC2, RDS, DynamoDB, Redshift, EFS and FSx backup policies, and view each backup policy details in Veeam Backup for AWS. You can also remove backup policies that you do not use anymore, export existing or import new backup policies.

Starting and Stopping Policies

You can start a backup policy manually, for example, if you want to create an additional restore point in the snapshot or backup chain and do not want to modify the configured backup policy schedule. You can also stop a backup policy if processing of an instance is about to take too long, and you do not want the policy to have an impact on the production environment during business hours.

To start or stop a backup policy, do the following:

1. Navigate to **Policies**.
2. Switch to the necessary tab and select the backup policy.
3. Click **Start** or **Stop**.

NOTE

When you run a backup policy manually, consider the following:

- The created restore points will be retained for the time period specified in the most frequent backup policy schedule.
- [Applies only to EC2 backup policies] If the backup policy stores backups in a backup repository with immutability settings enabled, the created restore points will be immutable for the time period determined based on the retention settings specified in the most frequent backup policy schedule. For more information, see [Immutability](#).

The screenshot shows the Veeam Backup for AWS console. The top header displays the Veeam logo, the text "Veeam Backup for AWS", the server time "Jun 14, 2024 3:36 PM", and the user "administrator Portal Administrator". The left sidebar contains navigation options: Infrastructure (Overview, Resources), Management (Policies, Protected Data, Session Logs). The main content area has tabs for EC2, Databases, File Systems, and VPC. The EC2 tab is selected, showing a search bar, a filter dropdown, and a toolbar with actions like Start, Stop, Enable, Add, Edit, Priority, View Info, Remove, Advanced, and Export to... Below this is a table of policies:

Priority	Policy	Snapshots	Backups	Archive	Last Run
1	policy-ec2-487569979969	Success	Success	Success	06/12/2024 1:58:26 PM
2	policy-ec2-407355668422	Never Executed	Not Configured	Not Configured	—

Below the table are sections for Instances and Sessions, each with a search bar and status filters.

Disabling and Enabling Policies

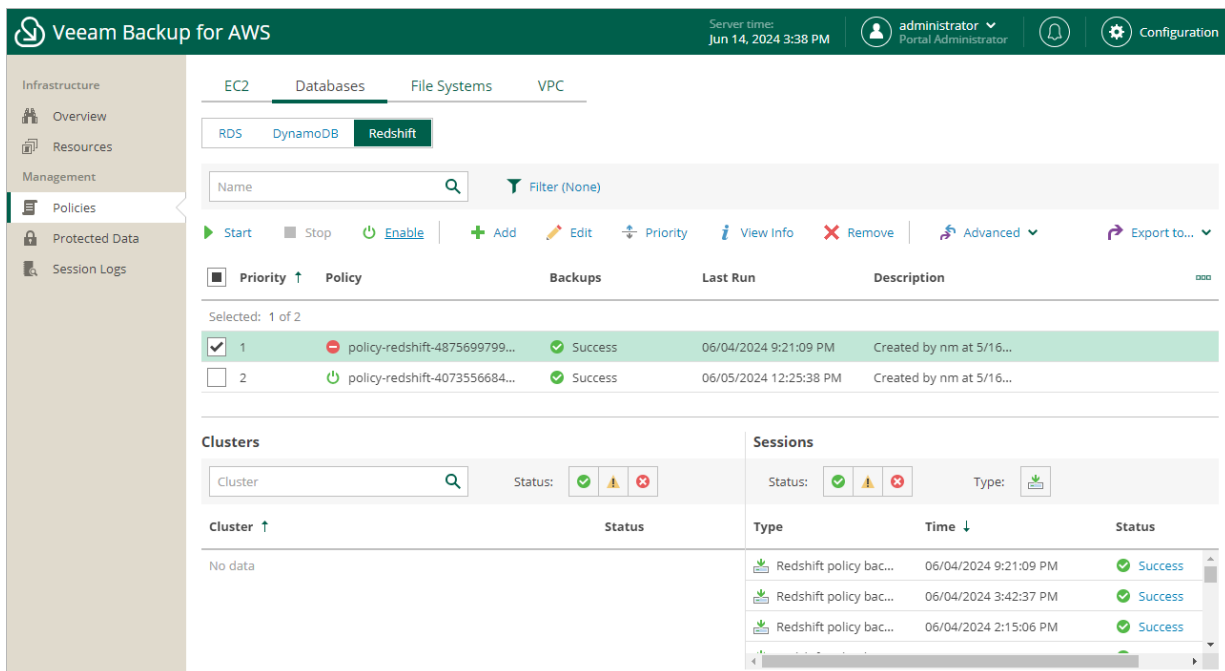
By default, Veeam Backup for AWS runs all created backup policies according to the specified schedules. However, you can temporarily disable a backup policy so that Veeam Backup for AWS does not run the backup policy automatically. You will still be able to [manually start](#) or enable the disabled backup policy at any time you need.

To enable or disable a backup policy, do the following:

1. Navigate to **Policies**.
2. Switch to the necessary tab and select the backup policy.
3. Click **Disable** or **Enable**.

NOTE

Disabling a backup policy does not affect the retention settings configured for the cloud-native snapshots, image-level and archived backups created by the policy. Veeam Backup for AWS will continue running retention sessions for the disabled backup policy and removing restore points according to the configured settings.



The screenshot shows the Veeam Backup for AWS console interface. The top navigation bar includes the Veeam logo, the product name 'Veeam Backup for AWS', the server time 'Jun 14, 2024 3:38 PM', and the user 'administrator Portal Administrator'. The left sidebar shows the navigation menu with 'Policies' selected. The main content area is titled 'Redshift' and contains a search bar, a filter dropdown, and a toolbar with actions like Start, Stop, Enable, Add, Edit, Priority, View Info, Remove, Advanced, and Export to... Below this is a table of backup policies. The table has columns for Priority, Policy, Backups, Last Run, and Description. Two policies are listed, both with a status of 'Success'. Below the table, there are sections for Clusters and Sessions. The Clusters section is empty, and the Sessions section shows three successful backup sessions for the Redshift policy.

Priority	Policy	Backups	Last Run	Description
1	policy-redshift-4875699799...	Success	06/04/2024 9:21:09 PM	Created by nm at 5/16...
2	policy-redshift-4073556684...	Success	06/05/2024 12:25:38 PM	Created by nm at 5/16...

Cluster	Status
No data	

Type	Time	Status
Redshift policy bac...	06/04/2024 9:21:09 PM	Success
Redshift policy bac...	06/04/2024 3:42:37 PM	Success
Redshift policy bac...	06/04/2024 2:15:06 PM	Success

Setting Policy Priority

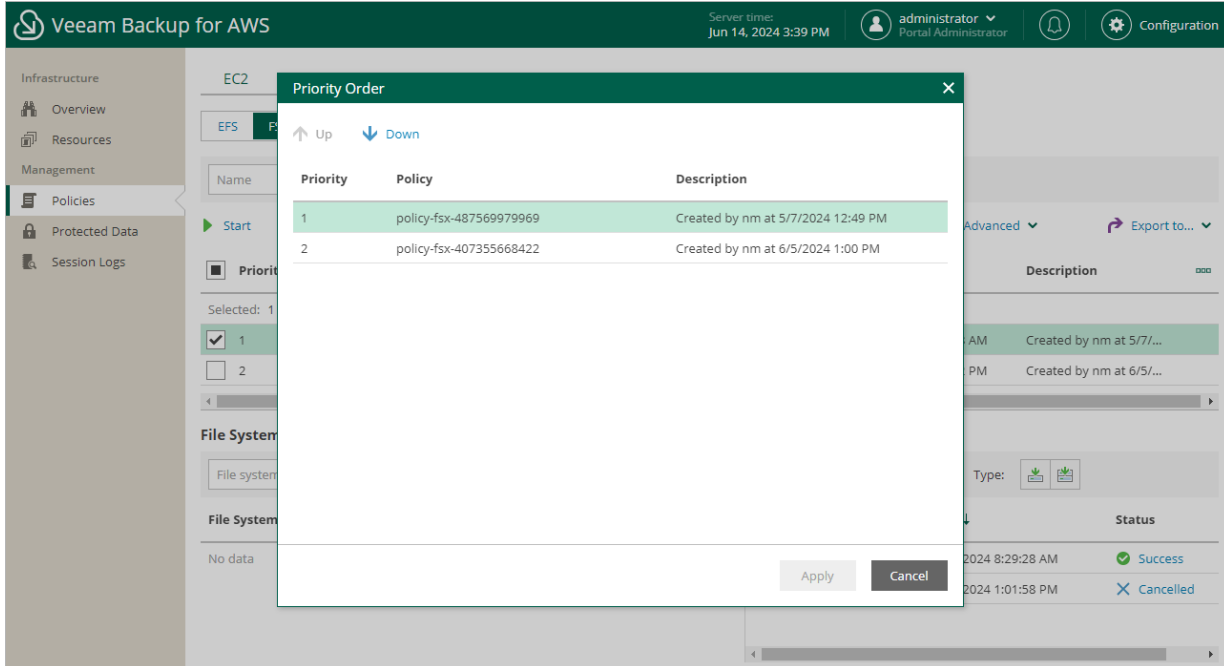
By default, Veeam Backup for AWS runs backup policies in the order you create them. However, you can set the backup policy priority manually:

1. Navigate to **Policies**.
2. Switch to the necessary tab and click **Policy Priority**.
3. In the **Priority Order** window, use the **Up** and **Down** arrows to set priority for backup policies, and click **Apply** to save the settings.

The first backup policy in the list will have the highest priority.

NOTE

If a resource is included into multiple backup policies, it will be processed only by the backup policy that has the highest priority.



Editing Policy Settings

For each backup policy, you can modify settings configured while creating the policy:

1. Navigate to **Policies**.
2. Switch to the necessary tab and select the backup policy whose settings you want to edit.
3. Click **Edit**. The **Edit Policy** wizard will open.
4. Edit backup policy settings as described in sections [Creating EC2 Backup Policies](#), [Creating RDS Backup Policies](#), [Creating DynamoDB Backup Policies](#), [Creating Redshift Backup Policies](#), [Creating EFS Backup Policies](#) or [Creating FSx Backup Policies](#).

TIP

To protect additional resources by a configured backup policy, you can either edit the resource list in the backup policy settings, or add resources to the backup policy on the **Resources** tab. To learn how to add resources on the **Resources** tab, see [Adding Resources to Policy](#).

The screenshot displays the Veeam Backup for AWS console interface. The top navigation bar includes the Veeam logo, server time (Jun 14, 2024 3:40 PM), user profile (administrator), and configuration settings. The left sidebar shows navigation options: Infrastructure, Overview, Resources, Management, Policies, Protected Data, and Session Logs. The main content area is divided into tabs for EC2, Databases, File Systems, and VPC. Under the File Systems tab, there are sub-tabs for EFS and FSx. A search bar and a filter dropdown are visible. Below, a table lists backup policies with columns for Priority, Policy, Backups, Backup Copy, Last Run, and Description. Two policies are shown, both with a 'Success' status. Below the table, there are sections for 'File Systems' and 'Sessions'. The 'Sessions' section shows a table with columns for Type, Time, and Status, listing two backup sessions: one successful and one cancelled.

Exporting and Importing Policies

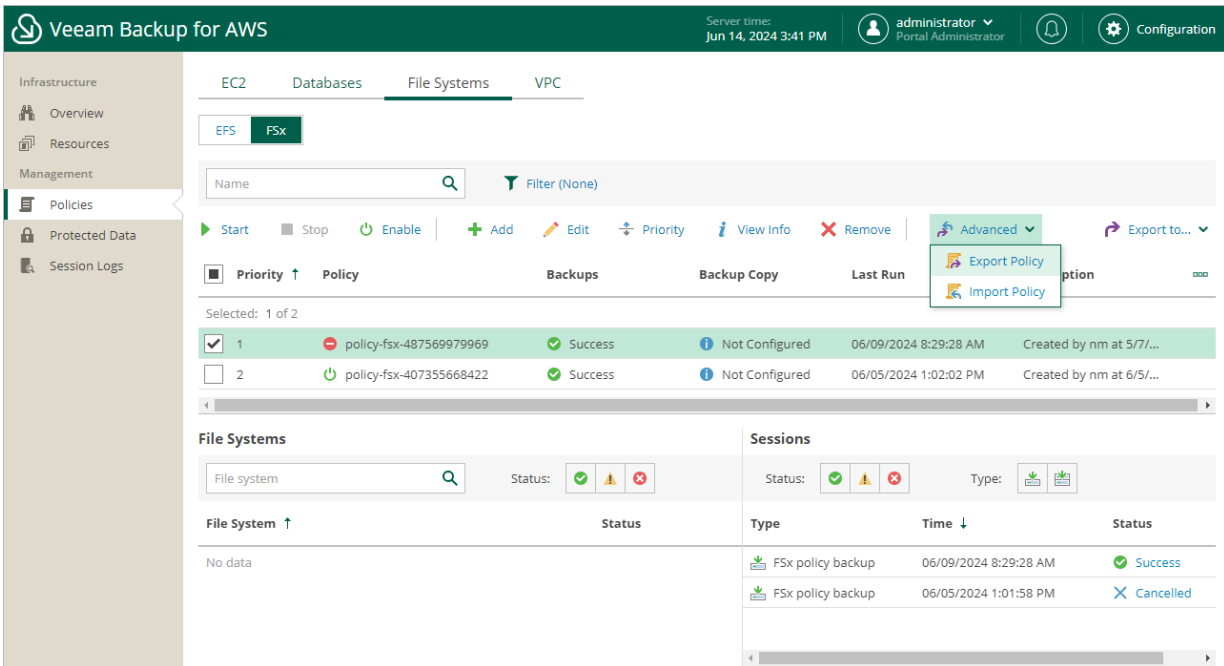
Veeam Backup for AWS allows you to use settings of an existing backup policy as a template for creating other backup policies. You can export a backup policy to a .JSON file, modify the necessary settings in the file, and then import the policy to the same or a different backup appliance.

Exporting Backup Policies

To export a backup policy to a .JSON file:

1. Navigate to **Policies**.
2. Switch to the necessary tab and select the backup policy whose settings you want to export.
3. Click **Advanced > Export Policy**.

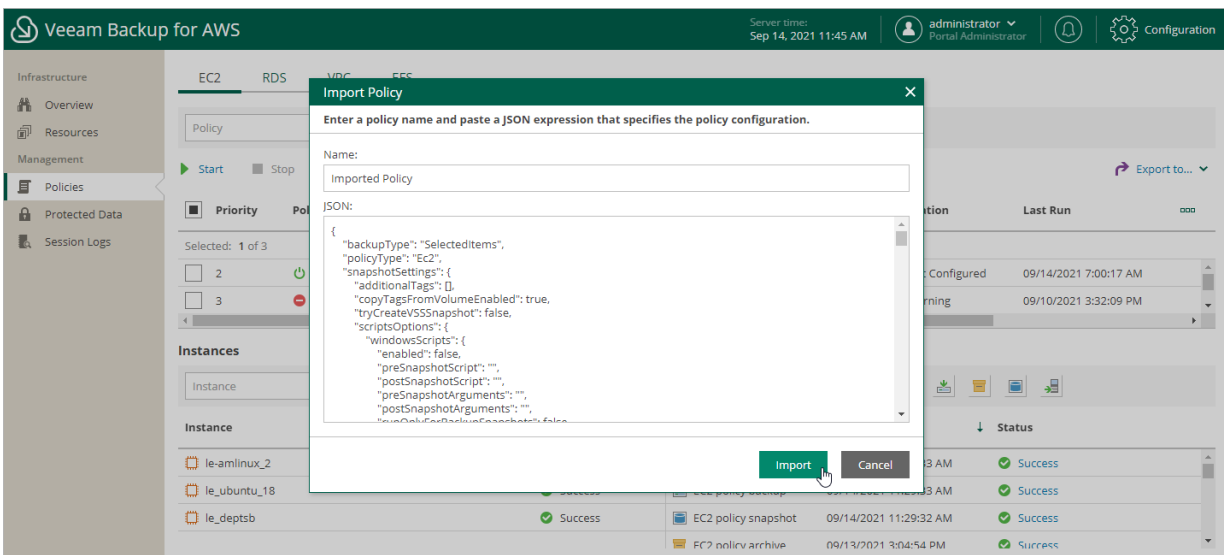
Veeam Backup for AWS will save the backup policy settings as a single .JSON file to the default download directory on the local machine.



Importing Backup Policies

To import a backup policy from a .JSON file:

1. Navigate to **Policies**.
2. Switch to the necessary tab and click **Advanced > Import Policy**.
3. In the **Import Policy** window, specify a name for the imported backup policy, paste the content of the necessary .JSON file, and click **Apply**.



Managing Backed-Up Data

The actions that you can perform with backed-up data depend on whether you access the data using the Veeam Backup & Replication console or the Veeam Backup for AWS Web UI.









Managing Backed-Up Data Using Console

To view and manage backed-up data, navigate to the **Backups** node of the **Home** view. The node displays information on all restore points created by backup appliances.

NOTE

You cannot remove created image-level backups and snapshots from the Veeam Backup & Replication console. To remove restore points of EC2 instances, RDS resources, DynamoDB tables, Redshift clusters, EFS file systems, FSx file systems and VPC configurations, open the backup [appliance Web UI](#) and follow the instructions provided in section [Managing Backed-up Data using Web UI](#).

When you expand the **Backups** node in the working area, you can see the following icons:

Icon	Protected Workload
	Indicates that the protected workload is an EC2 instance.
	Indicates that the protected workload is an DB instance.
	Indicates that the protected workload is an Aurora DB cluster.
	Indicates that the protected workload is a DynamoDB table.
	Indicates that the protected workload is a Redshift cluster.
	Indicates that the protected workload is an EFS file system.
	Indicates that the protected workload is an FSx file system.
	Indicates that the protected workload is a VPC configuration.

The **Backups** node contains 4 subnodes:

- The **Snapshots** subnode displays information on cloud-native snapshots of the protected EC2 instances and RDS resources, as well as information on cloud-native backups of the protected Redshift clusters, DynamoDB tables, EFS and FSx file systems:
 - *<appliance_name>* nodes show snapshots or backups created manually on backup appliances and snapshots or backups imported to the backup appliances from AWS Regions specified in backup policy settings.
 - *<backup_policy_name>* nodes show snapshots or backups created by backup policies.

To learn how Veeam Backup for AWS creates cloud-native snapshots of EC2 instances and RDS resources, as well as cloud-native backups of DynamoDB tables, Redshift clusters, EFS file systems and FSx file systems, see sections [EC2 Backup](#), [RDS Backup](#), [DynamoDB Backup](#), [Redshift Backup](#), [EFS Backup](#) and [FSx Backup](#).

- The **External Repository** subnode displays information on image-level backups of the protected EC2 instances and RDS resources that are stored in standard backup repositories, as well as backups of VPC configurations that are stored on backup appliances.
 - *<backup_policy_name>* nodes show backups of EC2 instances and RDS resources created by backup policies.
 - *<aws_account_name>* nodes show VPC configuration backups created for specific AWS accounts.

To learn how Veeam Backup for AWS creates image-level backups of EC2 instances and RDS resources, as well as VPC configuration backups, see sections [EC2 Backup](#), [RDS Backup](#) and [VPC Configuration Backup](#).

NOTE

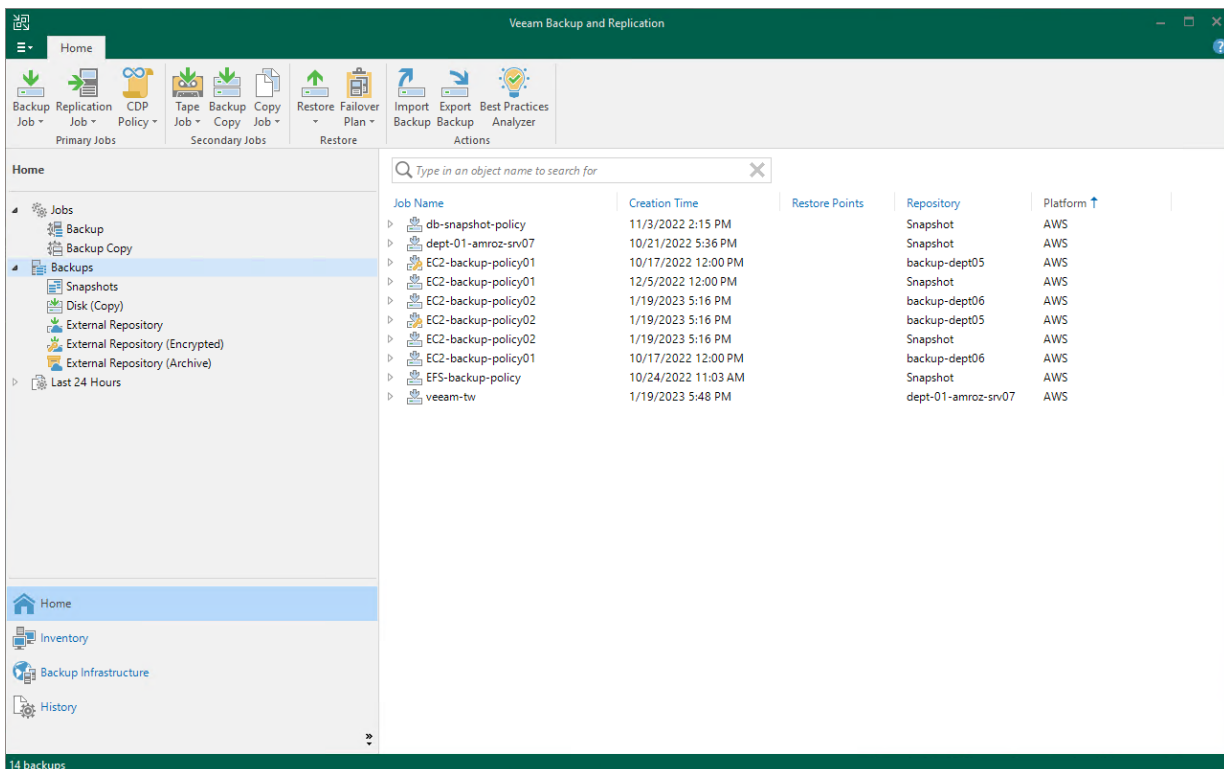
If a backup chain was originally encrypted and then got decrypted by Veeam Backup & Replication, the backup chain will be marked with the **Key** icon.

- The **External Repository (Encrypted)** subnode displays information on encrypted image-level backups of EC2 instances and RDS resources that are stored in standard backup repositories and that have not been decrypted yet, which means either that you have not specified the decryption password or that the specified password is invalid.

To learn how to decrypt backups, see [Decrypting Backups](#).

- The **External Repository (Archive)** subnode displays information on image-level backups of EC2 instances and RDS resources that are stored in archive backup repositories.

To learn how Veeam Backup for AWS creates archive backups, see [EC2 Archive Backup Chain](#) and [RDS Archive Backup Chain](#).



Decrypting Backups

Veeam Backup & Replication automatically decrypts backup files stored in repositories either using passwords that you specify when adding these repositories to the backup infrastructure or using KMS keys automatically detected by Veeam Backup & Replication. If you do not specify decryption passwords or if Veeam Backup & Replication does not have permissions to access KMS keys, the backup files remain encrypted.

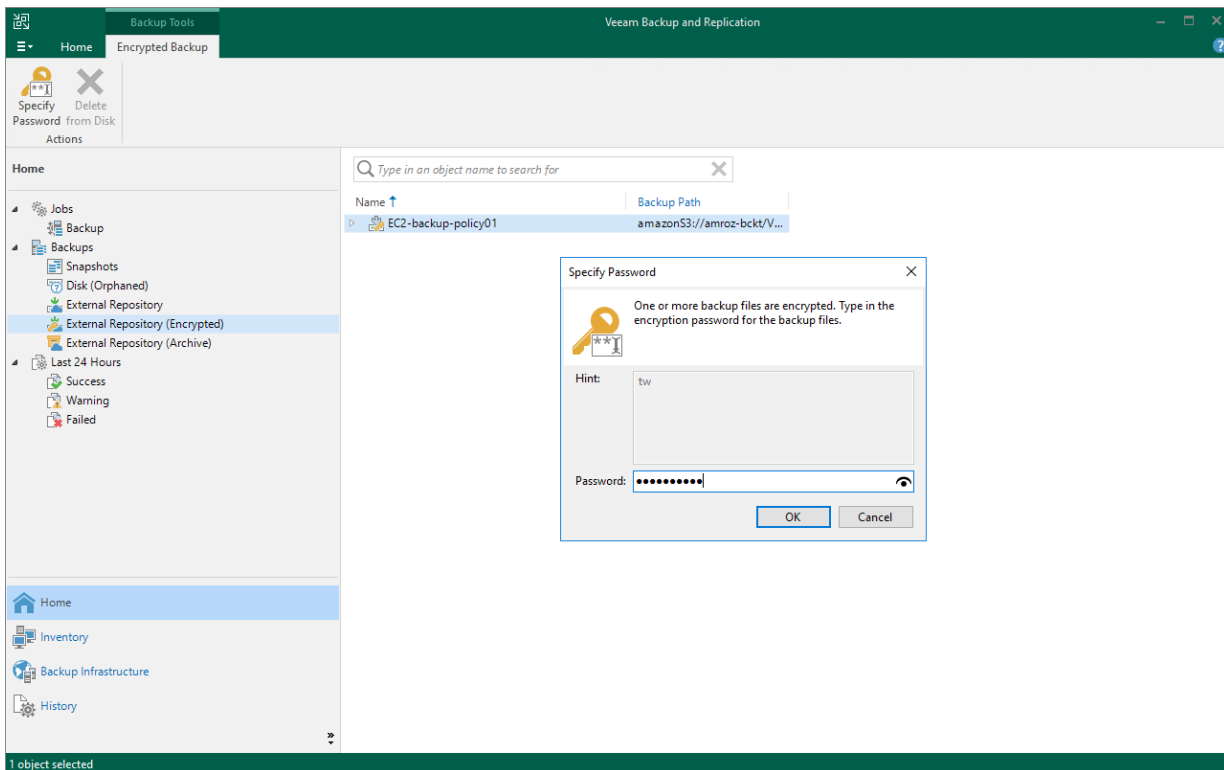
- To decrypt backup files encrypted using a KMS key, make sure that the IAM user specified when [creating a new repository](#) or [adding an existing repository](#) is assigned permissions required to access KMS keys. For more information on the required permissions, see [Plug-in Permissions](#).
- To decrypt backup files encrypted using a password, do the following:
 - a. In the Veeam Backup & Replication console, open the **Home** view.
 - b. Navigate to **Backups > External Repository (Encrypted)**.
 - c. Expand the backup policy that protects an AWS resource whose image-level backup you want to decrypt, select the backup chain that belongs to the resource and click **Specify Password** on the ribbon.

Alternatively, you can right-click the necessary backup chain and select **Specify password**.

TIP

To decrypt all backups created by a backup policy, right-click the policy and select **Specify Password**.

- d. In the **Specify Password** window, enter the password that was used to encrypt the data stored in the target repository.



Managing Backed-Up Data Using Web UI

Veeam Backup for AWS stores information on all protected AWS resources in the configuration database. Even if a resource is no longer protected by any configured backup policy and even if the resource no longer exists in AWS, information on the backed-up data will not be deleted from the database until Veeam Backup for AWS automatically removes all restore points associated with this resource according to the retention settings saved in the backup metadata. You can also remove the restore points manually on the **Protected Data** page.

NOTE

Veeam Backup for AWS does not include restore points created manually in backup and snapshot chains, and does not apply the configured retention policy settings to these restore points. This means that the restore points are kept in your AWS environment unless you remove them manually, as described in sections [Removing EC2 Snapshots Created Manually](#), [Removing RDS Snapshots Created Manually](#), [Removing DynamoDB Backups Created Manually](#), [Removing Redshift Backups Created Manually](#), [Removing EFS Backups Created Manually](#) and [Removing FSx Backups Created Manually](#).

EC2 Data

After a backup policy successfully creates a restore point of an EC2 instance according to the specified schedule, or after you create a snapshot of an EC2 instance manually, Veeam Backup for AWS adds the instance to the resource list on the **Protected Data** page.

For each backed-up EC2 instance, Veeam Backup for AWS creates a record in the configuration database with the following set of properties:

- **Instance** – the name of the EC2 instance.
- **Policy** – the name of the backup policy that processed the EC2 instance.
- **Restore Points** – the number of restore points created for the EC2 instance.

To view the list of restore points, click the link in the **Restore Points** column. The **Available Restore Points** window will display information on each restore point, including the following: the date when the restore point was created, the type of the restore point, the region where the restore point is stored, the state of the restore point (for image-level backups), the name and storage class of the backup repository where the restore point is stored, and the configured retention policy settings (*D* – daily, *W* – weekly, *M* – monthly or *Y* – yearly).

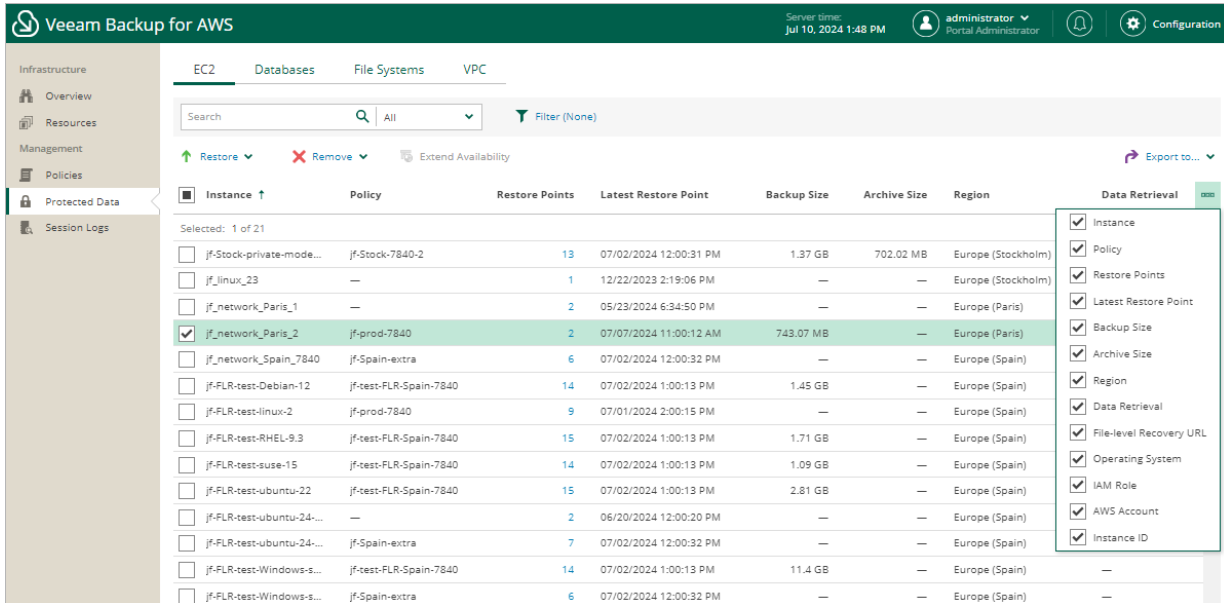
- **Latest Restore Point** – the date and time of the latest restore point that was created for the EC2 instance.
- **Backup Size** – the size of all backups created for the selected EC2 instance stored in standard repositories.
- **Archive size** – the size of all backups created for the selected EC2 instance stored in archive backup repositories.
- **Region** – the AWS Region in which the EC2 instance resides.
- **Data Retrieval** – shows whether any of the archived restore points of the EC2 instance is retrieved.
- **File-level Recovery URL** – a link to the file-level recovery browser.

The link appears when the file-level recovery session is started for the selected EC2 instance. The link contains a DNS name of the worker instance hosting the file-level recovery browser and authentication information used to access this worker instance.

- **Operating System** – the operating system running on the EC2 instance.
- **IAM Role** – the IAM Role used to back up the EC2 instance.
- **AWS Account** – the AWS account to which the EC2 instance belongs.
- **Instance ID** – the AWS ID of the EC2 instance.

On the **Protected Data** page, you can also perform the following actions:

- Remove restore points if you no longer need them. For more information, see sections [Removing EC2 Backups and Snapshots](#) and [Removing EC2 Snapshots Created Manually](#).
- Retrieve image-level backups stored in archive backup repositories. For more information, see [Retrieving EC2 Data From Archive](#).
- Restore data of backed-up EC2 instances. For more information, see [EC2 Restore Using Web UI](#).



Removing EC2 Backups and Snapshots

Veeam Backup for AWS applies the [configured retention policy settings](#) to automatically remove cloud-native snapshots, snapshot replicas and image-level backups created by backup policies. If necessary, you can also remove the backed-up data manually.

IMPORTANT

Do not delete backup files from Amazon S3 buckets in the AWS Management Console. If some file in a backup chain is missing, you will not be able to roll back EC2 instance data to the necessary state.

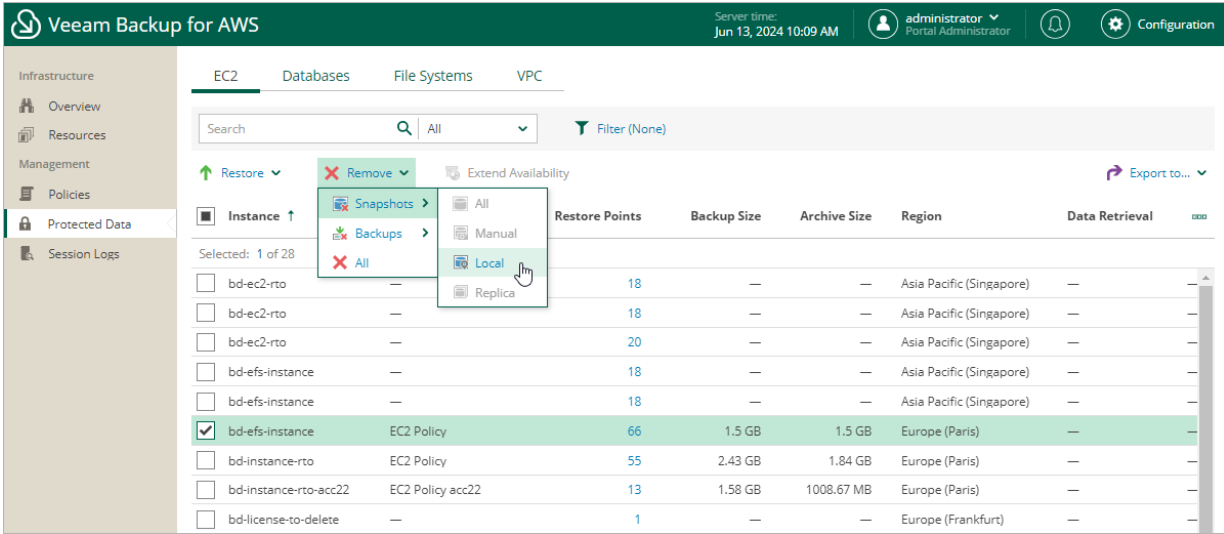
To remove backed-up data manually, do the following:

1. Navigate to **Protected Data > EC2**.
2. Select EC2 instances whose data you want to remove.
3. Click **Remove** and select either of the following options:
 - **Snapshots > All** – to remove all cloud-native snapshots and snapshot replicas created for the selected EC2 instances both by backup policies and manually.
 - **Snapshots > Manual** – to remove cloud-native snapshots created for the selected EC2 instances manually.

If you want to remove only specific cloud-native snapshots, follow the instructions provided in section [Removing Snapshots Created Manually](#).

 - **Snapshots > Local** – to remove cloud-native snapshots created for the selected EC2 instances by backup policies.
 - **Snapshots > Replica** – to remove snapshot replicas created for the selected EC2 instances by backup policies.
 - **Backups > All** – to remove all backups created for the selected EC2 instances.
 - **Backups > Standard** – to remove all standard backups created for the selected EC2 instances.

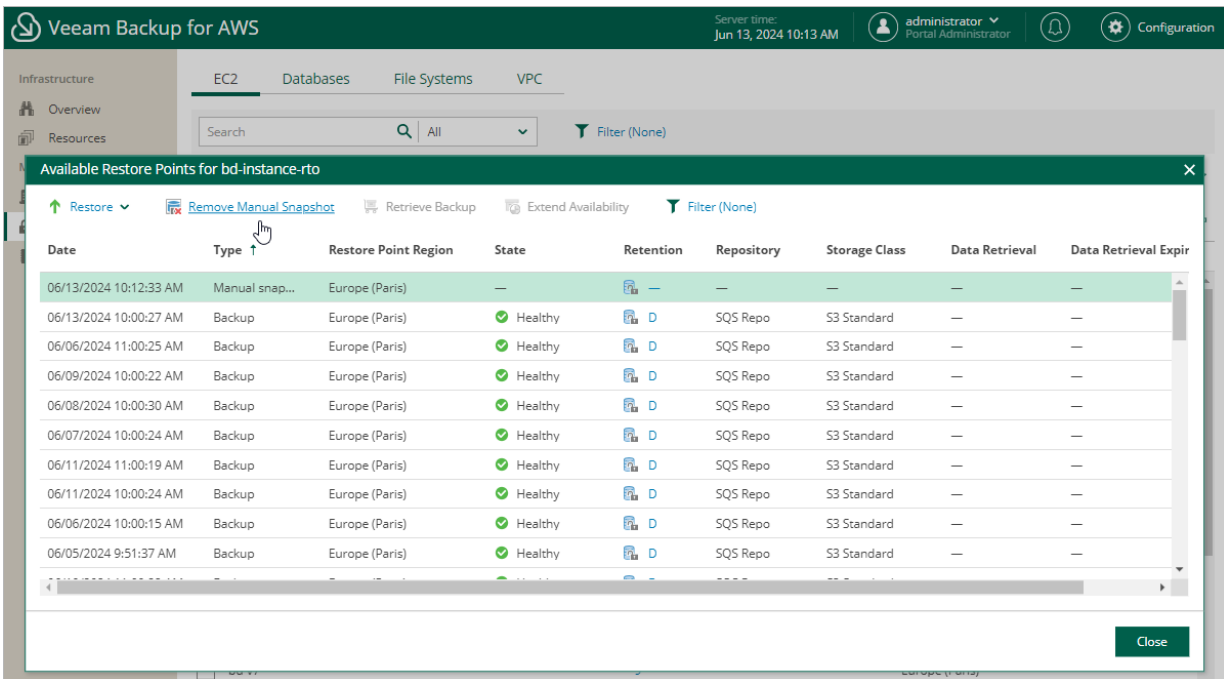
- **Backups > Archived** – to remove all archived backups created for the selected EC2 instances.
- **All** – to remove all cloud-native snapshots, snapshot replicas, and image-level backups created for the selected EC2 instances both by backup policies and manually.



Removing EC2 Snapshots Created Manually

To remove all cloud-native snapshots created for an EC2 instance manually, follow the instructions provided in the [Removing EC2 Backups and Snapshots](#) section. If you want to remove a specific snapshot created manually, do the following:

1. Navigate to **Protected Data > EC2**.
2. Select the necessary instance, and click the link in the **Restore Points** column.
3. In the **Available Restore Points** window, select a snapshot that you want to remove, and click **Remove Manual Snapshot**.



Retrieving EC2 Data From Archive

Backups stored in archive backup repositories are not immediately accessible. If you want to restore an EC2 instance from a backup that is stored in a repository of the S3 Glacier Flexible Retrieval or S3 Glacier Deep Archive storage class, you must first retrieve the archived data. During the data retrieval process, a temporary copy of the archived data is created in an Amazon S3 bucket where the repository is located. This copy is stored in the S3 standard storage class for a period of time that you specify when launching the data retrieval process. If the time period expires while a restore operation is still running, Veeam Backup for AWS automatically extends the period to keep the retrieved data available for 1 more day. You can also [extend the availability period manually](#).

To retrieve archived data, you can launch the data retrieval process either from the [Data Retrieval wizard](#) before you begin a restore operation, or directly from the [Restore wizard](#). When you retrieve archived data, you can choose one of the following options:

- **Expedited** – the most expensive option. The retrieved data is available within 1-5 minutes.
Amazon does not support this option for data stored in the S3 Glacier Deep Archive storage class. For more information, see [AWS Documentation](#).
- **Standard** – the recommended option. The retrieved data is available within 3-5 hours for data stored in the S3 Glacier Flexible Retrieval storage class and within 12 hours for data stored in the S3 Glacier Deep Archive storage class.
- **Bulk** – the least expensive option. The retrieved data is available within 5-12 hours for data stored in the S3 Glacier Flexible Retrieval storage class and within 48 hours for data stored in the S3 Glacier Deep Archive storage class.
- **Standard accelerated** – the option that is less expensive than the **Expedited** option. The retrieved data is available within 15-30 minutes for data stored in the S3 Glacier Flexible Retrieval storage class.
With this option enabled, Veeam Backup for AWS leverages the [S3 Batch Operations functionality](#) to retrieve the archived data.

TIP

Before you enable the **Standard accelerated** option, it is recommended that you check whether the IAM role specified to access the archive backup repository has all the permissions required to perform data retrieval operations using the S3 Batch Operations functionality, as described in section [Checking IAM Role Permissions](#).

If some of the IAM role permissions required to perform data retrieval operations using the S3 Batch Operations functionality are missing, Veeam Backup for AWS will use the **Standard** option to retrieve data.


For more information on archive retrieval options, see [AWS Documentation](#).

Retrieving Data Manually

To retrieve archived data of an EC2 instance, do the following:

1. Navigate to **Protected Data > EC2**.
2. Select the necessary instance, and click the link in the **Restore Points** column.

3. In the **Available Restore Points** window, select a restore point that contains archived data you want to retrieve, and click **Retrieve Backup**. The **Data Retrieval** wizard will open.



4. At the **Settings** step of the wizard, specify the following settings:
 - a. In the **Retrieval mode** section, select the **retrieval option** that Veeam Backup for AWS will use to retrieve the data.
 - b. In the **Availability period** section, specify the number of days for which you want to keep the data available for restore operations.

You will be able to **manually extend data availability** later if required.

TIP

If you want to receive an email notification when the data availability period is about to expire, select the **Send notification email** check box and specify the number of hours before the expiration time when the notification will be sent.

The screenshot shows the 'Data Retrieval' wizard in Veeam Backup for AWS. The current step is 'Configure data retrieval settings'. The interface includes a top navigation bar with the Veeam logo, server time (Oct 12, 2023 12:37 PM), and user information (administrator, Portal Administrator). A left sidebar shows 'Settings' and 'Summary' tabs. The main content area contains the following settings:

- Retrieval mode:** Four radio button options: Expedited, Standard accelerated, Standard (selected), and Bulk. Each option includes a brief description of its performance and cost characteristics.
- Availability period:** A section with a text input field set to '3' days, a 'Send email notification' checkbox checked and set to '2' hours before data expires, and a 'Notify when data retrieval completes' checkbox checked.

At the bottom of the wizard, there are 'Next' and 'Cancel' buttons.

5. At the **Summary** step of the wizard, review configuration information and click **Finish**.

The screenshot shows the 'Data Retrieval' wizard in Veeam Backup for AWS, now at the 'Review configured settings' step. The 'Summary' tab is selected in the left sidebar. The main content area displays a summary of the configured settings:

- Retrieval settings:** A summary bar indicating '1 will be retrieved'.
- Retrieval mode:** A summary bar showing 'Retrieval mode: Standard'.
- Availability period:** A summary bar showing 'Availability period: 3 days' and 'Email notification: Enabled (2 hours before data expiration)'.

Below the summary, there is an information icon and a message: 'After you click Finish, you can track data retrieval progress in the notification area.' At the bottom of the wizard, there are 'Previous', 'Finish', and 'Cancel' buttons.

IMPORTANT

If you cancel the Data Retrieval session, or the Veeam Backup for AWS service is restarted while the Data Retrieval session is still running, AWS will retrieve data anyway and keep it for the specified availability period. However, Veeam Backup for AWS will not be able to access the retrieved data.


Extending Data Availability

To extend time for which you want to keep retrieved data available for restore operations:

1. Navigate to **Protected Data > EC2**.
2. Select the EC2 instance for which you want to extend availability of the retrieved data.
3. Click **Extend Availability**.

Alternatively, click the link in the **Restore Points** column. In the **Data Retrieval** window, select the restore point that contains the retrieved data, and click **Extend Availability**.

4. In the **Extend Data Availability Period** window, specify the number of days for which you want to keep the data available for restore operations, and click **Extend**.



RDS Data

After a backup policy successfully creates a restore point of an RDS resource according to the specified schedule, or after you create a snapshot of an RDS resource manually, Veeam Backup for AWS adds the resource to the resource list on the **Protected Data** page.

For each backed-up RDS resource, Veeam Backup for AWS creates a record in the configuration database with the following set of properties:


- **Instance** – the name of the DB instance or Aurora DB cluster.
- **Policy** – the name of the backup policy that processed the DB instance or Aurora DB cluster.
- **Restore Points** – the number of restore points created for the DB instance or Aurora DB cluster.

To view the list of restore points, click the link in the **Restore Points** column. The **Available Restore Points** window will display information on each restore point, including the following: the date when the restore point was created, the type of the restore point, the region where the restore point is stored, the state of the restore point (for image-level backups), the name and storage class of the backup repository where the restore point is stored, and the configured retention policy settings (*D* – daily, *W* – weekly, *M* – monthly or *Y* – yearly).

- **Latest Restore Point** – the date and time of the latest restore point that was created for the DB instance or Aurora DB cluster.
- **Backup Size** – the size of all backups created for the PostgreSQL DB instance stored in standard backup repositories.
- **Archive size** – the size of all backups created for the PostgreSQL DB instance stored in archive backup repositories.
- **Engine** – a database engine of the DB instance or Aurora DB cluster.
- **Instance Size** – the size of the DB instance storage.
- **AWS Account** – the AWS account to which the DB instance or Aurora DB cluster belongs.
- **Instance ID** – the AWS ID of the DB instance or Aurora DB cluster.
- **Region** – the AWS Region in which the DB instance or Aurora DB cluster resides.

On the **Protected Data** page, you can also perform the following actions:

- Remove restore points if you no longer need them. For more information, see sections [Removing RDS Backups and Snapshots](#) and [Removing RDS Snapshots Created Manually](#).
- Restore data of backed-up RDS resources. For more information, see [RDS Restore Using Web UI](#).



Removing RDS Backups and Snapshots

Veeam Backup for AWS applies the [configured retention policy settings](#) to automatically remove cloud-native snapshots and snapshot replicas and image-level backups created by backup policies. If necessary, you can also remove the backed-up data manually.

IMPORTANT

Consider the following:

- Do not delete backup files from Amazon S3 buckets in the AWS Management Console. If some file in a backup chain is missing, you will not be able to roll back DB instance or Aurora DB cluster data to the necessary state.
- In Veeam Backup for AWS, you can remove only snapshots created by the Veeam backup service. To delete AWS Snapshots (DB instance snapshots and DB cluster snapshots created in AWS), use [Amazon Management Console](#).


To remove backed-up data manually, do the following:

1. Navigate to **Protected Data > Databases > RDS**.
2. Select RDS resources whose data you want to remove.
3. Click **Remove** and select either of the following options:
 - **Snapshots > All** – to remove all cloud-native snapshots and snapshot replicas created for the selected RDS resources both by backup policies and manually.
 - **Snapshots > Manual Snapshots** – to remove cloud-native snapshots created for the selected RDS resources manually.

If you want to remove only specific cloud-native snapshots, follow the instructions provided in section [Removing Snapshots Created Manually](#).

- **Snapshots > Local Snapshots** – to remove cloud-native snapshots created for the selected RDS resources by backup policies.

- **Snapshots > Replica** – to remove snapshot replicas created for the selected RDS resources by backup policies.
- **Backups > All** – to remove all backups created for the selected RDS resources.
- **Backups > Standard** – to remove all standard backups created for the selected RDS resources.
- **Backups > Archived** – to remove all archived backups created for the selected RDS resources.
- **All** – to remove all cloud-native snapshots, snapshot replicas, and image-level backups created for the selected RDS resources both by backup policies and manually.




Removing RDS Snapshots Created Manually

To remove all cloud-native snapshots created for a DB instance or an Aurora DB cluster manually, follow the instructions provided in the [Removing RDS Backups and Snapshots](#) section. If you want to remove a specific snapshot created manually, do the following:

1. Navigate to **Protected Data > Databases > RDS**.
2. Select the necessary resource, and click the link in the **Restore Points** column.

3. In the **Available Restore Points** window, select a snapshot that you want to remove, and click **Remove Manual Snapshot**.



DynamoDB Data

After a backup policy successfully creates a restore point of a DynamoDB table according to the specified schedule, or after you create a backup of a DynamoDB table manually, Veeam Backup for AWS adds the table to the resource list on the **Protected Data** page.

For each backed-up DynamoDB table, Veeam Backup for AWS creates a record in the configuration database with the following set of properties:

- **Name** – the name of the DynamoDB table.
- **Table ID** – the AWS ID of the of the DynamoDB table.
- **Policy** – the name of the backup policy that processed the DynamoDB table.
- **Restore Points** – the number of restore points created for the DynamoDB table.

To view the list of restore points, click the link in the **Restore Points** column. The **Available Restore Points** window will display information on each restore point, including the following: the date when the restore point was created, the size and type of the restore point, the backup vault where the restore point is stored, and the configured retention policy settings (*D*– daily, *W*– weekly, *M*– monthly or *Y*– yearly).

- **Latest Restore Point** – the date and time of the latest restore point that was created for the DynamoDB table.
- **Backup Size** – the size of all backups created for the DynamoDB table stored in backup vaults.
- **Region** – the AWS Region in which the DynamoDB table resides.
- **AWS Account** – the AWS account to which the DynamoDB table belong.

On the **Protected Data** page, you can also perform the following actions:

- Remove restore points if you no longer need them. For more information, see sections [Removing DynamoDB Backups](#) and [Removing DynamoDB Backups Created Manually](#).
- Restore data of backed-up DynamoDB tables. For more information, see [DynamoDB Restore Using Web UI](#).

The screenshot shows the Veeam Backup for AWS interface. The top navigation bar includes the Veeam logo, the text 'Veeam Backup for AWS', the server time 'Jun 13, 2024 10:38 AM', and the user 'administrator Portal Administrator'. The left sidebar shows a navigation menu with 'Protected Data' selected. The main content area is titled 'Databases' and contains a table of DynamoDB tables. The table has columns for Name, Policy, Restore Points, Latest Restore Point, Backup Size, Region, and AWS Account. A context menu is open over the table, showing options to export data to various formats: Name, Table ID, Policy, Restore Points, Latest Restore Point, Backup Size, Region, and AWS Account.

Name	Policy	Restore Points	Latest Restore Point	Backup Size	Region	AWS Account
bd-milan-table	DynamoDB	332	06/13/2024 10:00:18 AM	0 Bytes	Europe (Milan)	
bd-singapore-en...	—	2	11/09/2023 10:28:17 AM	0 Bytes	Asia Pacific (Singap...	
bd-singapore-ta...	—	2	10/19/2023 2:16:06 PM	0 Bytes	Asia Pacific (Singap...	
bd-singapore-ta...	—	1	10/16/2023 1:42:11 PM	12 Bytes	Asia Pacific (Singap...	
bd-singapore-ta...	—	2	11/09/2023 4:31:57 PM	78 Bytes	Asia Pacific (Singap...	
khromeev-test	—	1	04/15/2024 9:14:07 AM	0 Bytes	US East (N. Virginia)	
nm-metadata-7	—	1	02/09/2024 9:34:12 AM	0 Bytes	Europe (London)	
nm-temp	—	1	05/10/2024 1:41:24 PM	0 Bytes	US East (N. Virginia)	487569979969 (v...
pi-paris-newtable	—	1	10/11/2023 8:47:14 AM	0 Bytes	Europe (Paris)	487569979969
pi-streams-resto...	—	1	10/10/2023 3:41:33 PM	0 Bytes	Europe (Ireland)	487569979969 (v...

Removing DynamoDB Backups

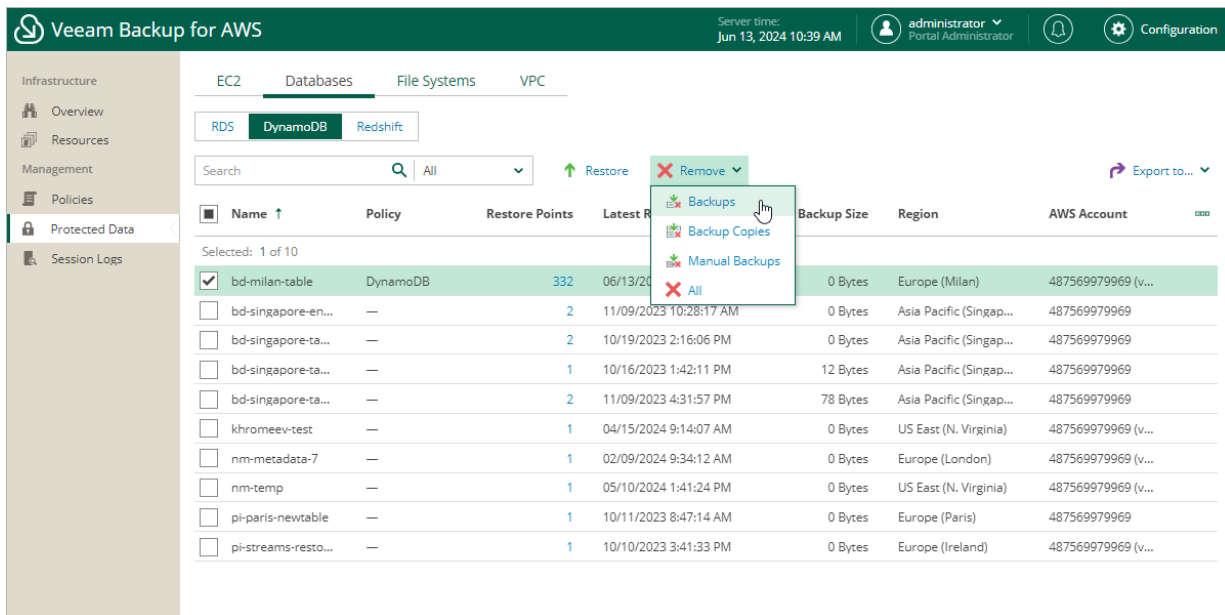
Veeam Backup for AWS applies the [configured retention policy settings](#) to automatically remove DynamoDB backups and backup copies created by backup policies. If necessary, you can also remove the backed-up data manually.

To remove backed-up data manually, do the following:

1. Navigate to **Protected Data > Databases > DynamoDB**.
2. Select DynamoDB table whose data you want to remove.
3. Click **Remove** and select either of the following options:
 - **Backups** – to remove DynamoDB backups created for the selected table by backup policies.
 - **Backup Copies** – to remove backup copies created for the selected table by backup policies.
 - **Manual Backups** – to remove DynamoDB backups created for the selected table manually.

If you want to remove only specific manual backup, follow the instructions provided in section [Removing DynamoDB Backups Created Manually](#).

- **All** – to remove all backups and backup copies created for the selected tables both by backup policies and manually.




Removing DynamoDB Backups Created Manually

To remove all backups created for a DynamoDB table manually, follow the instructions provided in the [Removing DynamoDB Backups](#) section. If you want to remove a specific DynamoDB backup created manually, do the following:

1. Navigate to **Protected Data > Databases > DynamoDB**.
2. Select the necessary table, and click the link in the **Restore Points** column.

3. In the **Available Restore Points** window, select a backup that you want to remove, and click **Remove Manual Backup**.



Redshift Data

After a backup policy successfully creates a restore point of a Redshift cluster according to the specified schedule, or after you create a backup of a Redshift cluster manually, Veeam Backup for AWS adds the cluster to the resource list on the **Protected Data** page.

For each backed-up Redshift cluster, Veeam Backup for AWS creates a record in the configuration database with the following set of properties:

- **Cluster** – the name of the Redshift cluster.
- **Policy** – the name of the backup policy that processed the Redshift cluster.
- **Restore Points** – the number of restore points created for the Redshift cluster.

To view the list of restore points, click the link in the **Restore Points** column. The **Available Restore Points** window will display information on each restore point, including the following: the date when the restore point was created, the size and type of the restore point, the backup vault where the restore point is stored, and the configured retention policy settings (*D*– daily, *W*– weekly, *M*– monthly or *Y*– yearly).

- **Latest Restore Point** – the date and time of the latest restore point that was created for the Redshift cluster.
- **Cluster ID** – the AWS ID of the Redshift cluster.
- **AWS Account** – the AWS account to which the Redshift cluster belong.
- **Backup Size** – the size of all backups created for the Redshift cluster stored in backup vaults.
- **Region** – the AWS Region in which the Redshift cluster resides.

On the **Protected Data** page, you can also perform the following actions:

- Remove restore points if you no longer need them. For more information, see sections [Removing Redshift Backups](#) and [Removing Redshift Backups Created Manually](#).
- Restore data of backed-up Redshift cluster. For more information, see [Redshift Restore Using Web UI](#).

The screenshot shows the Veeam Backup for AWS interface. The top navigation bar includes the Veeam logo, the product name, the server time (Jul 8, 2024 3:17 PM), the user (administrator), and a Configuration button. The left sidebar shows the navigation menu with 'Protected Data' selected. The main content area is titled 'Redshift' and contains a table with the following columns: Cluster, Policy, Restore Points, Latest Restore Point, Backup Size, AWS Acco..., Cluster ID, and Region. A context menu is open over the table, showing options: Cluster, Policy, Restore Points, Latest Restore Point, Backup Size, AWS Account, Cluster ID, and Region, all of which are checked.

Cluster	Policy	Restore Points	Latest Restore Point	Backup Size	AWS Acco...	Cluster ID	Region
jf-cluster-relocatio...	—	1	06/21/2024 3:59:55 ...	74 MB	942676447...	27021684-4	
jf-custom-encrypt...	—	2	06/20/2024 12:39:36...	159 MB	942676447...	8b2806de-5	
jf-custom-to-default	—	1	06/20/2024 12:39:36...	86 MB	942676447...	de4e83ae-2	
jf-default-encrypt...	—	2	06/20/2024 12:40:46...	90 MB	942676447...	5809ec99-c	
jf-default-encrypt...	—	1	06/20/2024 12:41:17...	48 MB	942676447...	2fe58613-d	
jf-default-to-custom	—	1	06/20/2024 12:39:36...	49 MB	942676447...	578dcf89-5	
jf-external-encrypt...	—	1	06/20/2024 12:41:47...	299 MB	942676447...	d78d147d-4	
jf-multi-color-stock	—	2	05/06/2024 1:15:41 ...	148 MB	942676447...	fd998be-22b...	US West (Oregon)
jf-notificators-rete...	—	1	05/31/2024 11:00:18...	66 MB	942676447...	232d762b-05...	Asia Pacific (Mumbai)
jf-reloc-cluster	—	1	05/06/2024 1:15:41 ...	37 MB	942676447...	6de3e638-eaf...	US West (Oregon)
jf-reloc-cluster	—	1	05/08/2024 9:14:47 ...	48 MB	942676447...	6652f9a3-f19...	US West (Oregon)
jf-to-delete-cluster	—	1	05/16/2024 10:12:39...	70 MB	942676447...	afcfd8b2-deb...	US West (Oregon)
jf-unencrypted-no...	—	2	06/20/2024 12:39:36...	88 MB	942676447...	d845eff7-e25...	US West (Oregon)
jf-unencrypted-to...	—	1	06/20/2024 12:40:46...	49 MB	942676447...	156e9b78-6a1...	US West (Oregon)
jf-unencrypted-to...	—	1	06/20/2024 12:39:36...	77 MB	942676447...	d422a191-f1d...	US West (Oregon)

Removing Redshift Backups


Veeam Backup for AWS applies the [configured retention policy settings](#) to automatically remove Redshift backups created by backup policies. If necessary, you can also remove the backed-up data manually.

To remove backed-up data manually, do the following:

1. Navigate to **Protected Data > Databases > Redshift**.
2. Select Redshift cluster whose data you want to remove.
3. Click **Remove** and select either of the following options:
 - **Backups** – to remove Redshift backups created for the selected cluster by backup policies.
 - **Manual Backups** – to remove Redshift backups created for the selected cluster manually.

If you want to remove only specific manual backup, follow the instructions provided in section [Removing Redshift Backups Created Manually](#).

- **All** – to remove all backups created for the selected clusters both by backup policies and manually.




Removing Redshift Backups Created Manually

To remove all backups created for a Redshift cluster manually, follow the instructions provided in the [Removing Redshift Backups](#) section. If you want to remove a specific Redshift backup created manually, do the following:

1. Navigate to **Protected Data > Databases > Redshift**.
2. Select the necessary cluster, and click the link in the **Restore Points** column.

- In the **Available Restore Points** window, select a backup that you want to remove, and click **Remove Manual Backup**.



EFS Data

After a backup policy successfully creates a restore point of an EFS file system according to the specified schedule, or after you create a backup of an EFS file system manually, Veeam Backup for AWS adds the file system to the resource list on the **Protected Data** page.

For each backed-up EFS file system, Veeam Backup for AWS creates a record in the configuration database with the following set of properties:

- **Name** – the name of the EFS file system.
- **Policy** – the name of the backup policy that processed the EFS file system.
- **Restore Points** – the number of restore points created for the EFS file system.

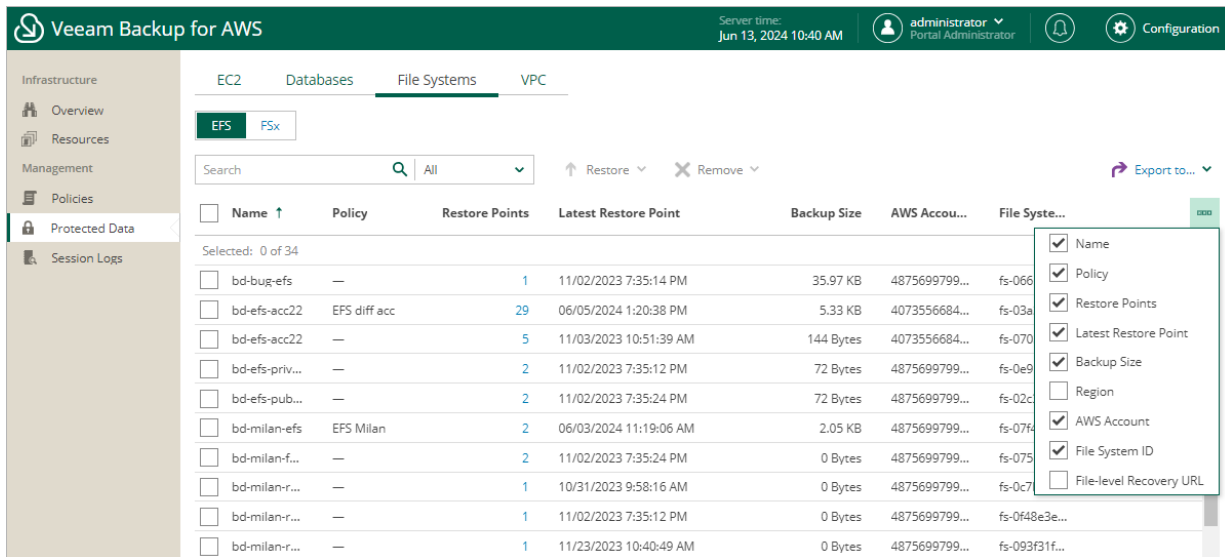
To view the list of restore points, click the link in the **Restore Points** column. The **Available Restore Points** window will display information on each restore point, including the following: the date when the restore point was created, the size and type of the restore point, the backup vault where the restore point is stored, and the configured retention policy settings (*D*– daily, *W*– weekly, *M*– monthly or *Y*– yearly).

- **Latest Restore Point** – the date and time of the latest restore point that was created for the EFS file system.
- **Backup Size** – the size of all backups created for the EFS file system.
- **Region** – the AWS Region in which the EFS file system resides.
- **AWS Account** – the AWS account to which the EFS file system belong.
- **File System ID** – the AWS ID of the EFS file system.
- **File-level Recovery URL** – a link to the file-level recovery browser.

The link appears when the restore session is started for the file-level recovery process. The link contains a public DNS name or an IP address of the backup appliance hosting the file-level recovery browser and authentication information used to access the appliance.

On the **Protected Data** page, you can also perform the following actions:

- Remove restore points if you no longer need them. For more information, see sections [Removing EFS Backups](#) and [Removing EFS Backups Created Manually](#).
- Restore data of backed-up EFS file system. For more information, see [EFS Restore Using Web UI](#).



Removing EFS Backups


Veeam Backup for AWS applies the [configured retention policy settings](#) to automatically remove EFS file system backups and backup copies created by backup policies. If necessary, you can also remove the backed-up data manually.

To remove backed-up data manually, do the following:

1. Navigate to **Protected Data > File Systems > EFS**.
2. Select EFS file systems whose data you want to remove.
3. Click **Remove** and select either of the following options:
 - **Backups** – to remove EFS backups created for the selected file systems by backup policies.
 - **Backup Copies** – to remove backup copies created for the selected file systems by backup policies.
 - **Manual Backups** – to remove EFS backups created for the selected file systems manually.

If you want to remove only specific manual backup, follow the instructions provided in section [Removing EFS Backups Created Manually](#).

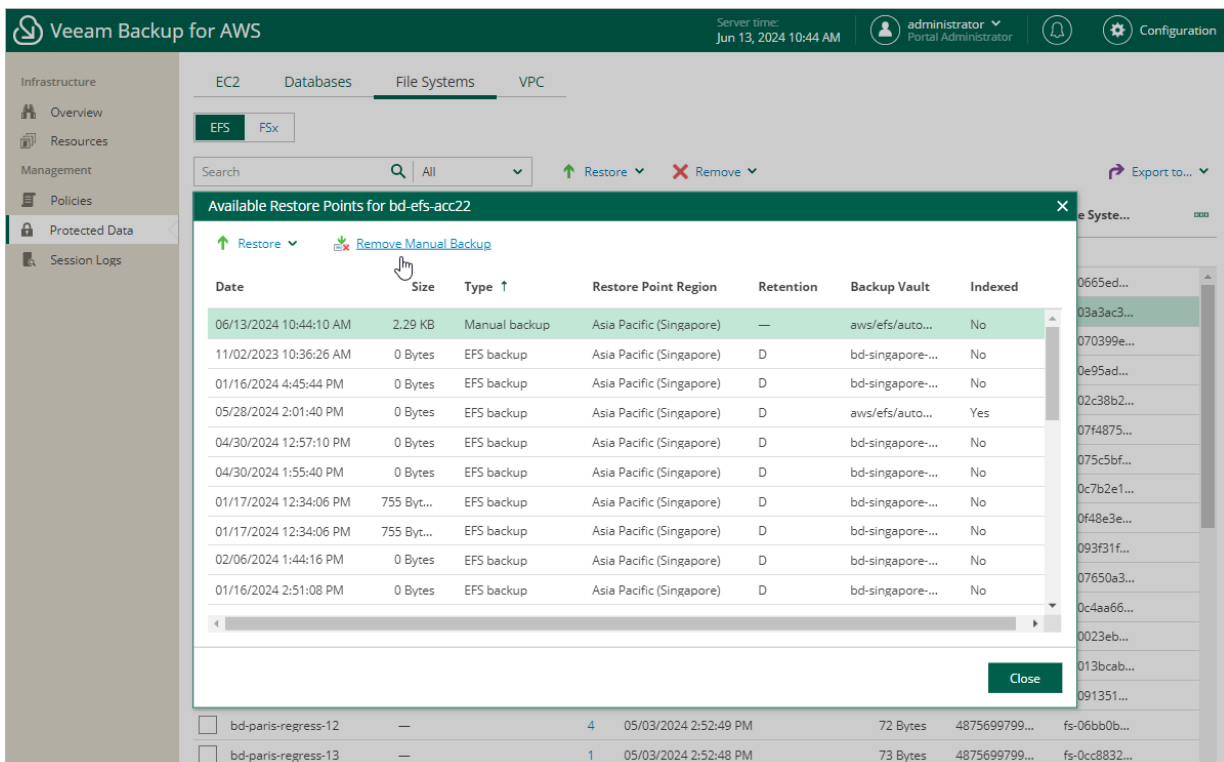
- **All** – to remove all backups and backup copies created for the selected file systems both by backup policies and manually.



Removing EFS Backups Created Manually

To remove all backups created for an EFS file system manually, follow the instructions provided in the [Removing EFS Backups](#) section. If you want to remove a specific EFS backup created manually, do the following:

1. Navigate to **Protected Data > File Systems > EFS**.
2. Select the necessary file system, and click the link in the **Restore Points** column.
3. In the **Available Restore Points** window, select a backup that you want to remove, and click **Remove Manual Backup**.



FSx Data

After a backup policy successfully creates a restore point of an FSx file system according to the specified schedule, or after you create a backup of an FSx file system manually, Veeam Backup for AWS adds the file system to the resource list on the **Protected Data** page.

For each backed-up FSx file system, Veeam Backup for AWS creates a record in the configuration database with the following set of properties:

- **Name** – the name of the FSx file system.
- **Policy** – the name of the backup policy that processed the FSx file system.
- **Restore Points** – the number of restore points created for the FSx file system.

To view the list of restore points, click the link in the **Restore Points** column. The **Available Restore Points** window will display information on each restore point, including the following: the date when the restore point was created, the size and type of the restore point, the backup vault where the restore point is stored, and the configured retention policy settings (*D*– daily, *W*– weekly, *M*– monthly or *Y*– yearly).

- **Latest Restore Point** – the date and time of the latest restore point that was created for the FSx file system.
- **Region** – the AWS Region in which the FSx file system resides.
- **AWS Account** – the AWS account to which the FSx file system belong.
- **File System ID** – the AWS ID of the FSx file system.

On the **Protected Data** page, you can also perform the following actions:

- Remove restore points if you no longer need them. For more information, see sections [Removing FSx Backups](#) and [Removing FSx Backups Created Manually](#).
- Restore data of backed-up FSx file system. For more information, see [FSx Restore Using Web UI](#).

The screenshot shows the Veeam Backup for AWS interface. The top navigation bar includes the Veeam logo, server time (Jun 13, 2024 10:46 AM), and user information (administrator, Portal Administrator). The left sidebar contains navigation options: Infrastructure, Overview, Resources, Management, Policies, Protected Data (selected), and Session Logs. The main content area is titled 'File Systems' and shows a table of backed-up FSx file systems. The table has columns for Name, Policy, Restore Points, Latest Restore Point, Region, AWS Account, and File System ID. A context menu is open over the File System ID column, showing options: Name, Policy, Restore Points, Latest Restore Point, Region, AWS Account, and File System ID. The table contains 12 rows of data, including file systems like 'pi-encrypted-opz', 'nm-zfs-single-az-1', and 'nm-zfs-multi-az'.

Name	Policy	Restore Points	Latest Restore Point	Region	AWS Acco...	File System ID
pi-encrypted-opz	—	1	04/18/2024 10:14:14 PM	Europe (Ireland)	487569979...	fs-0b701902
nm-zfs-single-az-1	—	1	04/02/2024 12:16:31 PM	Canada (Central)	487569979...	fs-0e752241
nm-zfs-single-az-1	—	2	04/02/2024 2:37:01 PM	Canada (Central)	487569979...	fs-0396e837
nm-zfs-multi-az	—	1	04/02/2024 3:11:49 PM	Canada (Central)	487569979...	fs-07dc8494
nm-zfs-multi-az	—	1	03/24/2024 10:02:58 AM	Canada (Central)	487569979...	fs-085afa0d
nm-zfs-frankfurt	—	1	03/18/2024 11:59:53 AM	Europe (Frankfurt)	407355668...	fs-0c9c72a1
nm-zfs-frankfurt	—	1	03/21/2024 3:26:29 PM	Europe (Frankfurt)	487569979...	fs-0cfc6caa11c2...
nm-zfs	—	1	03/20/2024 12:56:05 PM	Canada (Central)	407355668...	fs-0f4555a2dec3...
nm-windows-sin...	—	1	04/05/2024 5:44:46 PM	Canada (Central)	487569979...	fs-0244634668b...
nm-windows-sin...	—	2	04/01/2024 9:04:55 AM	Canada (Central)	487569979...	fs-0e5305d27adf...

Removing FSx Backups

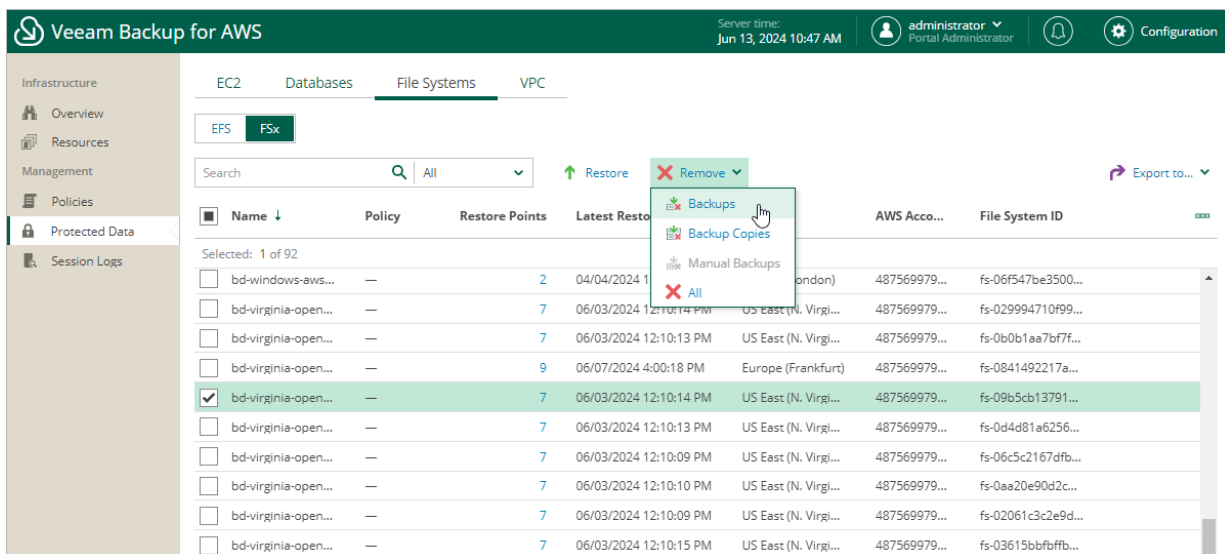
Veeam Backup for AWS applies the [configured retention policy settings](#) to automatically remove FSx file system backups and backup copies created by backup policies. If necessary, you can also remove the backed-up data manually.

To remove backed-up data manually, do the following:

1. Navigate to **Protected Data > File Systems > FSx**.
2. Select FSx file systems whose data you want to remove.
3. Click **Remove** and select either of the following options:
 - **Backups** – to remove FSx backups created for the selected file systems by backup policies.
 - **Backup Copies** – to remove backup copies created for the selected file systems by backup policies.
 - **Manual Backups** – to remove FSx backups created for the selected file systems manually.

If you want to remove only specific manual backup, follow the instructions provided in section [Removing FSx Backups Created Manually](#).

- **All** – to remove all backups and backup copies created for the selected file systems both by backup policies and manually.

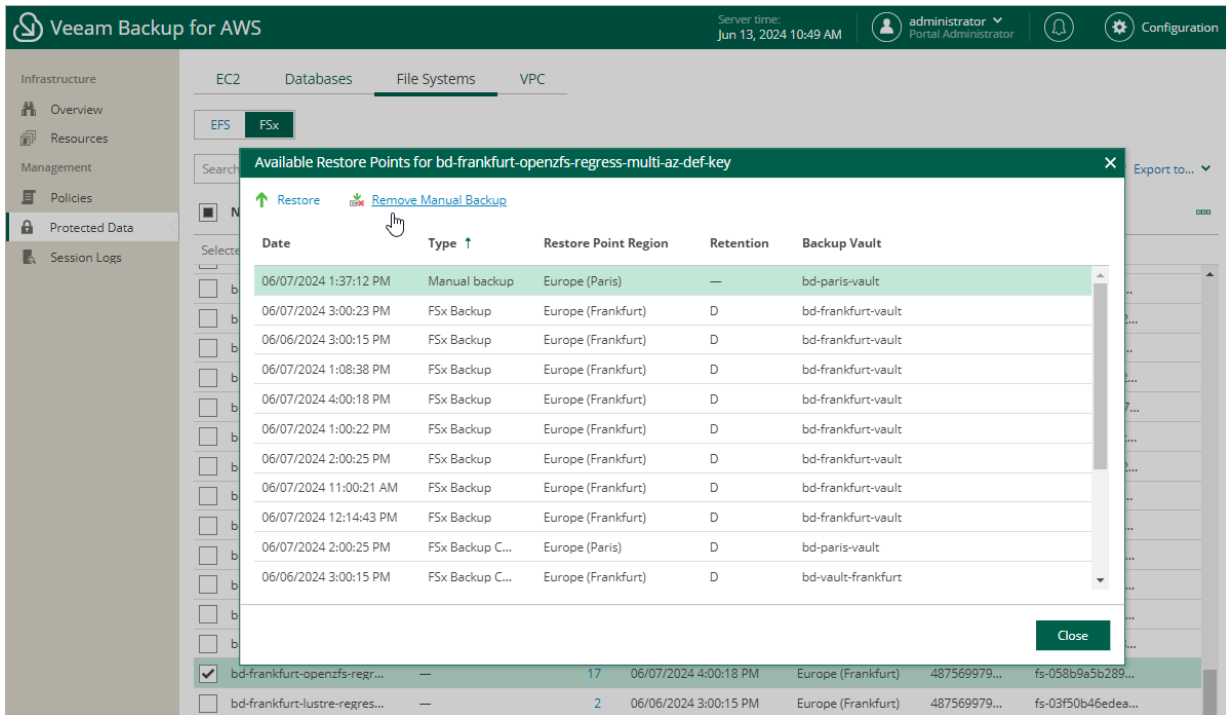


Removing FSx Backups Created Manually

To remove all backups created for an FSx file system manually, follow the instructions provided in the [Removing FSx Backups](#) section. If you want to remove a specific FSx backup created manually, do the following:

1. Navigate to **Protected Data > File Systems > FSx**.
2. Select the necessary file system, and click the link in the **Restore Points** column.

- In the **Available Restore Points** window, select a backup that you want to remove, and click **Remove Manual Backup**.



VPC Configuration Data

After the VPC Configuration Backup policy successfully creates a restore point for the VPC configuration of an AWS Region within an AWS account, the configuration record is automatically added to the resource list on the **Protected Data** page.

For each protected AWS Region within the AWS account, Veeam Backup for AWS creates a configuration record in the database with the following set of properties:

- **AWS Account** – the name of the AWS account whose IAM role was used to collect VPC configuration data.
- **Region** – the AWS Region whose VPC configuration data is backed up.
- **Latest Backup** – the date and time of the latest created restore point.
- **Latest Changes** – the summary of changes in the VPC configuration in comparison with the previous restore point.
- **Restore Points** – the total number of restore points created for the VPC configuration.

On the **Protected Data** page, you can perform the following actions:

- Compare the attributes of the current VPC configuration with the attributes stored in a backup. For more information, see [Comparing VPC Configuration Backups](#).
- Export the backed-up VPC configuration data to an AWS CloudFormation template. For more information, see [Exporting VPC Configuration](#).
- Remove restore points if you no longer need them. For more information, see [Removing VPC Configuration Backups](#).
- Restore data of backed-up VPC configurations. For more information, see [VPC Configuration Restore Using Web UI](#).

The screenshot displays the Veeam Backup for AWS web interface. The top navigation bar includes the Veeam logo, the product name 'Veeam Backup for AWS', the server time 'Jun 13, 2024 10:54 AM', and user information 'administrator Portal Administrator'. The main content area is divided into a left sidebar with navigation options (Infrastructure, Overview, Resources, Management, Policies, Protected Data, Session Logs) and a main panel. The main panel is currently showing the 'VPC' configuration page. It features a search bar for 'Account' and action buttons for 'Restore', 'Export', 'Compare', 'Remove', and 'Export to...'. Below this is a table with columns for 'AWS Account', 'Region', 'Latest Backup', 'Latest Changes', and 'Restore Points'. Two rows of data are visible, both for 'Europe (London)'. The 'Latest Backup' and 'Latest Changes' columns contain truncated text. Below the table is a 'Configuration Details' section with a search bar and a table listing individual VPC configurations. This table has columns for 'Name', 'ID', 'Type', 'Modification Date', and 'State'. The 'State' column shows various icons representing the configuration's status, such as 'Created', 'Modified', and 'Deleted'.

AWS Account	Region	Latest Backup	Latest Changes	Restore Points
487569979969...	Europe (London)	06/05/2024 1:08:47 PM	1 Vpc, 5 Subnet, 1 Endpoint, 4 ...	1
407355668422...	Europe (London)	06/05/2024 1:08:47 PM	3 Vpc, 5 Subnet, 1 Endpoint, 4 ...	1

Name	ID	Type	Modification Date	State
—	acl-0c50cdf985307297c	Acl	06/05/2024 1:08:47 PM	Created
—	acl-0c520464	Acl	06/05/2024 1:08:47 PM	Modified
—	eipalloc-099e33a52a0b9eb7f	ElasticIp	06/05/2024 1:08:47 PM	Deleted
nm-vb	eipalloc-0bac3dcf98d4f600c	ElasticIp	06/05/2024 1:08:47 PM	Modified
for VBR	eipalloc-0e7d2468e0be2d227	ElasticIp	06/05/2024 1:08:47 PM	Deleted
nm-monitoring	eipalloc-0f1aeb54a02ba2d74	ElasticIp	06/05/2024 1:08:47 PM	Created
—	eipalloc-0fcfa9241b1fb2426	ElasticIp	06/05/2024 1:08:47 PM	Deleted

Comparing VPC Configuration Backups

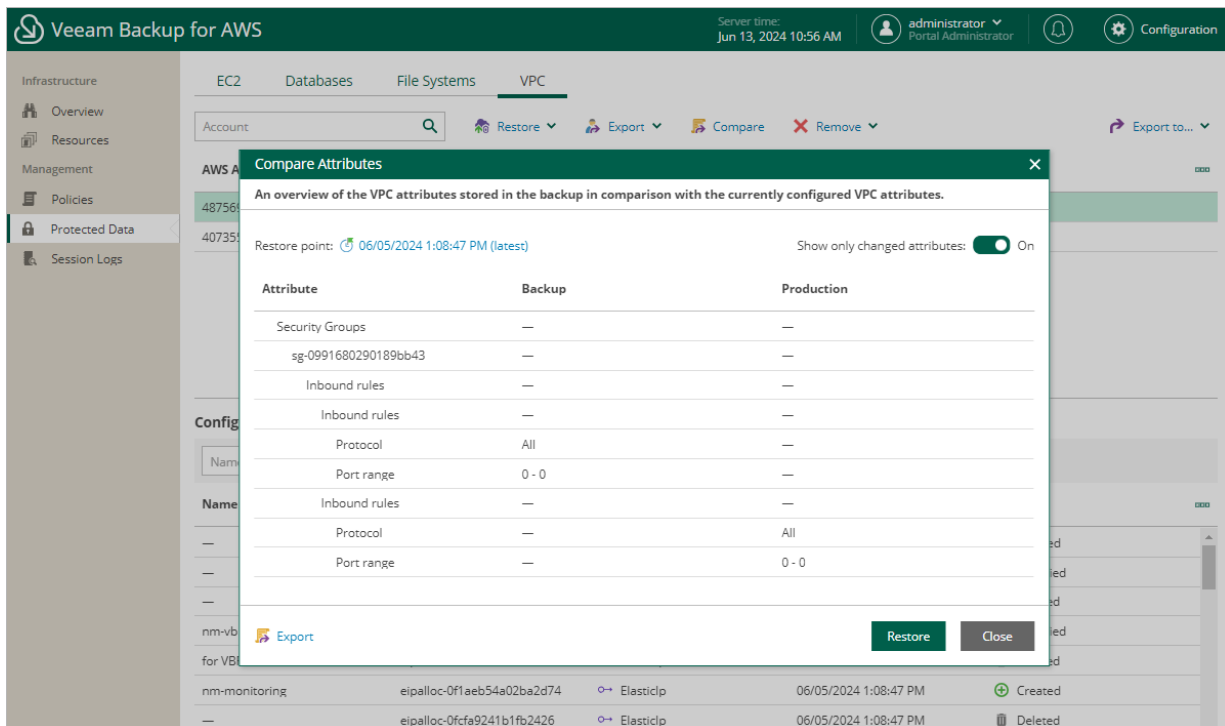
You can compare the attributes of the current Amazon VPC configuration to the attributes of a backed-up Amazon VPC configuration. To do that:

1. Navigate to **Protected Data > VPC**.
2. Select the necessary configuration record.
3. Click **Compare**.

By default, Veeam Backup for AWS uses the most recent valid restore point. However, you can compare the VPC configuration data to an earlier state. To do that, click the **Restore point** link in the **Compare Attributes** window.

TIP

You can use the selected restore point to restore or export the VPC configuration. To do that, click either **Restore** or **Export**, and follow the instructions provided in section [Performing Entire Configuration Restore](#) or [Performing Entire Configuration Export](#).



Exporting VPC Configuration

You can export backed-up VPC configuration data to an AWS CloudFormation template in the JSON format using one of the following options:

- [Perform the entire VPC configuration export.](#)
- [Perform the selected VPC configuration items export.](#)

Performing Entire Configuration Export

You can export the entire VPC configuration and restore it from the CloudFormation template to the original location or to a new location.

IMPORTANT

If you plan to restore the exported VPC configuration, consider that restore to a new location is not supported for the following VPC configuration items:

- Client VPN endpoints.
- Customer gateways and load balancer listeners that use authentication certificates.
- In route tables, for core networks and routes to AWS Outpost local gateways, network interfaces, instances and carrier gateways.

To export the entire VPC configuration to a CloudFormation template, do the following:

1. [Launch the VPC Export wizard.](#)
2. [Select a restore point and VPCs to export.](#)
3. [Specify an IAM identity for export.](#)
4. [Choose an export mode.](#)
5. [Configure mapping for Availability Zones.](#)
6. [Review settings for VPC peering connections.](#)
7. [Specify an Amazon S3 bucket where the Cloud Formation template must be placed.](#)
8. [Specify a reason for export.](#)
9. [Review export settings.](#)

Step 1. Launch VPC Export Wizard

To launch the **VPC Export** wizard, do the following:

1. Navigate to **Protected Data > VPC**.
2. Select the configuration record for an AWS Region whose VPC configuration you want to restore.
3. Click **Export > Entire VPC**.

The screenshot shows the Veeam Backup for AWS interface. The top navigation bar includes the Veeam logo, server time (Jun 18, 2024 10:34 AM), and user information (administrator, Portal Administrator). The left sidebar shows the navigation menu with 'Protected Data' selected. The main content area is titled 'VPC' and contains a table of VPC configurations. The 'Export' dropdown menu is open, showing options for 'Entire VPC' and 'Selected Items'. The 'Entire VPC' option is highlighted by the mouse cursor.

Account	Region	Modification Date	Test Changes	Restore Points
407355668422 (...)	Asia Pacific (Sydney)	06/18/2024 10:00:04 AM	No changes detected	13247
407355668422 (...)	Israel (Tel Aviv)	06/18/2024 10:00:04 AM	No changes detected	13247
407355668422 (...)	Europe (Paris)	06/18/2024 10:00:04 AM	No changes detected	13246
407355668422 (...)	Asia Pacific (Jakarta)	06/18/2024 10:00:04 AM	No changes detected	13247
407355668422 (...)	Asia Pacific (Melbourne)	06/18/2024 10:00:04 AM	No changes detected	13247
407355668422 (...)	US East (N. Virginia)	06/18/2024 10:00:04 AM	No changes detected	13247
407355668422 (...)	Europe (Frankfurt)	06/18/2024 10:00:04 AM	No changes detected	13246
407355668422 (...)	Asia Pacific (Osaka)	06/18/2024 10:00:04 AM	No changes detected	13247

Name	ID ↑	Type	Modification Date	State
khromeev-ACL-privateVPC	acl-03632553dc5cdacba	Acl	05/07/2024 4:58:58 PM	Modified
—	acl-052aafd2ab43748af	Acl	02/19/2024 11:00:27 AM	Created
—	acl-0bd6b9f5e1c1d5df9	Acl	12/21/2023 4:00:49 PM	Deleted
—	acl-0c30d0dfba51f81c8	Acl	02/19/2024 11:00:27 AM	Created
—	acl-803052fd	Acl	02/19/2024 11:00:27 AM	Created

Step 2. Select Restore Point

At the **Export List** step of the wizard, select the VPC whose configuration you want to export and a restore point that will be used to export the selected VPC configuration. By default, Veeam Backup for AWS uses the most recent valid restore point. However, you can export the VPC configuration data to an earlier state.

To select a restore point, do the following:

1. In the **Choose restore point** section, click the link to the right of **Restore point**.
2. In the **Available restore points** window, select the necessary restore point and click **Apply**.
3. In the **Choose VPCs to export** section, select VPCs whose configuration you want to export.

The screenshot shows the Veeam Backup for AWS interface during the VPC Export process. The main window is titled 'VPC Export' and has a sidebar with navigation options: Export List, Account, Export Mode, Target, Reason, and Summary. The 'Export List' section is active, showing 'Choose restore point' with a restore point of '06/18/2024 10:05:05 AM' and 'Choose VPCs to export' with a search box and a table of VPCs. A dialog box titled 'Available restore points' is open, displaying a table of restore points with columns for 'Date' and 'Changed Objects'. The second row, '06/18/2024 10:30:09 AM', is selected. The dialog also shows 'Page 1 of 67' and 'Apply' and 'Cancel' buttons.

Date	Changed Objects
06/18/2024 10:35:11 AM	No changes detected
06/18/2024 10:30:09 AM	No changes detected
06/18/2024 10:25:08 AM	No changes detected
06/18/2024 10:20:08 AM	No changes detected
06/18/2024 10:15:07 AM	No changes detected
06/18/2024 10:10:06 AM	No changes detected
06/18/2024 10:05:05 AM	No changes detected
06/18/2024 10:00:04 AM	No changes detected
06/18/2024 9:55:04 AM	No changes detected
06/18/2024 9:50:03 AM	No changes detected
06/18/2024 9:45:02 AM	No changes detected

Step 3. Specify IAM Identity


At the **Account** step of the wizard, specify an IAM role whose permissions Veeam Backup for AWS will use to perform the export operation. For more information on permissions required for the IAM role, see [VPC Configuration Restore IAM Permissions](#).

To specify an IAM role for export, select the necessary IAM role from the list. For an IAM role to be displayed in the list, it must be added to Veeam Backup for AWS as described in section [Adding IAM Roles](#). If you have not added the necessary IAM role to Veeam Backup for AWS beforehand, you can do it without closing the **VPC Export** wizard. To do that, click **Add** and complete the [Add IAM Role](#) wizard.

IMPORTANT

Consider the following:

- Make sure that the specified IAM role belongs to an AWS account to which you plan to restore the VPC configuration.
- It is recommended that you check whether the selected IAM role has all the permissions required to perform the operation. If some permissions of the IAM role are missing, the export operation will fail to complete successfully. To run the IAM role permission check, click **Check Permissions** and follow the instructions provided in section [Checking IAM Role Permissions](#).



Step 4. Choose Export Mode

At the **Export Mode** step of the wizard, choose whether you plan to restore the exported VPC configuration to the original or to a custom location. If you select the **Export to a new location** option, specify the target AWS Region where the VPC configuration will be restored.

IMPORTANT

Before you choose the export mode, consider the following:

- If you plan to restore the exported VPC configuration to the original location – when you restore the VPC configuration from the CloudFormation template, all exported VPC configuration items will be newly created in the source AWS Region. If there are any already existing items with the same names in the current VPC configuration, the restored items will be created with new IDs, but with the same names.
- If you plan to restore the exported VPC configuration to a custom location – the source and target AWS Regions may have different lists of the supported AWS services. In this case, when you restore the VPC configuration from the CloudFormation template, VPC endpoints created using an AWS service that is not available in the target AWS Region will not be restored.

The screenshot shows the Veeam Backup for AWS interface. The top bar includes the Veeam logo, the text "Veeam Backup for AWS", the server time "Jun 18, 2024 10:39 AM", and user information "administrator Portal Administrator". The main content area is titled "VPC Export" and contains a sidebar with navigation options: "Export List", "Account", "Export Mode" (selected), "Availability Zones", "Peering Connection", "Target", "Reason", and "Summary". The "Export Mode" section is active, showing two radio button options: "Export to the original location" (unselected) and "Export to a new location" (selected). Below the selected option is a dropdown menu showing "Europe (Paris)". At the bottom of the wizard, there are three buttons: "Previous", "Next", and "Cancel".

Step 5. Configure Availability Zone Mapping

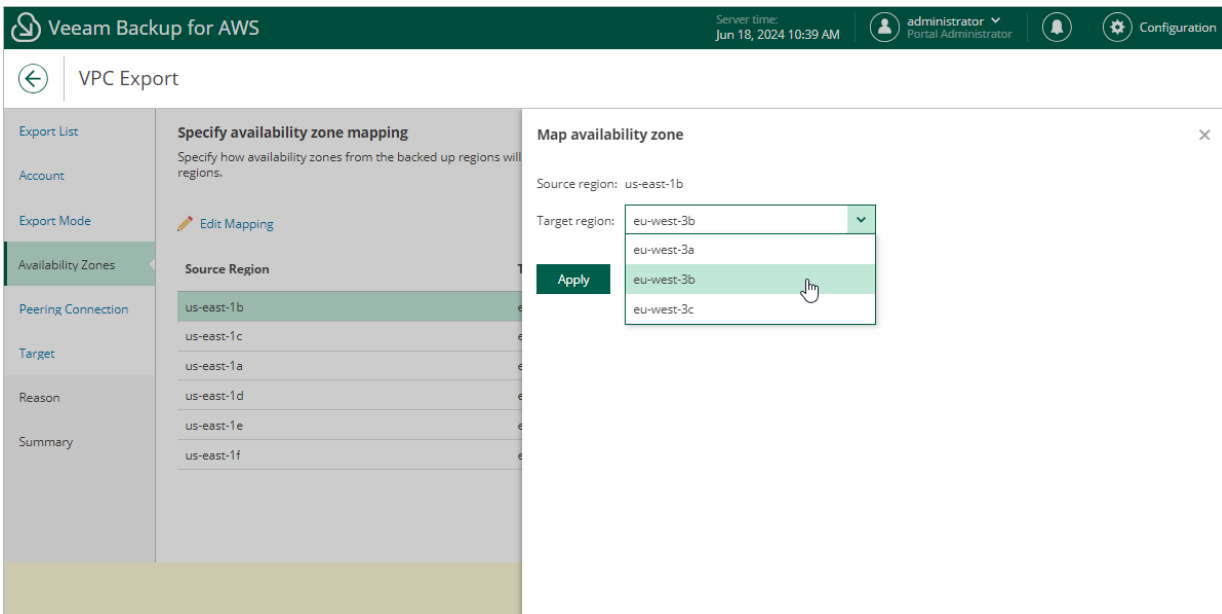
[This step applies only if you have selected the **Export to a new location** option at the **Export Mode** step of the wizard]

At the **Availability Zones** step of the wizard, for each source Availability Zone, choose an Availability Zone in the target AWS Region where VPC configuration items of the source Availability Zone will be restored:

1. Choose an Availability Zone from the list and click **Edit Mapping**.
2. In the **Map availability zone** window, select the target Availability Zone from the **Target region** drop-down list.
3. Click **Apply**.

IMPORTANT

The source and target AWS Regions may have different number of Availability Zones. In this case, Veeam Backup for AWS will automatically change subnet configuration for transit gateway VPC attachments, VPC endpoints and load balancers. After restoring, you can modify the subnet configuration manually in the AWS Management Console. To learn how to modify subnet configuration for VPC networking components, see [AWS Documentation](#).



Step 6. Review Peering Connection Settings

[This step applies only if you have selected the **Export to a new location** option at the **Export Mode** step of the wizard]

At the **Peering Connection** step of the wizard, review VPC peering connection settings. You cannot modify the VPC peering connection settings for the exported VPC. By default, Veeam Backup for AWS will export VPC peering connections as follows:

- If you export both VPCs between which you have created a peering connection, Veeam Backup for AWS will create a peering connection between the exported VPCs in the target AWS Region.
- If you export a VPC that has a peering connection to a VPC in the same AWS Region, Veeam Backup for AWS will create an inter-region peering connection between the exported VPC in the target AWS Region and the VPC with which the source VPC is peered in the source AWS Region.
- If you export a VPC that has a peering connection to a VPC in another AWS Region, Veeam Backup for AWS will create an inter-region peering connection between the exported VPC in the target AWS Region and the VPC with which the source VPC is peered in the other AWS Region.

NOTE

VPC peering connections will have the *Pending Acceptance* status after restoring from the exported CloudFormation template. To accept the restored VPC peering connections, use the AWS Management Console. For more information, see [AWS Documentation](#).

The screenshot displays the 'Review peering connection settings' step in the Veeam Backup for AWS wizard. The interface includes a navigation pane on the left with options like 'Export List', 'Account', 'Export Mode', 'Availability Zones', 'Peering Connection', 'Target', 'Reason', and 'Summary'. The main content area shows a table of peering connections with the following data:

Name	ID	Requested VPC	Accepted VPC
khrom-peering-Oregon	pcx-0a668833a27d6e0b32	vpc-0fc3c68b2a07ab3d5	vpc-02cdc004a3c9
khrom-worker-virginia-zu...	pcx-0d524f42bdb37a241	vpc-0fc3c68b2a07ab3d5	vpc-0c44a6d3b63e
khrom-VPN-private0fc3c6...	pcx-047403f99b88515ba	vpc-0fc3c68b2a07ab3d5	vpc-0cefb9ea62dd
khrom-VPN-Virg-S3-Endp...	pcx-02e4a9044f23bcf47	vpc-0cefb9ea62dd324c4	vpc-05504c381de1

At the bottom of the wizard, there are three buttons: 'Previous', 'Next', and 'Cancel'.

Step 7. Specify Amazon S3 Bucket

At the **Target** step of the wizard, specify an Amazon S3 bucket where Veeam Backup for AWS will save the CloudFormation template with the exported VPC configuration data.

Choose whether you want to save the template in the root folder of the selected Amazon S3 bucket or to create a new folder for the template.

NOTE

If you enable the [private network deployment](#) functionality, Veeam Backup for AWS will still use the public `s3.<region>.amazonaws.com` endpoint to export VPC configuration.

The screenshot shows the 'VPC Export' wizard in the Veeam Backup for AWS console. The interface is divided into a left sidebar and a main content area. The sidebar contains navigation links: 'Export List', 'Account', 'Export Mode', 'Availability Zones', 'Peering Connection', 'Target' (highlighted), 'Reason', and 'Summary'. The main content area is titled 'Specify target location' and includes the following elements:

- A header bar with the Veeam logo, 'Veeam Backup for AWS', server time 'Jun 18, 2024 10:43 AM', and user information 'administrator Portal Administrator'.
- A back arrow and the title 'VPC Export'.
- A left sidebar with navigation options.
- A main section titled 'Specify target location' with the instruction: 'To perform the export operation, specify an S3 bucket where the created CloudFormation template will be stored.'
- A 'Bucket:' field with a dropdown menu showing 'ak-virg-v7'.
- Two radio button options: 'Use root folder' (unselected) and 'Create new folder:' (selected).
- A text input field for the folder name, containing 'sb_dept01'.
- An information icon and a note: 'For more information on how to import CloudFormation templates using S3, see AWS Documentation.'
- A bottom navigation bar with 'Previous', 'Next', and 'Cancel' buttons.

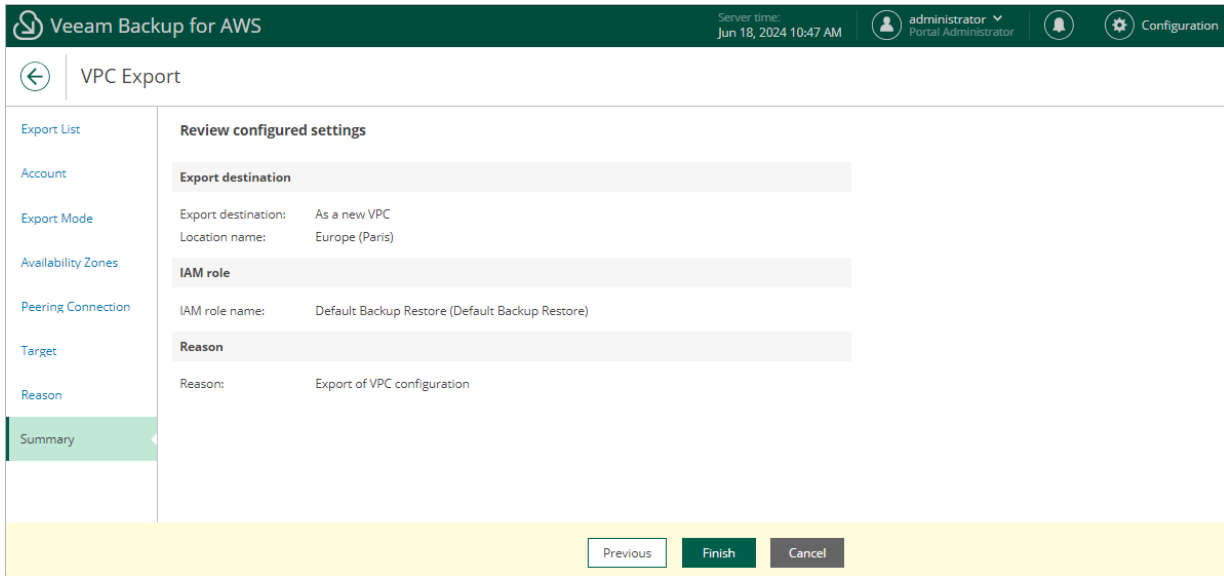
Step 8. Specify Export Reason

At the **Reason** step of the wizard, specify a reason for the export of the VPC configuration. The information you provide will be saved in the session history and you can reference it later.

The screenshot shows the 'VPC Export' wizard in Veeam Backup for AWS. The interface includes a top navigation bar with the Veeam logo, 'Veeam Backup for AWS', server time 'Jun 18, 2024 10:46 AM', and user information 'administrator Portal Administrator'. A left sidebar lists the wizard steps: Export List, Account, Export Mode, Availability Zones, Peering Connection, Target, Reason (highlighted), and Summary. The main area is titled 'Export reason' and contains the instruction 'Specify a reason for performing the export operation.' Below this is a text input field with the value 'Export of VPC configuration'. At the bottom, there are three buttons: 'Previous', 'Next', and 'Cancel'.

Step 9. Finish Working with Wizard

At the **Summary** step of the wizard, review summary information and click **Finish**.



The screenshot shows the 'VPC Export' wizard in the 'Summary' step. The interface includes a top navigation bar with the Veeam logo, 'Veeam Backup for AWS', server time 'Jun 18, 2024 10:47 AM', and user 'administrator Portal Administrator'. A left sidebar lists steps: Export List, Account, Export Mode, Availability Zones, Peering Connection, Target, Reason, and Summary (highlighted). The main area is titled 'Review configured settings' and contains three sections: 'Export destination' (As a new VPC, Location name: Europe (Paris)), 'IAM role' (IAM role name: Default: Backup Restore (Default: Backup Restore)), and 'Reason' (Reason: Export of VPC configuration). At the bottom, there are 'Previous', 'Finish', and 'Cancel' buttons.

Performing Selected Items Export

NOTE

If you export only specific VPC configuration items, you will not be able to choose a location. By default, Veeam Backup for AWS will create a CloudFormation template to restore to the original location.

When you restore the exported items from the CloudFormation template, all exported VPC configuration items will be newly created in the source AWS Region. If there are any already existing items with the same names in the current VPC configuration, the restored items will be created with new IDs, but with the same names.

To export specific VPC configuration items to a CloudFormation template, do the following:

1. [Launch the VPC Export wizard.](#)
2. [Select a restore point and VPCs to export.](#)
3. [Specify an IAM identity for export.](#)
4. [Specify an Amazon S3 bucket where the Cloud Formation template must be placed.](#)
5. [Specify a reason for the export.](#)
6. [Finish working with the wizard.](#)

Step 1. Launch VPC Export Wizard

To launch the **VPC Export** wizard, do the following.

1. Navigate to **Protected Data > VPC**.
2. Select the configuration record for an AWS Region whose VPC configuration you want to restore.
3. Click **Export > Selected items**.

The screenshot shows the Veeam Backup for AWS interface. The top navigation bar includes the Veeam logo, server time (Jun 18, 2024 10:48 AM), and user information (administrator, Portal Administrator). The left sidebar shows the navigation menu with 'Protected Data' selected. The main content area is titled 'VPC' and contains a table of VPC configurations. The 'Export' menu is open, showing 'Entire VPC' and 'Selected Items' options. The 'Selected Items' option is highlighted. Below the table is a 'Configuration Details' section with a search bar and a table of configuration details.

AWS Account	Region	Backup	Latest Changes	Restore Points
407355668422 (v...	Middle East (UAE)	07/10/2024 1:30:11 PM	No changes detected	17148
407355668422 (v...	Canada West (Calgary)	07/10/2024 1:30:11 PM	No changes detected	16903
407355668422 (v...	Europe (London)	07/10/2024 1:30:11 PM	No changes detected	17148
407355668422 (v...	US East (Ohio)	07/10/2024 1:30:11 PM	No changes detected	17147
407355668422 (v...	Asia Pacific (Sydney)	07/10/2024 1:30:11 PM	No changes detected	17148
407355668422 (v...	Israel (Tel Aviv)	07/10/2024 1:30:11 PM	No changes detected	17148
407355668422 (v...	Europe (Paris)	07/10/2024 1:30:11 PM	No changes detected	17147
407355668422 (v...	Asia Pacific (Jakarta)	07/10/2024 1:30:11 PM	No changes detected	17148
407355668422 (v...	Asia Pacific (Melbourne)	07/10/2024 1:30:11 PM	No changes detected	17148
407355668422 (v...	US East (N. Virginia)	07/10/2024 1:30:11 PM	No changes detected	17148
407355668422 (v...	Europe (Frankfurt)	07/10/2024 1:30:11 PM	No changes detected	17147
407355668422 (v...	Asia Pacific (Osaka)	07/10/2024 1:30:11 PM	No changes detected	17148

Configuration Details

Name	ID	Type	Modification Date	State
------	----	------	-------------------	-------

Page 1 of 9

Step 2. Select Restore Point

At the **Export List** step of the wizard, select the VPC configuration items you want to export and a restore point that will be used to export the selected VPC configuration items. By default, Veeam Backup for AWS uses the most recent valid restore point. However, you can export the VPC configuration data to an earlier state.

1. To select the restore point:
 - a. In the **Choose restore point** section, click the link to the right of **Restore point**.
 - b. In the **Available restore points** window, select the necessary restore point and click **Apply**.
2. To select the VPC configuration items:
 - a. In the **Create export list** section, select the type of VPC configuration item you want to export and click **Edit**.
 - b. In the **Edit export list** window, click **Add to Export List**.
 - c. In the **Item List** window, select check boxes next to the items that you want to export, and click **Add**.
 - d. In the **Edit export list** window, review the restore list and click **Apply**.

IMPORTANT

When performing the export operation, Veeam Backup for AWS does not validate the export list. If any of the VPC configuration items on which the selected items depend are missing from the current VPC configuration, the restore of the selected VPC configuration items from the created CloudFormation template will fail.

The screenshot shows the Veeam Backup for AWS interface. The top navigation bar includes the Veeam logo, 'Veeam Backup for AWS', server time 'Jun 18, 2024 10:51 AM', and user 'administrator Portal Administrator'. The main window is titled 'VPC Export' and is divided into three panes. The left pane shows a sidebar with 'Export List' selected. The middle pane, 'Choose restore point', shows a restore point of '06/18/2024 10:05:05 AM'. The right pane, 'Edit export list (Security Group)', shows a search bar for 'Name or ID' and 'State'. Below the search bar is a table with columns 'Name', 'ID', 'Type', and 'State'. One item is selected: 'khrom-vpn-endpoints-SG' with ID 'sg-00dd0292e4d41bd7d' and Type 'Security Group'. The 'State' column shows a 'Modified' icon. At the bottom of the right pane are 'Apply' and 'Cancel' buttons.

Name	ID	Type	State
khrom-vpn-endpoints-SG	sg-00dd0292e4d41bd7d	Security Group	Modified

Step 3. Specify IAM Identity

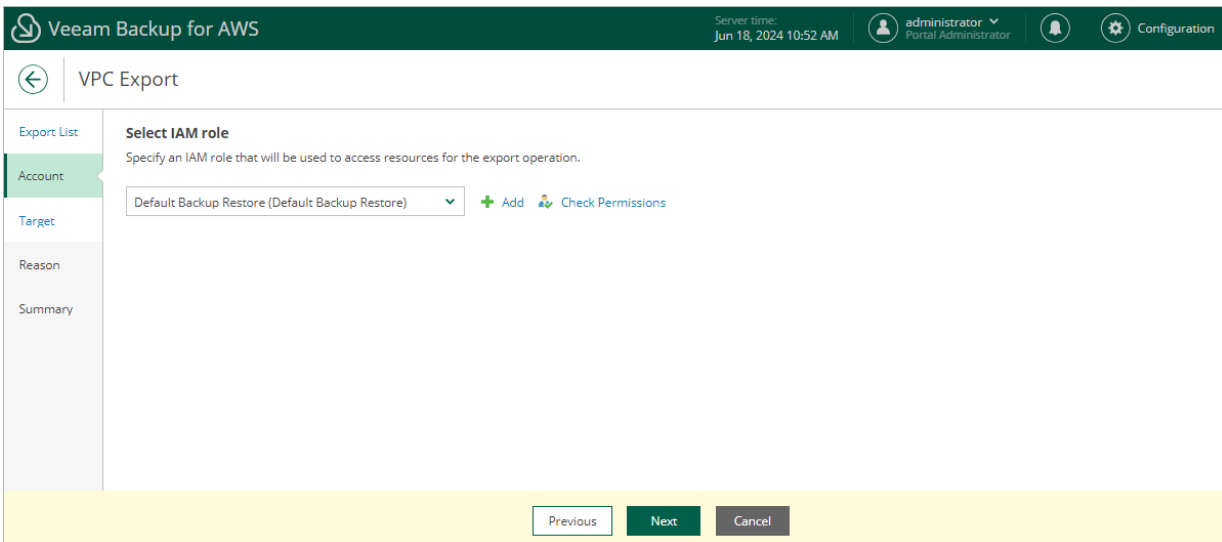
At the **Account** step of the wizard, specify an IAM role whose permissions Veeam Backup for AWS will use to perform the export operation. For more information on permissions required for the IAM role, see [VPC Configuration Restore IAM Permissions](#).

To specify an IAM role for export, select the necessary IAM role from the list. For an IAM role to be displayed in the IAM Role list, it must be added to Veeam Backup for AWS as described in section [Adding IAM Roles](#). If you have not added the necessary IAM role to Veeam Backup for AWS beforehand, you can do it without closing the VPC Restore wizard. To do that, click **Add** and complete the [Add IAM Role](#) wizard.

IMPORTANT

Consider the following:

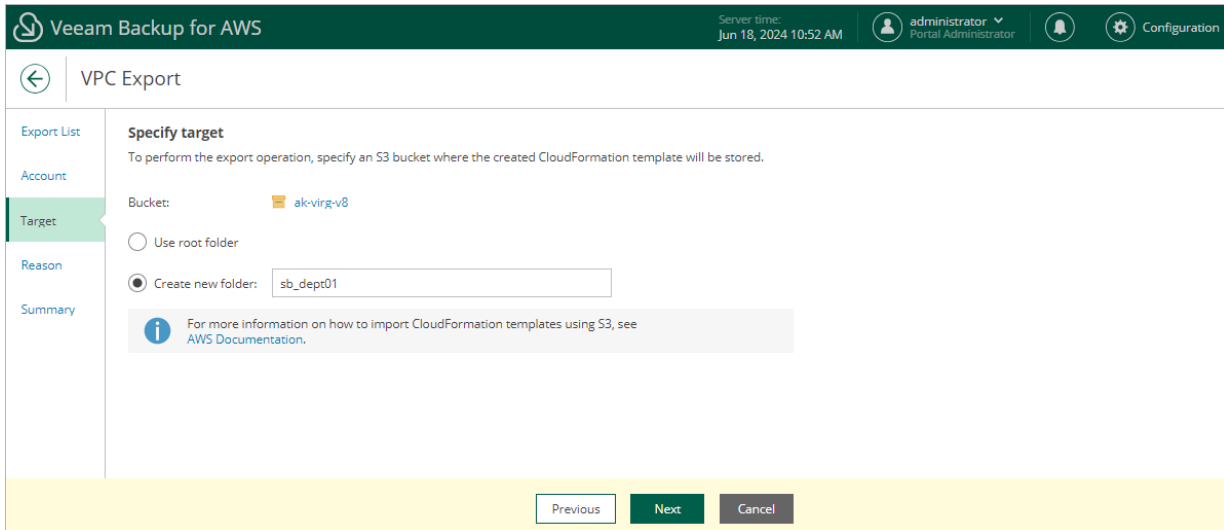
- Make sure that the specified IAM role belongs to an AWS account to which you plan to restore the VPC configuration items.
- It is recommended that you check whether the selected IAM role has all the permissions required to perform the operation. If some permissions of the IAM role are missing, the export operation will fail to complete successfully. To run the IAM role permission check, click **Check Permissions** and follow the instructions provided in section [Checking IAM Role Permissions](#).



Step 4. Specify Amazon S3 Bucket

At the **Target** step of the wizard, specify an Amazon S3 bucket where Veeam Backup for AWS will save the CloudFormation template with the exported VPC configuration items.

Choose whether you want to save the template in the root folder of the selected Amazon S3 bucket or to create a new folder for the template.



The screenshot shows the 'VPC Export' wizard in the Veeam Backup for AWS interface. The 'Target' step is active, where the user specifies the S3 bucket and folder for the CloudFormation template. The bucket is 'ak-virg-v8'. The user has selected 'Create new folder' with the name 'sb_dept01'. The interface includes a navigation pane on the left with options: Export List, Account, Target (selected), Reason, and Summary. The top bar shows the Veeam logo, server time (Jun 18, 2024 10:52 AM), and user information (administrator Portal Administrator). The bottom bar has 'Previous', 'Next', and 'Cancel' buttons.

Veeam Backup for AWS

Server time: Jun 18, 2024 10:52 AM

administrator Portal Administrator

Configuration

VPC Export

Export List

Account

Target

Reason

Summary

Specify target

To perform the export operation, specify an S3 bucket where the created CloudFormation template will be stored.

Bucket: ak-virg-v8

Use root folder

Create new folder: sb_dept01

For more information on how to import CloudFormation templates using S3, see AWS Documentation.

Previous Next Cancel

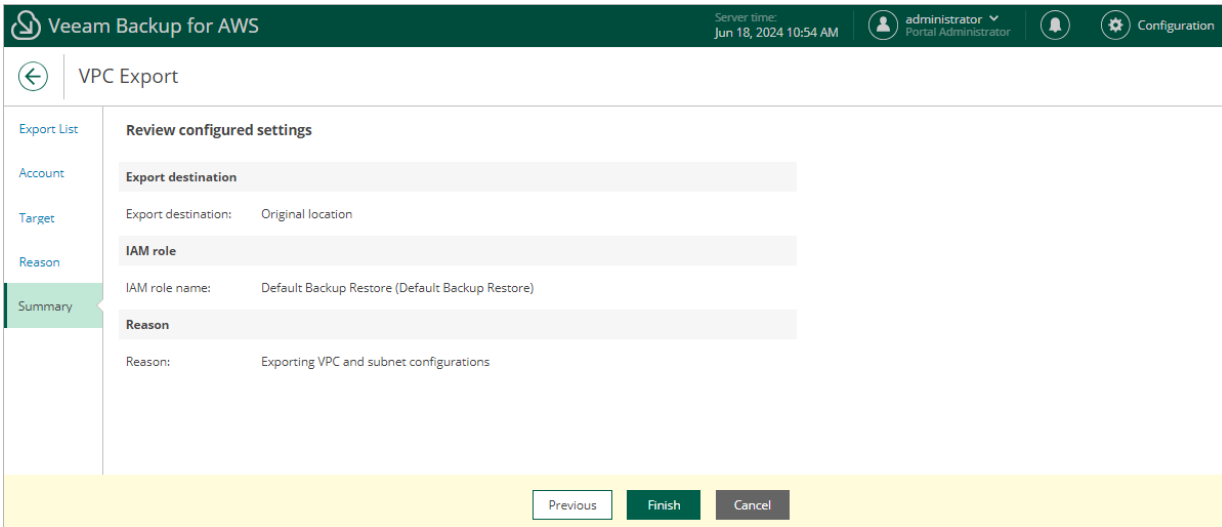
Step 5. Specify Export Reason

At the **Reason** step of the wizard, specify a reason for the export of the VPC configuration items. The information you provide will be saved in the session history and you can reference it later.

The screenshot shows the Veeam Backup for AWS interface. The top navigation bar includes the Veeam logo, the text "Veeam Backup for AWS", the server time "Jun 18, 2024 10:53 AM", the user "administrator Portal Administrator", and a "Configuration" link. The main content area is titled "VPC Export" and features a left-hand navigation pane with steps: "Export List", "Account", "Target", "Reason" (highlighted), and "Summary". The "Reason" step is active, displaying the heading "Export reason" and the instruction "Specify a reason for performing the export operation." Below this, a text input field contains the text "Exporting VPC and subnet configurations". At the bottom of the wizard, there are three buttons: "Previous", "Next" (highlighted in green), and "Cancel".

Step 6. Finish Working with Wizard

At the **Summary** step of the wizard, review summary information and click **Finish**.



Removing VPC Configuration Backups

Veeam Backup for AWS applies the [configured retention policy settings](#) to automatically remove VPC configuration backups and backup copies created by the VPC Configuration Backup policy. If necessary, you can also remove these backups manually – from the configuration database or from the repository. Keep in mind, that:

- If a backup is removed from the repository but still exists in the configuration database, you will be able to use this backup to restore the VPC configuration data.
- If a backup is removed from the configuration database but still exists in the repository, you will be able to use this backup to restore the VPC configuration data – but you will first have re-add the backup repository to Veeam Backup for AWS as described in section [Adding Backup Repositories Using Web UI](#).

To remove backed-up data manually, do the following:

1. Navigate to **Protected Data > VPC**.
2. Select the configuration record for which you want to remove the backed-up data.
Each configuration record contains a whole set of all virtual network configuration backups created for an AWS account and an AWS Region. Note that you cannot remove individual virtual network configuration items or specific backups.
3. Click **Remove** and select either of the following options:
 - **Backups** – to remove all VPC configuration backups for the selected configuration record from the Veeam Backup for AWS database.
 - **Backup Copies** – to remove all VPC configuration backups of all AWS Regions within selected AWS account from the backup repository specified in the [target settings](#) of the VPC Configuration Backup policy.

Veeam Backup for AWS Server time: Jun 13, 2024 10:55 AM administrator Portal Administrator Configuration

Infrastructure EC2 Databases File Systems VPC

Account Restore Export Compare Remove Export to...

AWS Account ↑	Region	Latest Backup	Latest Changes	Restore Points
487569979969...	Europe (London)	06/05/2024 1:08:47 PM	1 Vpc, 5 Subnet, 1 Endpoint, 4 ...	1
407355668422...	Europe (London)	06/05/2024 1:08:47 PM	3 Vpc, 5 Subnet, 1 Endpoint, 4 ...	1

Configuration Details

Name or ID Filter (None) State: + - ✎

Name	ID ↑	Type	Modification Date	State
—	ad-0c50cdf985307297c	Acl	06/05/2024 1:08:47 PM	Created
—	ad-0c520464	Acl	06/05/2024 1:08:47 PM	Modified
—	eipalloc-093e33a52a0b9eb7f	Elasticip	06/05/2024 1:08:47 PM	Deleted
nm-vb	eipalloc-0bac3dcf98d4f600c	Elasticip	06/05/2024 1:08:47 PM	Modified
for VBR	eipalloc-0e7d2468e0be2d227	Elasticip	06/05/2024 1:08:47 PM	Deleted
nm-monitoring	eipalloc-0f1aeb54a02ba2d74	Elasticip	06/05/2024 1:08:47 PM	Created
—	eipalloc-0fcfa9241b1fb2426	Elasticip	06/05/2024 1:08:47 PM	Deleted

Performing Restore

In various disaster recovery scenarios, you can perform the following restore operations using backed-up data:

- [Restore of EC2 instances](#) – restore EC2 instances from cloud-native snapshots, snapshot replicas or image-level backups to the original location or to a new location.
- [Restore of RDS resources](#) – restore DB instances and Aurora DB clusters (from cloud-native snapshots, snapshot replicas) and DB instance databases (from image-level backups) to the original location or to a new location.
- [Restore of DynamoDB tables](#) – restore DynamoDB tables from backups to the original location or to a new location.
- [Restore of Redshift clusters](#) – restore Redshift clusters from backups to their original location.
- [Restore of EFS file systems](#) – restore file systems from backups to the original location or to a new location.
- [Restore of FSx file systems](#) – restore file systems from backups to the original location or to a new location.
- [Restore of VPC configurations](#) – restore VPC configurations from VPC configuration backups to the original location or to a new location.
- [Instant Recovery](#) – immediately restore EC2 instances from image-level backups to VMware vSphere and Hyper-V environments, and to Nutanix AHV clusters.
- [EC2 instance disk export](#) – restore volume disks and convert them to disks of the VMDK, VHD or VHDX format.
- [EC2 instance disk publish](#) – publish point-in-time volume disks and copy the necessary files and folders to the target server.
- [Restore to Microsoft Azure](#) – restore EC2 instances from image-level backups to Microsoft Azure as Azure VMs.
- [Restore to Google Cloud](#) – restore EC2 instances from image-level backups to Google Cloud as VM instances.
- [Restore to Nutanix AHV](#) – restore EC2 instances from image-level backups to Nutanix AHV as Nutanix AHV VMs.

EC2 Restore

The actions that you can perform with restore points of EC2 instances depend on whether you access the restore points using the Veeam Backup & Replication console or the Veeam Backup for AWS Web UI.

EC2 Restore Using Console

Veeam Backup & Replication offers the following restore operations:

- [Instance restore](#) – restore an entire EC2 instance.
- [Guest OS file recovery](#) – restore individual files and folders of an EC2 instance.
- [Application restore](#) – restore applications such as Microsoft Active Directory, Microsoft Exchange, Microsoft SharePoint and Microsoft SQL Server.

You can restore EC2 instance data to the most recent state or to any available restore point.

IMPORTANT

You can use restore points stored in standard backup repositories to perform all the listed recovery operations, while restore points stored in archive backup repositories can only be used to perform restore of EC2 to the original or to a new location.

Performing Instance Restore

In case of a disaster, you can restore an entire EC2 instance from a cloud-native snapshot, a snapshot replica or an image-level backup. Veeam Backup & Replication allows you to restore one or more EC2 instances at a time, to the original location or to a new location.

How Instance Restore Works

To restore EC2 instances from cloud-native snapshots, Veeam Backup & Replication uses [native AWS capabilities](#). To restore EC2 instances from image-level backups, Veeam Backup & Replication uses different algorithms depending on whether a backup appliance is added to the backup infrastructure:

- If the backup appliance is connected to the backup server, Veeam Backup & Replication uses the restore algorithm described in section [Entire EC2 Restore](#).
- If the backup appliance is not connected to the backup server, Veeam Backup & Replication uses the restore algorithm described in the Veeam Backup & Replication User Guide, section [How Restore to Amazon EC2 Works](#).

NOTE

Restore to AWS Outposts is available only in the Veeam Backup for AWS Web UI. To learn how to perform restore to Outposts, see [Before You Begin](#).

How to Perform Instance Restore

To restore an EC2 instance, do the following:

1. [Launch the Restore to Amazon EC2 wizard](#).
2. [Select a restore point](#).
3. [Choose a restore mode](#).
4. [Select an AWS Region](#).

5. Specify instance type and enable encryption.
6. Specify a new name for the instance.
7. Configure network settings.
8. Specify a restore reason.
9. Finish working with the wizard.

Step 1. Launch Restore to Amazon EC2 Wizard

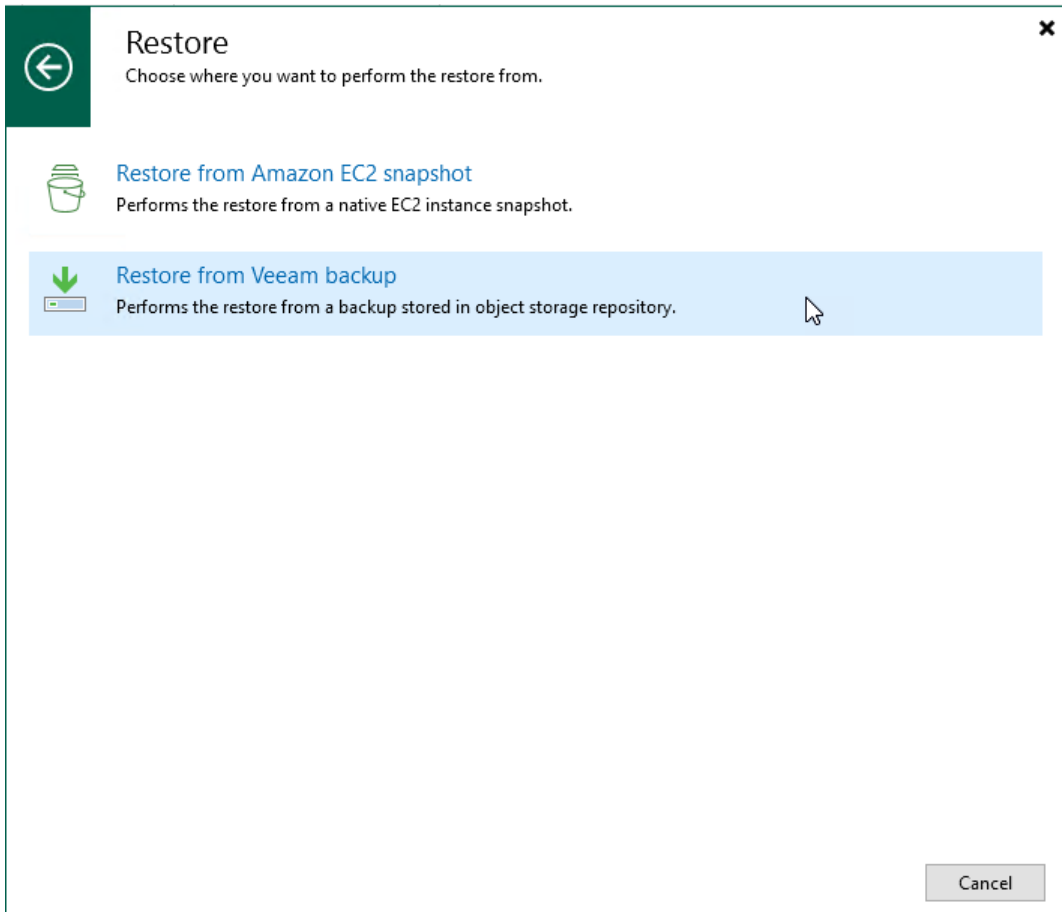
To launch the **Restore to Amazon EC2** wizard, do the following:

1. In the Veeam Backup & Replication console, open the **Home** view.
2. Navigate to **Backups > Snapshots** if you want to restore from a cloud-native snapshot, or to **Backups > External Repository** if you want to restore from an image-level backup.
3. In the working area, expand the backup policy that protects an EC2 instance that you want to restore, select the necessary instance and click **Amazon EC2** on the ribbon.

Alternatively, you can right-click the instance and select **Restore to Amazon EC2**.

TIP

You can also launch the **Restore to Amazon EC2** wizard from the **Home** tab. To do that, click **Restore** and select **AWS**. Then, in the **Restore** window, select **Amazon EC2 > Entire machine restore > Restore to public cloud > Restore to Amazon EC2** and, depending on whether you want to restore from a backup or a snapshot, click either **Restore from Amazon EC2 snapshot** or **Restore from Veeam backup**.



Step 2. Select Restore Point

At the **Instance** step of the wizard, choose a restore point that will be used to restore the selected EC2 instance. By default, Veeam Backup & Replication uses the most recent valid restore point. However, you can restore the instance data to an earlier state.

To select a restore point, do the following:

1. In the **Instance** list, select the EC2 instance and click **Point**.
2. In the **Restore Points** window, expand the backup policy that protects the EC2 instance, select the necessary restore point and click **OK**.

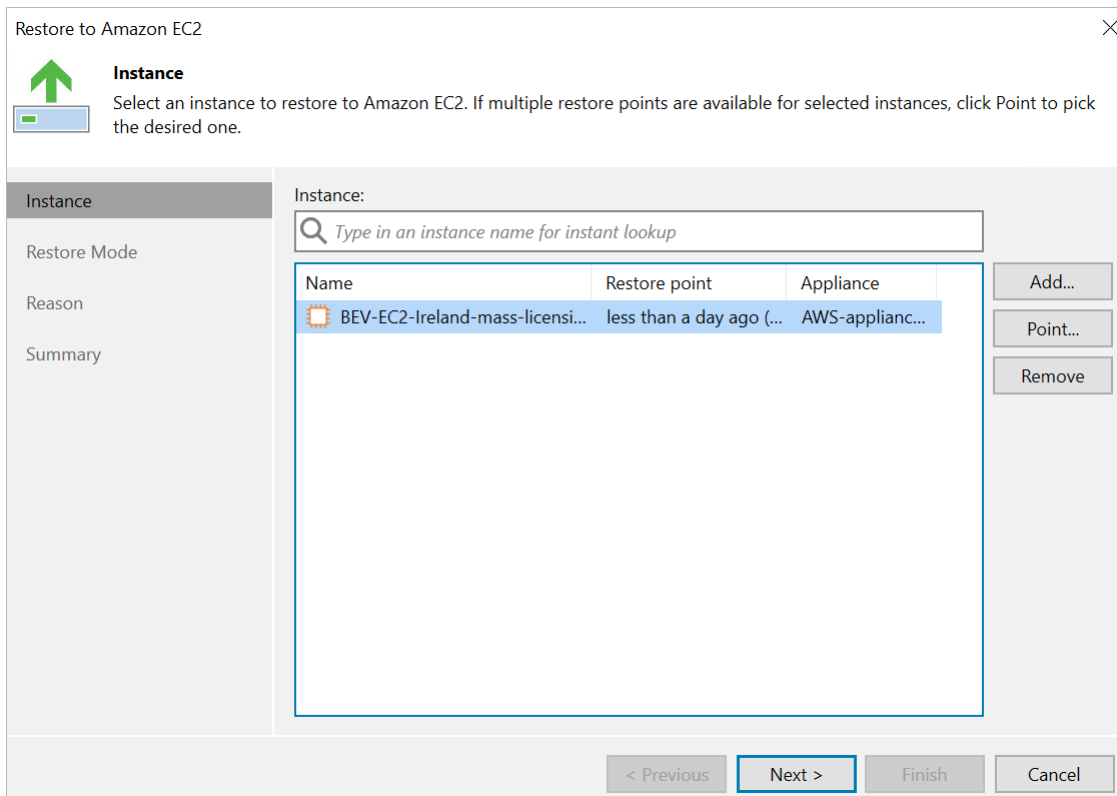
To help you choose a restore point, Veeam Backup & Replication provides the following information on each available restore point:

- **Job** – the name of the backup policy that created the restore point and the date when the restore point was created.
- **Type** – the type of the restore point.
- **Location** – the AWS Region or repository where the restore point is stored.

TIP

You can use the wizard to restore multiple instances at a time. To do that, click **Add**, select more EC2 instances to restore and select a restore point for each of them.

Note that if you want to restore an EC2 instance from a backup that is stored in a repository of the S3 Glacier Flexible Retrieval or S3 Glacier Deep Archive storage class, you must first retrieve the archived data. That is why Veeam Backup & Replication will open the **Retrieve Backup** wizard if the selected restore point is stored in an archive backup repository. To learn how to complete the wizard and retrieve the archived data, see [Retrieving Data from Archive](#).



Retrieving Data from Archive

Backups stored in archive backup repositories are not immediately accessible. If you want to restore an EC2 instance from a backup that is stored in an archive backup repository, you must first retrieve the archived data.

During the data retrieval process, a temporary copy of the archived data is created in an Amazon S3 bucket where the archive backup repository is located. This copy is stored in the S3 standard storage class for a period of time that you specify when launching the data retrieval process. If the time period expires while a restore operation is still running, Veeam Backup for AWS automatically extends the period to keep the retrieved data available for 1 more day. You can also [extend the availability period manually](#).

Retrieving Data

To retrieve data from an archived restore point, complete the **Retrieve Backup** wizard:

1. At the **Retrieval Mode** step of the wizard, choose the retrieval mode that Veeam Backup & Replication will use to retrieve the archived data:

- **Expedited** – the most expensive mode. If you choose this mode, the retrieved data will be available within 1-5 minutes.

Note that this mode is not supported for data stored in the S3 Glacier Deep Archive storage class.

- **Standard** – the recommended mode. If you choose this mode, the retrieved data will be available within 3-5 hours for data stored in the Amazon S3 Glacier Flexible Retrieval storage class and within 12 hours for data stored in the Amazon S3 Glacier Deep Archive storage class.

- **Bulk** – the least expensive mode. If you choose this mode, the retrieved data will be available within 5-12 hours for data stored in the Amazon S3 Glacier Flexible Retrieval storage class and within 48 hours for data stored in the Amazon S3 Glacier Deep Archive storage class.

- **Standard accelerated** – the option that is less expensive than the Expedited option. The retrieved data is available within 15-30 minutes for data stored in the S3 Glacier Flexible Retrieval storage class. With this option enabled, Veeam Backup for AWS leverages the S3 Batch Operations functionality to retrieve the archived data.

Before you enable this mode, it is recommended that you check whether the IAM role specified to access the archive backup repository has all the permissions required to perform data retrieval operations. For more information, see [Managing Backup Repositories](#).

For more information on archive retrieval options, see [AWS Documentation](#).

2. At the **Availability Period** step of the wizard, specify the number of days for which you want to keep the data available for restore operations.

The data will be available during the day when the retrieval process completes plus the specified number of days. Each day starts at 12:00 AM (UTC) and ends at 11:59 PM (UTC). For example, if the data retrieval finishes at 3:00 PM (UTC) on June 6, and the availability period is set to 1 day, the data will be available till 11:59 PM (UTC) on June 7.


TIP

If you want to receive an email notification when data is about to expire, select the **Enable e-mail notifications** check box and choose when you want to be notified (that is, the number of hours remaining until data expiration). To learn how to configure global email notification settings, see the Veeam Backup & Replication User Guide, section [Configuring Global Email Notification Settings](#).

3. At the **Summary** step of the wizard, review summary information and click **Finish**.

The retrieved data will be displayed in the **Home** view under the **Data Retrieval** node.

After you complete the **Retrieve Backup** wizard, you will be able to proceed with the **Restore to Amazon EC2** wizard. However, the restore process will start only after the data is retrieved.



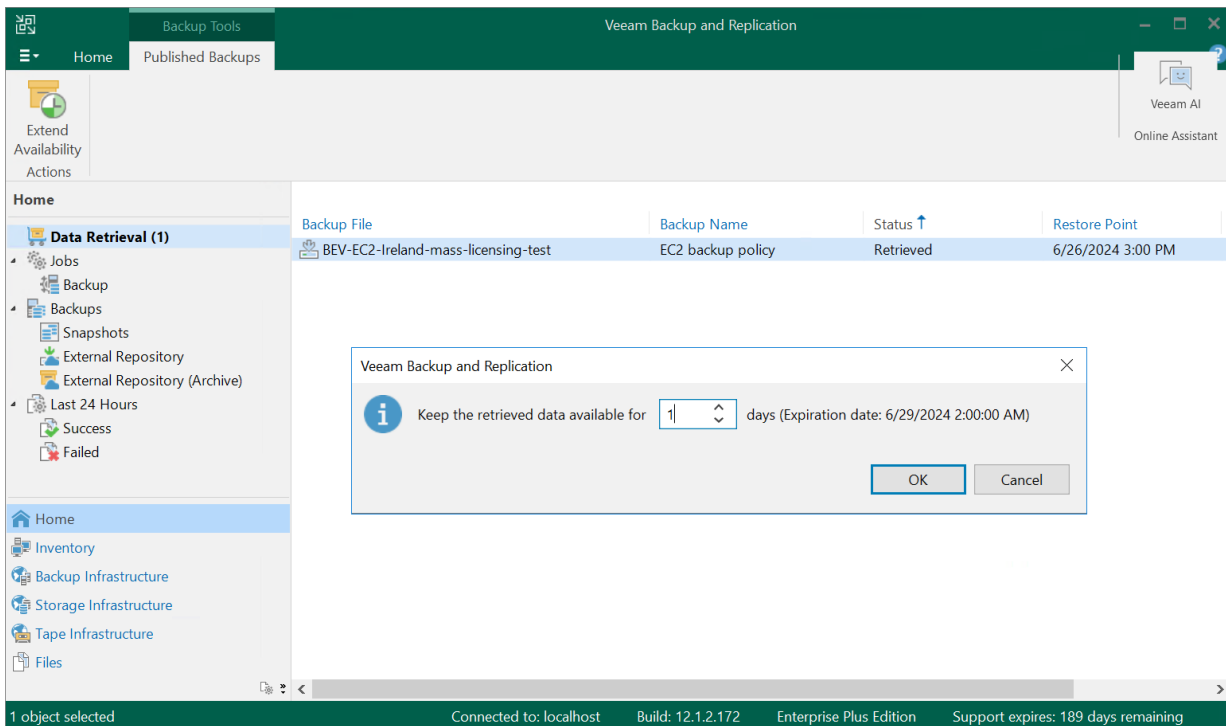
Extending Data Availability

To extend time for which you want to keep retrieved data available for restore operations:

1. In the Veeam Backup & Replication console, open the **Home** view.
2. Navigate to **Data Retrieval** node.
3. Select an EC2 instance for which you want to extend availability of the retrieved data and click **Extend Availability** on the ribbon.

Alternatively, you can right-click the EC2 instance and click **Extend availability**.

4. In the opened window, specify the number of days for which you want to keep the data available for restore operations, and click **OK**.



Step 3. Choose Restore Mode

At the **Restore Mode** step of the wizard, do the following:

1. Choose whether you want to restore the selected EC2 instance to the original or to a new location.

NOTE

If you choose to restore to the original location, consider the following:

- An IAM role that will be used to perform the restore operation must belong to an AWS account where the selected restore point was created.
- The source EC2 instance will be automatically powered off and removed from AWS after the restore process completes successfully.
- If private IP addresses that were assigned to the source EC2 instance are in use by the source or any other EC2 instance, the restored EC2 instance will be assigned new private IP addresses.

2. Click **Pick an IAM identity to use** to select an IAM identity whose permissions will be used to perform the restore operation:

- To specify an IAM role, select the **IAM role** option and choose the necessary IAM role from the **IAM role** drop-down list.

For an IAM role to be displayed in the list of available roles, it must be added to the backup appliance as described in section [Adding IAM Roles](#).

- To specify access keys of an IAM user, select the **Temporary access key** option, and use the **Access key** and **Secret key** fields to provide the access key ID and the secret access key.

For more information on the permissions that the IAM role or IAM user must have to perform the restore operation, see [EC2 Restore IAM Permissions](#).

NOTE

By default, to perform restore operations, Veeam Backup & Replication uses permissions of either the IAM role that was used to protect the source EC2 instance, or the IAM role that was used to update information on restore points created for the instance while rescanning the AWS infrastructure.

3. [Applies only if you restore EC2 instances from image-level backups and have chosen the **IAM role** option] Choose whether you want Veeam Backup & Replication to deploy worker instances processing EC2 instance data in [production accounts](#).


By default, Veeam Backup & Replication launches worker instances in the backup account. However, you may want to instruct Veeam Backup & Replication to launch worker instances in a production account, for example, if the EBS volumes of the restored EC2 instance are encrypted with the default AWS managed key. To do that, click **Configure** and select an IAM role that will be attached to the worker instances and used by Veeam Backup & Replication to communicate with these instances. The selected IAM role must belong to the same account to which the IAM role specified for the restore operation belongs, and must be assigned permissions listed in section [Worker IAM Permissions](#).

For an IAM role to be displayed in the list of available roles, it must be added to the backup appliance as described in section [Adding IAM Roles](#).

IMPORTANT

If you instruct Veeam Backup & Replication to deploy worker instances in a production account, you must assign additional permissions to the IAM role that will be used to perform the restore operation. For more information on the required permissions, see [EC2 Restore IAM Permissions](#).

Restore to Amazon EC2 ✕



Restore Mode

Specify whether selected EC2 instances should be restored back to the original location, or to a new location or with different settings.

Restore to the original location
Quickly initiate restore of the selected EC2 instance to its original location, with the original name and settings. This option minimizes the chance of user input error.


Restore to a new location, or with different settings
Customize the restored EC2 instance location, and change its settings. The wizard will automatically populate all controls with the original EC2 instance settings as the defaults.

[Pick an IAM identity to use](#)

Worker ✕

Deploy workers in production account with the following IAM role attached:

Worker-deployment-role ▼

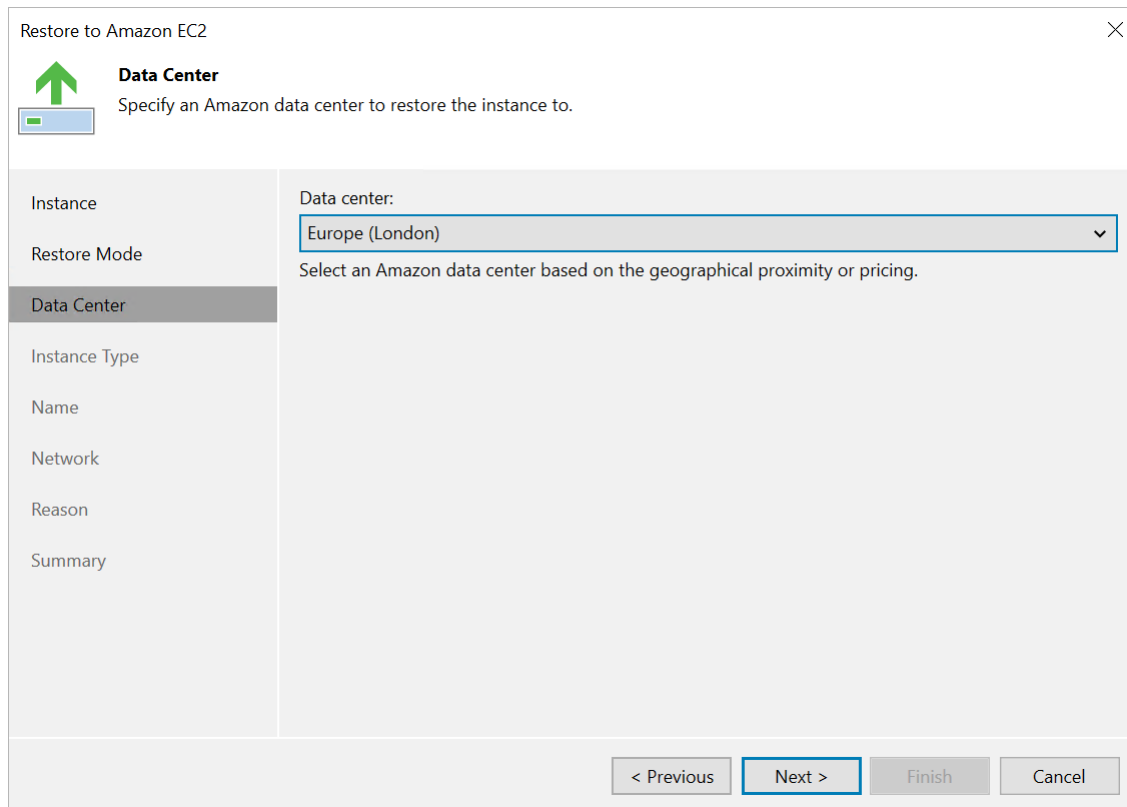
 Enable worker deployment in the production account to be able to restore instances whose volumes are encrypted using default AWS managed keys.

Step 4. Select Region

[This step applies only if you have selected the **Restore to a new location, or with different settings** option at the **Restore Mode** step of the wizard]

At the **Data Center** step of the wizard, select an AWS Region where the restored EC2 instance will operate.

If the selected location differs from the original location of the EC2 instance, Veeam Backup & Replication will raise a warning notifying that the locations do not match. Click **Yes** to acknowledge the warning. Otherwise, you will not be able to proceed with the wizard.



The screenshot shows the 'Restore to Amazon EC2' wizard window. The title bar reads 'Restore to Amazon EC2' with a close button (X) on the right. Below the title bar is a green upward-pointing arrow icon and the text 'Data Center' in bold. Underneath, it says 'Specify an Amazon data center to restore the instance to.' There is a small blue square icon with a white minus sign to the left of the text.

The main area is divided into two sections. On the left is a vertical navigation pane with the following items: 'Instance', 'Restore Mode', 'Data Center' (which is highlighted with a dark grey background), 'Instance Type', 'Name', 'Network', 'Reason', and 'Summary'. On the right, under the heading 'Data center:', there is a dropdown menu currently showing 'Europe (London)'. Below the dropdown, it says 'Select an Amazon data center based on the geographical proximity or pricing.'

At the bottom of the window, there are four buttons: '< Previous', 'Next >' (which is highlighted with a blue border), 'Finish', and 'Cancel'.

Step 5. Specify Instance Type and Enable Encryption

[This step applies only if you have selected the **Restore to a new location, or with different settings** option at the **Restore Mode** step of the wizard]

At the **Instance Type** step of the wizard, you can configure settings for the restored EC2 instance. To do that, select the instance and do the following:

- If you want to specify a new machine type for the restored EC2 instance, click **Type** and select the necessary type in the **Instance Type** window.
For the list of all existing EC2 instance types, see [AWS Documentation](#).
- If you want to change the encryption settings of the restored EC2 instance, click **Encryption** and do the following in the **Disk Encryption** window:
 - Select the **Preserve the original encryption settings** option if you do not want to encrypt the EBS volumes or want to apply the original encryption scheme of the source EC2 instance.

NOTE

You will not be able to select the **Preserve the original encryption settings** option if the AWS KMS key that was used to encrypt EBS volumes of the source instance is not available in the region to which the EC2 instance will be restored.

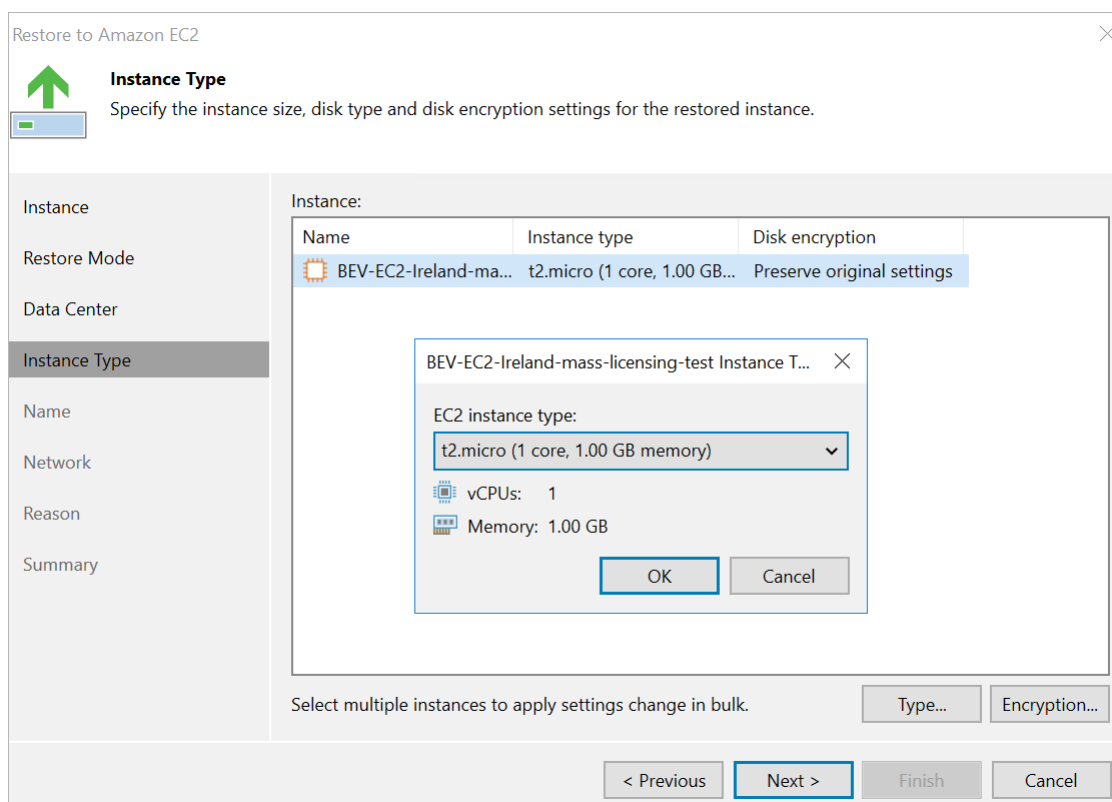
- Select the **Use the following encryption key** option if you want to encrypt the restored EBS volumes of the processed EC2 instance with an AWS KMS key. Then, choose the necessary KMS key from the list.

For a KMS key to be displayed in the list of available encryption keys, it must be stored in the AWS Region selected at [step 4](#) of the wizard, and the IAM role specified for the restore operation must have permissions to access the key. For more information on KMS keys, see [AWS Documentation](#).

TIP

If the necessary KMS key is not displayed in the list, or if you want to use a KMS key from an AWS account other than the AWS account to which the specified IAM role belongs, you can specify the Amazon Resource Number (ARN) of the key in the **Use the following encryption key** field.

For Veeam Backup for AWS to be able to encrypt the restored EBS volumes using the provided KMS key, either the IAM role or user specified for the restore operation, or the IAM role used to create the restore point selected at [step 2](#) of the wizard must have permissions to access the key.



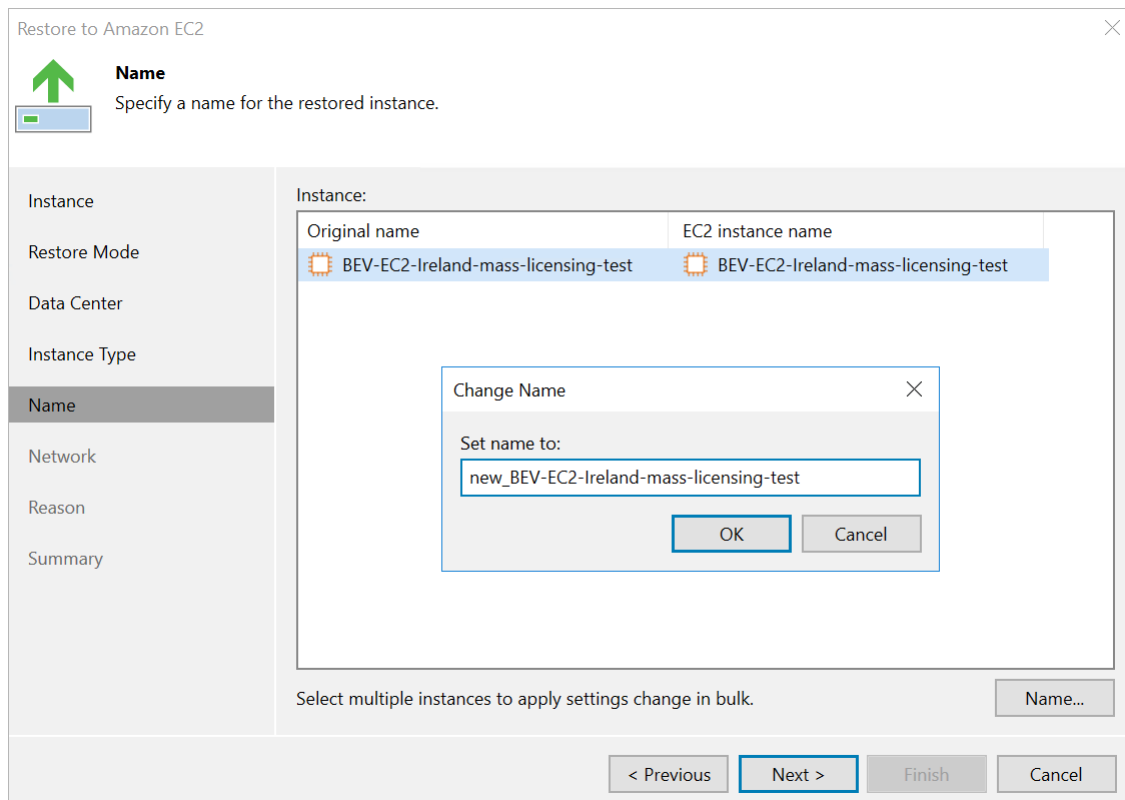
Step 6. Specify Instance Name

[This step applies only if you have selected the **Restore to a new location, or with different settings** option at the **Restore Mode** step of the wizard]

At the **Name** step of the wizard, you can specify a new name for the restored EC2 instance.

TIP

You can specify a single prefix or suffix and add it to the names of multiple restored EC2 instances. To do that, select the necessary instances and click **Name**. In the **Change Name** window, select the **Add prefix** or **Add suffix** check box, and provide the text that you want to add. Then, click **OK**.

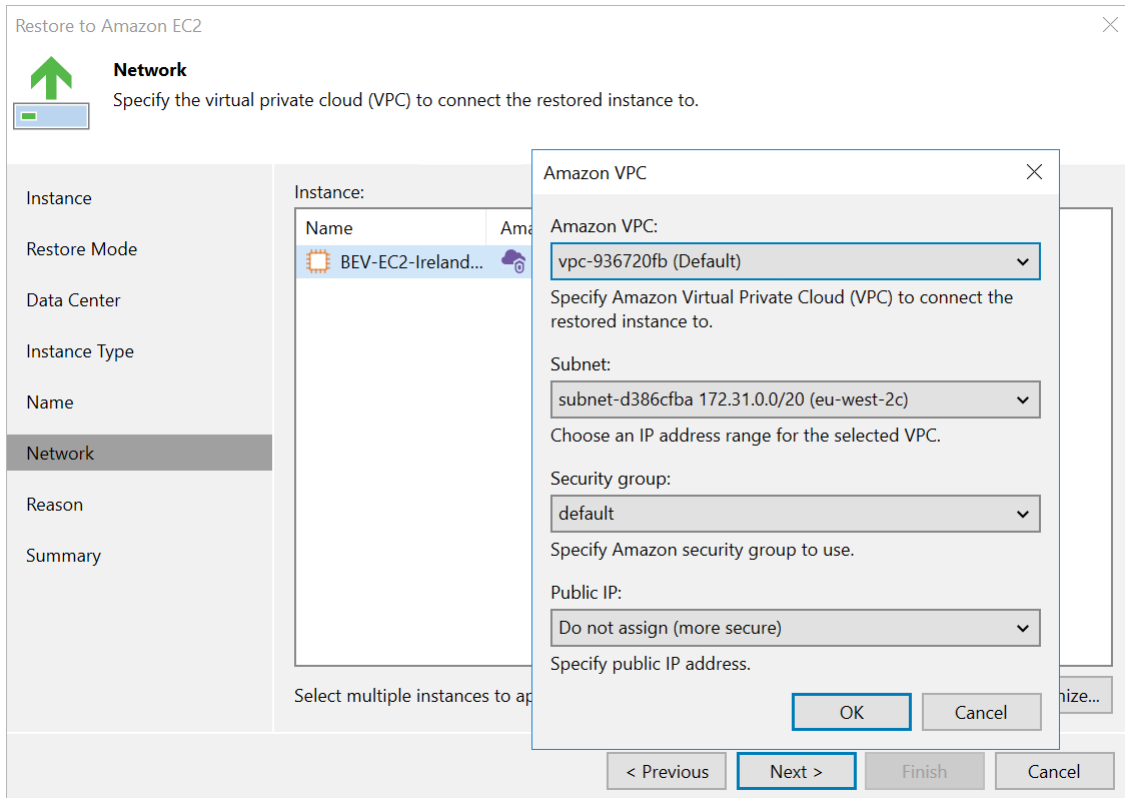


Step 7. Configure Network Settings

[This step applies only if you have selected the **Restore to a new location, or with different settings** option at the **Restore Mode** step of the wizard]

At the **Network** step of the wizard, you can select an Amazon VPC network to which the instance will be connected, a subnet in which the instance will be launched, and a security group that will be associated with the instance. To do that, select the EC2 instance and click **Customize**. You can also choose whether you want Veeam Backup & Replication to assign a public IP address to the restored instance.

For an Amazon VPC, subnet and security group to be displayed in the lists of available network specifications, they must be created in the AWS Region specified at [step 4](#) of the wizard, as described in [AWS Documentation](#).



Step 8. Specify Restore Reason

At the **Reason** step of the wizard, specify a reason for restoring the EC2 instance. The information you provide will be saved in the session history and you can reference it later.

The screenshot shows a wizard window titled "Restore to Amazon EC2" with a close button (X) in the top right corner. On the left side, there is a vertical navigation pane with the following items: "Instance", "Restore Mode", "Data Center", "Instance Type", "Name", "Network", "Reason" (which is highlighted with a dark grey background), and "Summary". Above the navigation pane, there is a green upward-pointing arrow icon and a small blue progress indicator. To the right of the arrow, the word "Reason" is written in bold. Below this, a text box contains the instruction: "Type in the reason for performing this restore operation. This information will be logged in the restore sessions history for later reference." The main area of the wizard is a large text input field with the label "Restore reason:" above it. The text "restore of failed EC2 instance" is entered into this field. Below the text field, there is a checkbox labeled "Do not show me this page again". At the bottom of the wizard, there are four buttons: "< Previous", "Next >", "Finish", and "Cancel".

Step 9. Finish Working with Wizard

At the **Summary** step of the wizard, review summary information and click **Finish**.

TIP

If you want to start the EC2 instance immediately after restore, select the **Power on target VM after restoring** check box.

Restore to Amazon EC2

Summary
You can copy the configuration information bellow for future reference.

Instance
Restore Mode
Data Center
Instance Type
Name
Network
Reason
Summary

Summary:
IAM role: EC2-restore-role
Data center: Europe (London)
Worker deployment in the production account: True
IAM role: Worker-deployment-role

Items:
Original instance name: BEV-EC2-Ireland-mass-licensing-test
EC2 instance name: new_BEV-EC2-Ireland-mass-licensing-test
Restore point: 6/26/2024 3:00:17 PM
EC2 instance type: t2.micro
VPC: vpc-936720fb (Default)
Subnet: subnet-d386cfba 172.31.0.0/20 (eu-west-2c)
Security group: default
Do not assign public IP address: True
KMS key: Preserve original settings

Power on target instance after restoring

< Previous Next > **Finish** Cancel

Performing Guest OS File Recovery

Veeam Backup & Replication allows you to use image-level backups to restore files and folders of various EC2 guest OS file systems from the Veeam Backup & Replication console. For more information, see the Veeam Backup & Replication User Guide, section [Guest OS File Recovery](#).

IMPORTANT

Guest OS File Recovery can be performed only using backup files stored in standard backup repositories for which you have specified access keys of an IAM user whose permissions are used to access the repository. To learn how to specify credentials for repositories, see sections [Creating New Repositories](#) and [Connecting to Existing Appliances](#).

You can also perform file-level recovery using the Veeam Backup for AWS Web UI. To learn how to recover files and folders to a local machine using file-level recovery browser, see [Performing File-Level Recovery](#).

Restoring from Microsoft Windows File Systems (FAT, NTFS or ReFS)

Before you start the restore operation, check the limitations and prerequisites described in the Veeam Backup & Replication User Guide, section [Requirements and Limitations](#).

To restore guest OS files and folders, do the following:

1. In the Veeam Backup & Replication console, open the **Home** view.
2. Navigate to **Backups > External Repository**.
3. Expand the backup policy that protects an EC2 instance whose files and folders you want to restore, select the necessary instance and click **Guest Files (Windows)** on the ribbon.
4. Complete the **File Level Restore** wizard as described in the Veeam Backup & Replication User Guide, section [Restoring VM Guest OS Files \(FAT, NTFS or ReFS\)](#).

Restoring from Linux, Unix and Other Supported File Systems

NOTE

You can restore files of Linux, Solaris, BSD, Novell Storage Services, Unix and Mac machines. For the list of supported file systems, see the Veeam Backup & Replication User Guide, section [Platform Support](#).

Before you start the restore operation, check the limitations and prerequisites described in the Veeam Backup & Replication User Guide, section [Requirements and Limitations](#).


To restore guest OS files and folders, do the following:

1. In the Veeam Backup & Replication console, open the **Home** view.
2. Navigate to **Backups > External Repository**.
3. Expand the backup policy that protects an EC2 instance whose files and folders you want to restore, select the necessary instance and click **Guest OS (Other)** on the ribbon.
4. Complete the **Guest File Restore** wizard as described in the Veeam Backup & Replication User Guide, section [Restoring VM Guest OS Files \(Multi-OS\)](#).

TIP

If the file system whose files and folders you want to restore is not included in the list of supported systems, do either of the following:

- Perform restore to the VMware vSphere environment using the Instant Disk Recovery technology. For more information, see the Veeam Backup & Replication User Guide, section [Restore from Other File Systems](#).
- Perform restore to the Microsoft Hyper-V environment using the Instant Recovery technology. For more information, see the Veeam Backup & Replication User Guide, section [Restore from Other File Systems](#).



Performing Application Restore

Veeam Backup & Replication provides auxiliary tools – Veeam Explorers – that allow you to restore application items directly from image-level backups of EC2 instances. You can restore items of the following applications: Microsoft Active Directory, Microsoft Exchange, Microsoft SharePoint, Microsoft SQL Server, Oracle and PostgreSQL. For more information on Veeam Explorers, see the [Veeam Explorers User Guide](#).

IMPORTANT

Application restore can be performed only using backup files stored in standard backup repositories for which you have specified access keys of an IAM user whose permissions are used to access the repository. To learn how to specify credentials for repositories, see sections [Editing Backup Repository Settings](#) and [Connecting to Existing Appliances](#).


To perform application restore, do the following:

1. In the Veeam Backup & Replication console, open the **Home** view.
2. Navigate to **Backups > External Repository**.

- Expand the backup policy that protects an EC2 instance whose application item you want to restore, select the necessary instance and click **Application Items** on the ribbon. Then, select the necessary application.
- In the restore wizard, select a backup that will be used to restore the application, specify a restore reason and click **Browse**.
- In the Veeam Explorer application, perform the steps described in the [Veeam Explorers User Guide](#).

IMPORTANT

The selected backup must be transactionally consistent. To learn how to create transactionally consistent backups, see [Creating EC2 Backup Policies](#).



EC2 Restore Using Web UI

Veeam Backup for AWS offers the following restore options:

- [Instance restore](#) – restores an entire EC2 instance.
- [Volume restore](#) – restores EBS volumes attached to an EC2 instance.
- [File-level recovery](#) – restores individual files and folders of an EC2 instance.

You can restore EC2 instance data to the most recent state or to any available restore point.

Performing EC2 Instance Restore

In case of a disaster, you can restore an entire EC2 instance from a cloud-native snapshot, snapshot replica or image-level backup. Veeam Backup for AWS allows you to restore one or more EC2 instances at a time, to the original location or to a new location.

NOTE

If you restore multiple EC2 instances that have the same EBS volume attached, Veeam Backup for AWS will restore one volume per each instance and enable the **Multi-Attach** option for every restored volume. To recover the source configuration, when the restore operation completes, manually delete extra EBS volumes in the AWS Management Console and attach the necessary volume to the instances.

For more information on Amazon EBS Multi-Attach, see [AWS Documentation](#).

How to Perform Instance Restore

To restore a protected EC2 instance, do the following:

1. [Check prerequisites and limitations](#).
2. [Launch the Instance Restore wizard](#).
3. [Select a restore point](#).
4. [Specify data retrieval settings for archived backups](#).
5. [Specify restore settings](#).
6. [Choose a restore mode](#).
7. [Enable encryption for EBS volumes](#).
8. [Specify EC2 instance settings](#).
9. [Configure network settings](#).
10. [Specify a restore reason](#).
11. [Finish working with the wizard](#).

Before You Begin

Before you restore EC2 instances, consider the following limitations:

- To restore an EC2 instance from a backup that is stored in an archive backup repository, you must retrieve the archived data first. You can either retrieve the archived data manually before you begin the restore operation, or launch the data retrieval process right from the **Restore** wizard. To learn how to retrieve data manually, see [Retrieving EC2 Data From Archive](#).
- When you restore an EC2 instance to a new location or with different settings, Veeam Backup for AWS will restore the instance with one network interface and will assign a new primary private IP address to the restored instance.
- If [stop protection](#) or [termination protection](#) are enabled for an EC2 instance, Veeam Backup for AWS will not be able to restore the instance and will raise an error notifying that you must disable stop protection or termination protection on the source instance.
- Veeam Backup for AWS does not support restore of IPv6 addresses, tags of Elastic IP addresses and prefixes assigned to Amazon EC2 network interfaces.
- When you restore an EC2 instance to the original location, Veeam Backup for AWS will restore the instance and all network interfaces that were attached to the source EC2 instance. However, consider the following:
 - If the Elastic IP address that was assigned to the source EC2 instance is still assigned to this EC2 instance, the address will be reassigned to the restored instance.
 - If the Elastic IP address is in use by any other EC2 instance, Veeam Backup for AWS will raise a warning. If you decide to proceed with the restore operation, the address will not be associated with the restored instance. Note that Veeam Backup for AWS will not allocate a new Elastic IP address to your AWS account.
 - If the Elastic IP address that was assigned to the source EC2 instance has been removed from AWS, Veeam Backup for AWS will attempt to restore this address using the native [AWS capabilities](#).
 - If private IP addresses that were assigned to the source EC2 instance are in use by the source or any other EC2 instance, Veeam Backup for AWS will raise a warning. If you decide to proceed with the restore operation, the restored EC2 instance will be assigned new private IP addresses.
 - If the source instance still exists in AWS, Veeam Backup for AWS will raise a warning. If you decide to proceed with the restore operation, the source EC2 instance, including all EBS volumes and all network interfaces attached to it, will be automatically deleted from AWS.


Note that EBS volumes excluded from the backup scope and volumes for which the `DeleteOnTermination` attribute is set to `false` will also be deleted from AWS.
- If you plan to restore an EC2 instance to an AWS Outpost, check the following prerequisites:
 - An IAM role you plan to specify for the restore operation must have the following permissions: `outposts:ListOutposts`, `outposts:GetOutpostInstanceTypes`. To grant the necessary permissions for the IAM role, use the AWS Management Console. For more information on how to grant permissions to an IAM role, see [AWS Documentation](#).
 - If an Outpost subnet is specified in the [worker instance network settings](#), restore of an EC2 instance to an AWS Region to which the AWS Outpost is connected may fail. The issue occurs if the default worker instance type is not supported for the AWS Outpost. To work around the issue, change the default worker profiles as described in section [Managing Worker Profiles](#).

Step 1. Launch Instance Restore Wizard

To launch the **Instance Restore** wizard, do the following.

1. Navigate to **Protected Data > EC2**.
2. Select the EC2 instance that you want to restore.
3. Click **Restore > Instance Restore**.

Alternatively, click the link in the **Restore Points** column. Then, in the **Available Restore Points** window, select the necessary restore point and click **Restore > Instance Restore**.



Step 2. Select Restore Point

At the **Instances** step of the wizard, select restore points to be used to perform the restore operation for each added instance. By default, Veeam Backup for AWS uses the most recent valid restore point. However, you can restore an EC2 instance to an earlier state.

IMPORTANT

If you select a restore point stored in an archive backup repository and the same restore point is also available in a standard backup repository, Veeam Backup for AWS will display the **Confirmation Restore** window. To proceed, choose whether you want to use the archived or standard restore point to perform the restore operation.

To select a restore point:

1. Select the EC2 instance.
2. Click **Restore Point**.
3. In the **Choose restore point** window, select the necessary restore point and click **Apply**.

To help you choose a restore point, Veeam Backup for AWS provides the following information on each available restore point:

- **Date** – the date when the restore point was created.
- **Type** – the type of the restore point:
 - *Snapshot* – a cloud-native snapshot created by a backup policy.
 - *Replica* – a snapshot replica created by a backup policy.
 - *Manual Snapshot* – a cloud-native snapshot created manually.
 - *Backup* – an image-level backup created by a backup policy.
 - *Archive* – an archived backup created by a backup policy.
- **State** – the state of the restore point (for image-level backups):
 - *Healthy* – the restore point has been verified by the health check session and reported to be healthy.
 - *Incomplete* – the restore point has been verified by the health check session and reported to be corrupted or incomplete.
- **Storage Class** – a storage class of the backup repository where the restore point is stored (for image-level backups).
- **Restore Point Region** – an AWS Region where the restore point is stored.

- **Account** – an AWS account where the restore point is stored.

The screenshot shows the Veeam Backup for AWS interface. The main window is titled "Instance Restore" and contains a "Choose instances to restore" panel on the left and a "Choose restore point" dialog on the right. The dialog displays a table of restore points with columns for Date, Type, State, Storage Class, and Restore Point R. The table lists several restore points, with the one dated 06/12/2024 1:58:34 PM and type Archive highlighted in green.

Date ↓	Type	State	Storage Class	Restore Point R
06/12/2024 1:58:34 PM	Snapshot	—	—	Europe (London)
06/12/2024 1:58:34 PM	Replica	—	—	Europe (Spain)
06/12/2024 1:58:34 PM	Backup	✓ Healthy	S3 Standard	Europe (London)
06/12/2024 1:58:34 PM	Archive	—	S3 Glacier Flexible ...	Europe (London)
06/12/2024 1:45:11 PM	Snapshot	—	—	Europe (London)
06/12/2024 1:45:11 PM	Replica	—	—	Europe (Spain)
06/12/2024 1:45:11 PM	Backup	✓ Healthy	S3 Standard	Europe (London)
06/12/2024 1:45:11 PM	Archive	—	S3 Glacier Flexible ...	Europe (London)

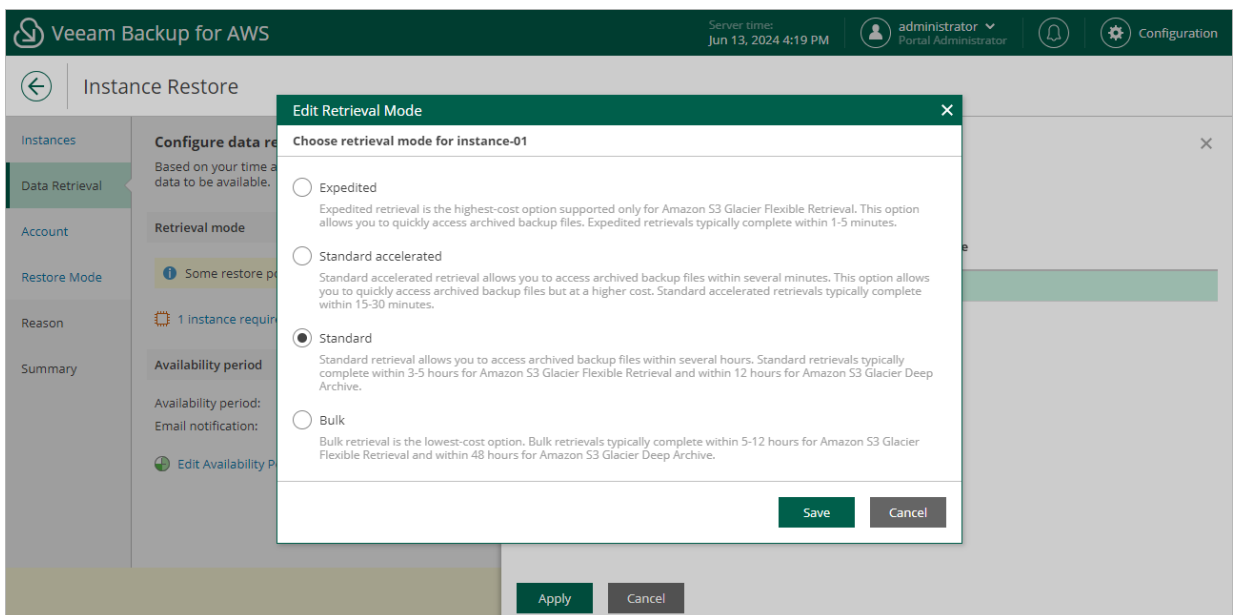
At the bottom of the dialog, there are "Apply" and "Cancel" buttons.

Step 3. Specify Data Retrieval Settings

[This step applies only if you have selected to restore from the archived restore point]

At the **Data Retrieval** step of the wizard, choose a retrieval mode and specify a period for which you want to keep the data available. To do that:

1. In the **Retrieval mode** section, click the link.
 - a. In the **Retrieval settings** window, for each processed EC2 instance, do the following:
 - i. Select an EC2 instance and click **Edit**.
 - ii. In the **Edit Retrieval Mode** window, select the retrieval mode that Veeam Backup for AWS will use to retrieve the archived data, and click **Save**. For more information on data retrieval modes, see [Retrieving EC2 Data From Archive](#).
 - b. To save changes made to the data retrieval settings, click **Apply**.



2. In the **Availability period** section, click **Edit Availability Period**.
 - a. In the **Availability settings** window, specify the number of days for which you want to keep the data available for restore operations.

IMPORTANT

If the time period expires while a restore operation is still running, the restore operation will fail. To work around the issue, you can instruct Veeam Backup for AWS to send an email notification when data is about to expire, and [manually extend the availability period](#) if required. To send the notification, select the **Send email notification** check box and choose when you want to be notified (that is, the number of hours remaining until data expiration).

b. To save changes made to the availability period settings, click **Apply**.

The screenshot shows the Veeam Backup for AWS interface. At the top, the header includes the Veeam logo, the text 'Veeam Backup for AWS', the server time 'Nov 16, 2023 11:55 AM', and the user 'administrator Portal Administrator'. The main content area is titled 'Instance Restore' and features a left-hand navigation menu with options: 'Instances', 'Data Retrieval' (highlighted), 'Account', 'Restore Mode', 'Reason', and 'Summary'. The 'Data Retrieval' section is expanded, showing 'Configure data retrieval settings' with a sub-section 'Availability period' that lists 'Availability period: 1 day' and 'Email notification: Disabled'. An 'Availability settings' dialog box is open on the right, containing the following text: 'Specify a time period for which you want the retrieved data to be available. If the time period expires while a restore operation is still running, the period will be automatically extended to keep the retrieved data available for 1 more day. You can also manually extend this period later if required.' Below this text are three settings: 'Keep data available for: 2 days', 'Send email notification: 1 hour before data expires', and 'Notify when data retrieval completes'. At the bottom of the dialog are 'Apply' and 'Cancel' buttons.

Step 4. Specify Restore Settings

At the **Account** step of the wizard, choose whether you want to use an IAM role or one-time access keys of an IAM user to allow Veeam Backup for AWS to perform the restore operation, and whether you want Veeam Backup for AWS to deploy worker instances in the production account. For information on the permissions that the IAM role or IAM user must have to perform the restore operation, see [EC2 Restore IAM Permissions](#).

IMPORTANT

Make sure that the specified IAM role or one-time access keys belong to an AWS account to which you plan to restore EC2 instances.

Specifying IAM Role

To specify an IAM role, select the **IAM role** option and choose the necessary IAM role from the list.

For an IAM role to be displayed in the **IAM Role** list, it must be added to Veeam Backup for AWS with the *Amazon EC2 Restore* operation selected as described in section [Adding IAM Roles](#). If you have not added the necessary IAM role to Veeam Backup for AWS beforehand, you can do it without closing the **Instance Restore** wizard. To do that, click **Add** and complete the [Add IAM Role](#) wizard.

IMPORTANT

It is recommended that you check whether the selected IAM role has all the permissions required to perform the operation. If some permissions of the IAM role are missing, the restore operation will fail to complete successfully. To run the IAM role permission check, click **Check Permissions** and follow the instructions provided in section [Checking IAM Role Permissions](#).

Specifying One-Time Access Keys

To specify one-time access keys, select the **Temporary access keys** option, and use the **Access key** and **Secret key** fields to provide the access key ID and the secret access key.

NOTE

Veeam Backup for AWS does not store one-time access keys in the configuration database.

Enabling Worker Deployment in Production Account

[This option applies only if you restore EC2 instances from image-level backups and have selected the **IAM role** option]

By default, Veeam Backup for AWS deploys worker instances used to perform restore operations in the [backup account](#). However, you can instruct Veeam Backup for AWS to deploy worker instances in a production account – that is, an account to which the EC2 instances will be restored. To do that, set the **Deploy workers in production account** toggle to *On*, and specify an IAM role that will be attached to the worker instances and used by Veeam Backup for AWS to communicate with these instances. The specified IAM role must belong to the same account to which the IAM role specified to perform the restore operation belongs, and must be assigned permissions listed in section [Worker Deployment Role Permissions in Production Accounts](#).

For an IAM role to be displayed in the **IAM role** list, it must be added to Veeam Backup for AWS with the *Production worker role* selected as described in section [Adding IAM Roles](#). If you have not added the necessary IAM role to Veeam Backup for AWS beforehand, you can do it without closing the **Add Policy** wizard. To do that, click **Add** and complete the **Add IAM Role** wizard.

IMPORTANT

Consider the following:

- If you instruct Veeam Backup for AWS to deploy worker instances in production accounts, you must assign additional permissions to the IAM role used to perform the restore operation. For more information on the required permissions, see [EC2 Restore IAM Permissions](#).
- It is recommended that you check whether both the IAM role specified in the **IAM role** section and the IAM role specified in the **Worker deployment** section have the required permissions. If some permissions of the IAM role are missing, the restore operation will fail to complete successfully. To run the IAM role permission check, click **Check Permissions** and follow the instructions provided in section [Checking IAM Role Permissions](#).

The screenshot shows the 'Instance Restore' wizard in Veeam Backup for AWS. The 'Account' step is active, showing options for IAM role and temporary access keys. The 'Worker deployment' section is also visible, with a toggle for 'Deploy workers in production account' set to 'On'. The interface includes a sidebar with navigation options like 'Instances', 'Data Retrieval', 'Account', 'Restore Mode', 'Reason', and 'Summary'. The top bar shows the server time as 'Jun 13, 2024 4:21 PM' and the user as 'administrator Portal Administrator'. At the bottom, there are 'Previous', 'Next', and 'Cancel' buttons.

Related Topics

- [Managing Worker Instances](#)
- [Managing Worker Configurations](#)

Step 5. Choose Restore Mode

At the **Restore Mode** step of the wizard, choose whether you want to restore the selected EC2 instance to the original or to a custom location. If you select the **Restore to new location, or with different settings** option, specify the target AWS Region where the restored EC2 instance will operate.

IMPORTANT

Consider the following:

- For Veeam Backup for AWS to be able to perform restore to the original location, the IAM role specified at the [Account](#) step of the wizard must belong to the AWS account to which the source EC2 instance belongs.
- Veeam Backup for AWS does not support restore to the original location if the source EC2 instance is still present in the location and [stop protection](#) or [termination protection](#) are enabled for the instance.

For more information on limitations and considerations, see [Before You Begin](#).

If you have AWS Outposts in your infrastructure, you can restore EC2 instances to an AWS Outpost. To do that:

1. Select the **Restore to new location, or with different settings** option.
2. From the drop-down list, select the AWS Region to which the AWS Outpost is connected.
3. Click the link to the right of **Select AWS Outpost**.
4. In the **Choose AWS Outpost** window, select the AWS Outpost where you want to restore the selected instances.
5. Click **Apply**.

NOTE

Consider the following:

- All objects residing in an AWS Outpost are encrypted.
- An AWS Outpost supports a limited list of instance types.

The screenshot shows the 'Instance Restore' wizard in Veeam Backup for AWS. The 'Restore Mode' step is active, showing two options: 'Restore to original location' (unselected) and 'Restore to new location, or with different settings' (selected). Below the selected option, a dropdown menu is set to 'Asia Pacific (Mumbai)'. The 'Select AWS Outpost' field is currently 'Not set...'. The interface includes a sidebar with navigation options like 'Instances', 'Data Retrieval', 'Account', 'Restore Mode', 'Encryption', 'Settings', 'Network', 'Reason', and 'Summary'. The top bar shows the server time as 'Jun 13, 2024 4:22 PM' and the user as 'administrator Portal Administrator'. At the bottom, there are 'Previous', 'Next', and 'Cancel' buttons.

Step 6. Enable Encryption

[This step applies only if you have selected the **Restore to new location, or with different settings** option at the **Restore Mode** step of the wizard]

At the **Encryption** step of the wizard, choose whether the restored EBS volumes of the processed EC2 instance will be encrypted with AWS KMS keys:

- If you do not want to encrypt the EBS volumes or want to apply the existing encryption scheme, select the **Use original encryption scheme** option.
- If you want to encrypt the EBS volumes, select the **Restore as encrypted instance** option and choose the necessary KMS key from the **Encryption key** list.

For a KMS key to be displayed in the list of available encryption keys, it must be stored in the AWS Region selected at [step 5](#) of the wizard and the IAM role or user specified for the restore operation at [step 4](#) of the wizard must have permissions to the key. For more information on KMS keys, see [AWS Documentation](#).

TIP

If the necessary KMS key is not displayed in the list, or if you want to use a KMS key from an AWS account other than the AWS account to which the specified IAM role belongs, you can select *Add custom key ARN* from the **Encryption key** drop-down list, and specify the amazon resource name (ARN) of the key in the **Add Custom Key ARN** window.

For Veeam Backup for AWS to be able to encrypt the restored EBS volumes using the provided KMS key, either the IAM role or user specified for the restore operation, or the IAM role used to create the restore point selected at [step 2](#) of the wizard must have permissions to access the key.

The screenshot shows the Veeam Backup for AWS Instance Restore wizard at the Encryption step. The interface includes a top navigation bar with the Veeam logo, server time (Jun 13, 2024 4:23 PM), user information (administrator, Portal Administrator), and a Configuration icon. The main content area is titled "Instance Restore" and features a left-hand navigation menu with options: Instances, Data Retrieval, Account, Restore Mode, Encryption (highlighted), Settings, Network, Reason, and Summary. The "Configure encryption settings" section contains the instruction: "Choose whether you want to use the original encryption scheme or encrypt the restored instances with a new key." Below this are two radio button options: "Use original encryption scheme" (unselected) and "Restore as encrypted instance" (selected). An "Encryption key:" dropdown menu is set to "aws/backup". An information icon with a link to a Veeam KB article is also present. At the bottom, there are "Previous", "Next", and "Cancel" buttons.

Step 7. Specify Instance Settings

The list of settings that you can configure for the restored EC2 instance depend on the option you choose at the **Choose Restore Mode** step of the wizard.

In This Section

- [Restoring to Original Location](#)
- [Restoring to New Location](#)

Restoring to Original Location

[This step applies only if you have selected the **Restore to original** option at the **Restore Mode** step of the wizard, and if the original Amazon machine image (AMI) that was used to launch the source instance has not been found]

At the **Settings** step of the wizard, select an AMI whose configuration will be used to launch the restored EC2 instance.

By default, Veeam Backup for AWS automatically chooses an AMI whose configuration is similar to the configuration of the restored instance. If Veeam Backup for AWS fails to choose an AMI automatically or you want to specify an AMI manually, click **Change AMI**. For an AMI to be displayed in the list of available AMIs, it must exist in the same AWS account and region in which the source instance resides.

TIP

When displaying the list of available AMIs, Veeam Backup for AWS applies a number of filtering criteria (such as *architecture*, *hypervisor*, *virtualization type*, *boot mode* and so on) to avoid misconfiguration issues. If the AMI that you want to use to launch the restored EC2 instance is not displayed in the list, you can try entering its ID in the search field of the **Instance settings** window – in this case, Veeam Backup for AWS will return the AMI without applying any filtering criteria.

The screenshot shows the Veeam Backup for AWS interface. The main window is titled "Instance Restore" and has a sidebar with navigation options: Instances, Data Retrieval, Account, Restore Mode, Encryption, Settings (selected), Network, Reason, and Summary. The "Configure restore settings" panel is active, showing a table with columns "Instance" and "Type". Two instances are listed: "instance-01" (t3.medium) and "instance-02" (t3.large). The "Instance settings" dialog is open, displaying a message: "The original AMI for the instance was al2023-ami-2023.4.20240611.0-kernel-6.1-x86_64 (ami-0b53285ea6c7a08...)". Below this is a search field for "AMI Name or ID" and a "Reset to Default and Apply" button. A table lists available AMIs with columns "Name", "ID", and "Description". The table contains several rows, with the second row highlighted in green. At the bottom of the dialog are "Apply" and "Cancel" buttons.

Name	ID	Description
(Apps4rent- H2O ai -AWS Linux ...	ami-0984e8b4e30c27f51	(Apps4rent- H2O ai -AWS Linux 2...
(SupportedImages) - LAMP Stack...	ami-0d1d179a1ddd51c7f	Linux Apache MySQL / MariaDB ...
(SupportedImages) - Amazon EC...	ami-0c382a1e38966b42d	al2023-ami-ecs-hvm-2023.0.202...
(SupportedImages) - Amazon EC...	ami-0ab28111256fd72d4	al2023-ami-ecs-hvm-2023.0.202...
(SupportedImages) - Amazon EC...	ami-05104d0d729326f48	al2023-ami-ecs-hvm-2023.0.202...
(SupportedImages) - Amazon EC...	ami-0fa7725872854dfe9	al2023-ami-ecs-hvm-2023.0.202...

Restoring to New Location

[This step applies only if you have selected the **Restore to new location, or with different settings** option at the **Restore Mode** step of the wizard]

TIP

As soon as you proceed to the **Settings** of the wizard, Veeam Backup for AWS will verify whether the original instance profiles attached to EC2 instances added to the restore session still exist in the AWS infrastructure, and whether the original IAM roles associated with these profiles have not been replaced or removed. If any of the conditions is not met, you will receive a warning in the **Instance Profile** column. To work around the issue, you can do either of the following:

- If an instance profile does not exist in the AWS infrastructure anymore, select another instance profile to be attached to the restored EC2 instance. Alternatively, proceed with the wizard to complete the restore operation without any profile attached.
- If the IAM role associated with an instance profile has been replaced or removed, select another instance profile that contains the required role, or replace the role associated with the profile in the AWS Management Console as described in [AWS Documentation](#).

At the **Settings** step of the wizard, do the following for each EC2 instance added to the restore session:

1. Select the EC2 instance.
2. If you want to specify a new name for the restored EC2 instance, click **Rename**.
3. If you want to configure type and profile settings for the restored EC2 instance, click **Edit**. For the list of all existing instance types, see [AWS Documentation](#).

To attach an instance profile to the restored EC2 instance or to replace the original instance profile that is already attached to the instance, set the **Configure instance profile** toggle to *On* and then select the necessary profile from the **Instance profile** drop-down list. If you set the **Instance profile** toggle to *Off*, the instance will be restored without any profile attached. For more information on instance profiles, see [AWS Documentation](#).

4. If you want to specify an Amazon machine image (AMI) for the restored EC2 instance, click **Change AMI**. For an AMI to be displayed in the list of available AMIs, it must exist in the target AWS account and AWS Region selected at [step 5](#) of the wizard.

Note that Veeam Backup for AWS automatically chooses either the AMI that was used to launch the source instance or an AMI whose configuration is similar to the configuration of the restored instance (if the original AMI has not been found).

TIP

When displaying the list of available AMIs, Veeam Backup for AWS applies a number of filtering criteria (such as *architecture*, *hypervisor*, *virtualization type*, *boot mode* and so on) to avoid misconfiguration issues. If the AMI that you want to use to launch the restored EC2 instance is not displayed in the list, you can try entering its ID in the search field of the **Instance settings** window – in this case, Veeam Backup for AWS will return the AMI without applying any filtering criteria.

Instance Restore

Instances

Data Retrieval

Account

Restore Mode

Encryption

Settings

Network

Reason

Summary

Configure restore settings

Specify settings for each instance that will be restored.

Edit Rename Change AMI View AMI Details

<input checked="" type="checkbox"/> Instance ↑	Type	Architecture
Selected: 1 of 2		
<input checked="" type="checkbox"/> instance-01	t3.medium	x86_64
<input type="checkbox"/> instance-02	t3.large	x86_64

Instance settings

Specify a type and profile for the restored instance.

Type Browse...

Configure instance profile On

Instance profile Browse...

IAM role ARN

Apply

Cancel

Step 8. Configure Network Settings

[This step applies only if you have selected the **Restore to new location, or with different settings** option at the **Restore Mode** step of the wizard]

At the **Network** step of the wizard, do the following for each EC2 instance in the list:

1. Select an EC2 instance and click **Edit**.
2. In the **Network settings** section of the opened window, choose an Amazon VPC to which the restored EC2 instance will be connected, select a subnet in which the EC2 instance will be launched and specify security groups that will be associated with the restored EC2 instance. To select security groups, click **Browse** to the right of **Security group**. Then, in the **Select Security Group** window, add security groups that must be associated with the instance, and click **Save**.

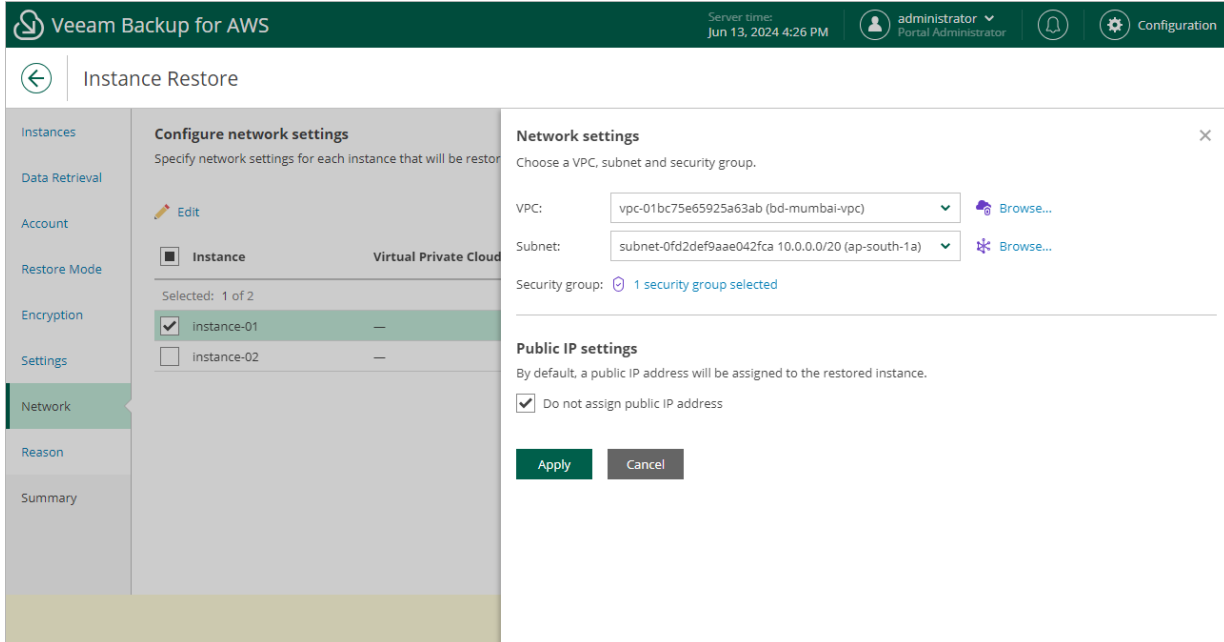
For a VPC, subnet and security group to be displayed in the lists of available network specifications, they must be created in the AWS Region specified at [step 5](#) of the wizard as described in [AWS Documentation](#).

If you restore EC2 instances to the AWS Outpost, for an Outpost subnet to be displayed in the **Subnet** drop-down list, choose the Amazon VPC that has one or more Outpost subnets.

IMPORTANT

When Veeam Backup for AWS backs up EC2 instances with IPv6 addresses assigned, it does not save the addresses. That is why when you restore these instances, IP addresses are assigned according to the settings specified in AWS for the subnet to which the instances are restored.

3. In the **Public IP** settings section of the opened window, choose whether you want Veeam Backup for AWS to assign a public IP address to the restored instance.



Related Resources

- [What Is Amazon VPC](#)
- [VPCs and Subnets](#)
- [Security Groups](#)

Step 9. Specify Restore Reason

At the **Reason** step of the wizard, you can specify a reason for restoring the EC2 instance. The information you provide will be saved in the session history and you can reference it later.

The screenshot shows the Veeam Backup for AWS interface. The top navigation bar includes the Veeam logo, the text "Veeam Backup for AWS", the server time "Jun 13, 2024 4:27 PM", the user "administrator Portal Administrator", and a "Configuration" link. The main content area is titled "Instance Restore" and features a left-hand navigation menu with the following items: Instances, Data Retrieval, Account, Restore Mode, Encryption, Settings, Network, Reason (highlighted in green), and Summary. The "Restore reason" section contains a text input field with the text "restoring failed EC2 instance". At the bottom of the wizard, there are three buttons: "Previous", "Next", and "Cancel".

Step 10. Finish Working with Wizard

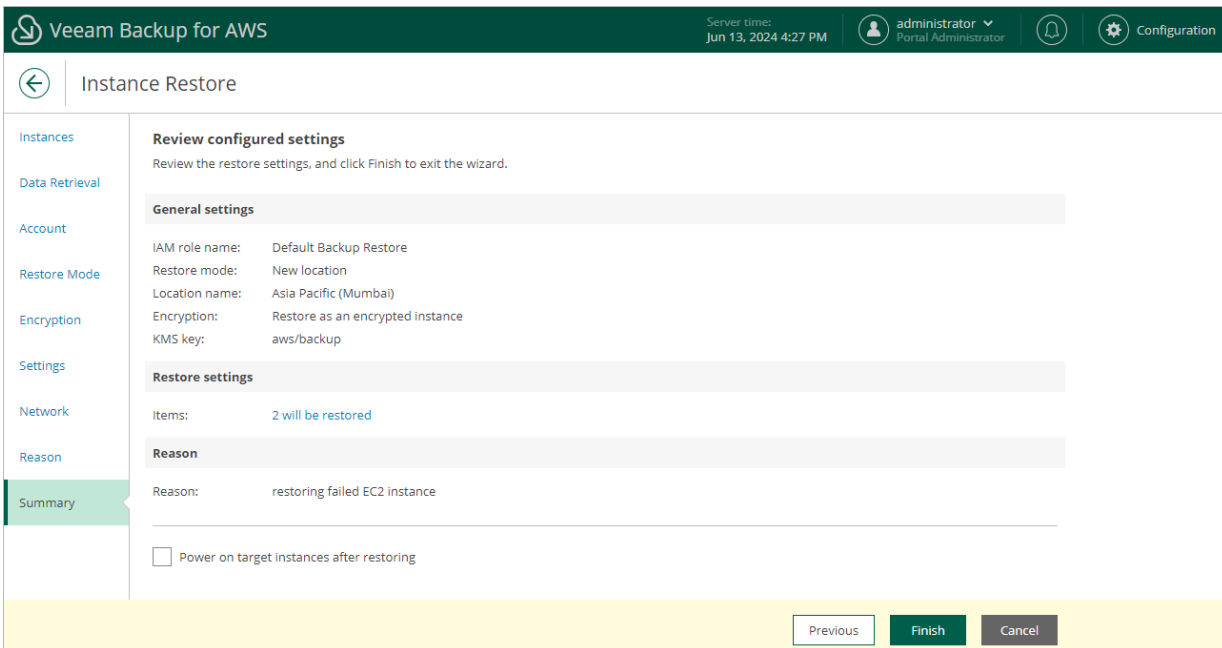
At the **Summary** step of the wizard, review summary information and click **Finish**.

TIP

If you want to start the restored EC2 instance as soon as the restore process completes, select the **Power on target instance after restoring** check box.

If you have selected the **Restore to original location** option at the **Restore Mode** step of the wizard, Veeam Backup for AWS will verify whether the original instance profile attached to the restored EC2 instance still exists in the AWS infrastructure, and whether the original IAM role associated with this profile has not been replaced, modified or removed. If any of the conditions is not met, you will receive a warning in the **Instance Profile Issue** window. To work around the issue, you can do either of the following:

- If the instance profile does not exist in the AWS infrastructure anymore, go back to [step 5](#), select the **Restore to new location, or with different settings** option and follow the instructions provided in section [Performing EC2 Instance Restore](#).
- If the IAM role associated with the instance profile has been replaced, modified or removed, go back to [step 5](#), select the **Restore to new location, or with different settings** option and choose another instance profile as described in section [Performing EC2 Instance Restore](#). Alternatively, click **Finish** to complete the restore session, and associate a new role in the AWS Management Console as described in [AWS Documentation](#).



Performing Volume Restore

In case a disaster strikes, you can restore corrupted EBS volumes of an EC2 instance from a cloud-native snapshot, snapshot replica or image-level backup. Veeam Backup for AWS allows you to restore EBS volumes to the original location or to a new location.

NOTE

Veeam Backup for AWS does not attach restored EBS volumes to any EC2 instances – the volumes are placed to the specified location as standalone EBS volumes.

How to Perform Volume Restore

To restore EBS volumes of a protected EC2 instance, do the following:

1. [Check prerequisites and limitations.](#)
2. [Launch the Volume Restore wizard.](#)
3. [Select a restore point.](#)
4. [Specify data retrieval settings for archived backups.](#)
5. [Specify restore settings.](#)
6. [Choose a restore mode.](#)
7. [Enable encryption for EBS volumes.](#)
8. [Specify the restored EBS volume name.](#)
9. [Specify a restore reason.](#)
10. [Finish working with the wizard.](#)

Before You Begin

To restore an EBS volume from a backup that is stored in the archive backup repository, the archived data must be retrieved first. You can retrieve the archived data manually before you begin the restore operation, or launch data retrieval from the **Restore** wizard. For more information on data retrieval, see [Retrieving EC2 Data From Archive](#).

If you plan to restore EBS volumes to an AWS Outpost, check the following prerequisites:


1. An IAM role you plan to specify for the restore operation must have the following permissions: `outposts:ListOutposts`, `outposts:GetOutpostInstanceTypes`. To grant the necessary permissions for the IAM role, use the [AWS Management Console](#).
2. If the Outpost subnet is specified in the [worker configuration settings](#), restore of EBS volumes to an AWS Region to which the AWS Outpost is connected may fail. The issue occurs if the default worker instance type is not supported for the AWS Outpost. In this case, change the default worker profiles as described in section [Managing Worker Profiles](#).

Step 1. Launch Volume Restore Wizard

To launch the **Volume Restore** wizard, do the following:

1. Navigate to **Protected Data > EC2**.
2. Select the EC2 instance whose EBS volumes you want to restore.
3. Click **Restore > Volume Restore**.

Alternatively, click the link in the **Restore Points** column. Then, in the **Available Restore Points** window, select the necessary restore point and click **Restore > Volume Restore**.



Step 2. Select Restore Point

At the **Instances** step of the wizard, select restore points to be used to perform the restore operation for each added instance. By default, Veeam Backup for AWS uses the most recent valid restore point. However, you can restore EBS volumes to an earlier state.

IMPORTANT

If you select a restore point stored in an archive backup repository and the same restore point is also available in a standard backup repository, Veeam Backup for AWS will display the **Confirmation Restore** window. To proceed, choose whether you want to use the archived or standard restore point to perform the restore operation.

To select a restore point:

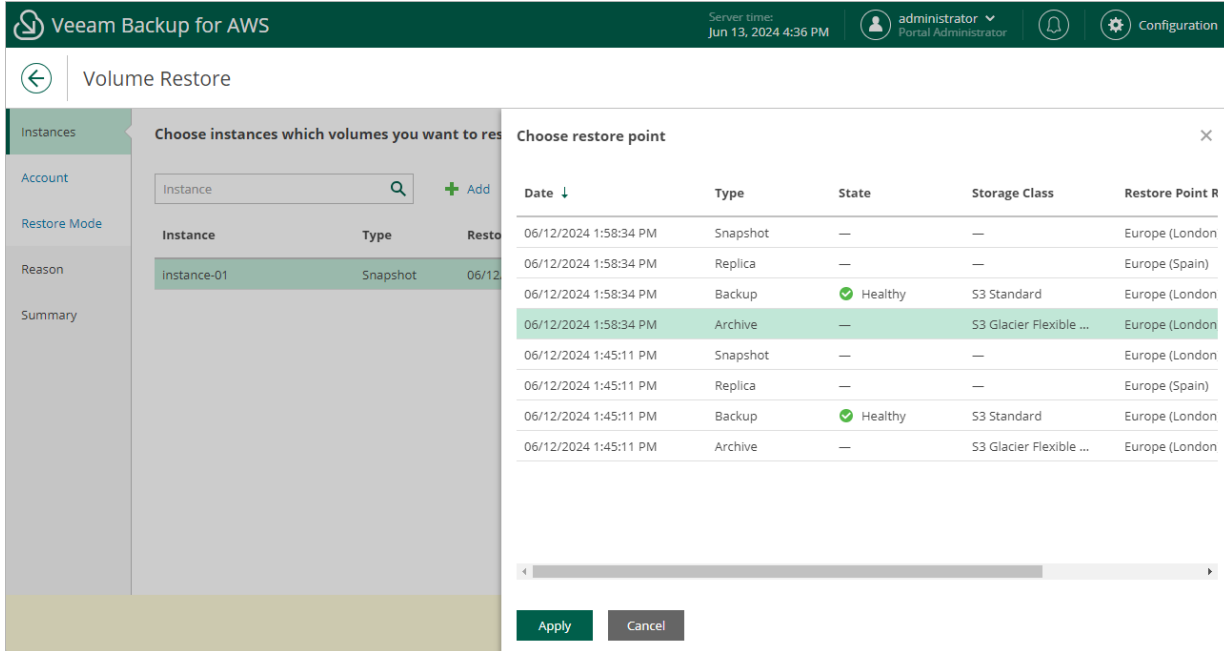
1. Select the EC2 instance.
2. Click **Restore Point**.
3. In the **Choose restore point** window, select the necessary restore point and click **Apply**.

To help you choose a restore point, Veeam Backup for AWS provides the following information on each available restore point:

- **Date** – the date when the restore point was created.
- **Size** – the size of the restore point.
- **Type** – the type of the restore point:
 - *Snapshot* – a cloud-native snapshot created by a backup policy.
 - *Replica* – a snapshot replica created by a backup policy.
 - *Manual Snapshot* – a cloud-native snapshot created manually.
 - *Backup* – an image-level backup created by a backup policy.
 - *Archive* – an archived backup created by a backup policy.
- **State** – the state of the restore point (for image-level backups):
 - *Healthy* – the restore point has been verified by the health check session and reported to be healthy.
 - *Incomplete* – the restore point has been verified by the health check session and reported to be corrupted or incomplete.
- **Storage Class** – a storage class of the backup repository where the restore point is stored (for image-level backups).
- **Restore Point Region** – an AWS Region where the restore point is stored (for cloud-native snapshots and snapshot replicas) or where the backup repository is located (for image-level backups).
- **IAM Role** – an IAM role used to create the restore point (for cloud-native snapshots and snapshot replicas).

TIP

If you want to restore only specific EBS volumes of the selected EC2 instances, you can exclude the unnecessary disks from the restore process. To do that, click **Exclusions** to open the **Specify exclusions** window, select check boxes next to the volumes that you do not want to restore, and click **Apply**.



The screenshot displays the Veeam Backup for AWS interface during a volume restore process. The main window is titled "Volume Restore". On the left, there is a sidebar with navigation options: "Instances", "Account", "Restore Mode", "Reason", and "Summary". The "Instances" section is active, showing a search bar and a "+ Add" button. Below the search bar, a table lists instances with columns "Instance", "Type", and "Restore". One instance, "Instance-01", is selected, with a "Snapshot" type and a restore date of "06/12/2024".

The "Choose restore point" panel is open, showing a table of restore points. The table has the following columns: "Date", "Type", "State", "Storage Class", and "Restore Point R". The table contains the following data:

Date ↓	Type	State	Storage Class	Restore Point R
06/12/2024 1:58:34 PM	Snapshot	—	—	Europe (London)
06/12/2024 1:58:34 PM	Replica	—	—	Europe (Spain)
06/12/2024 1:58:34 PM	Backup	✓ Healthy	S3 Standard	Europe (London)
06/12/2024 1:58:34 PM	Archive	—	S3 Glacier Flexible ...	Europe (London)
06/12/2024 1:45:11 PM	Snapshot	—	—	Europe (London)
06/12/2024 1:45:11 PM	Replica	—	—	Europe (Spain)
06/12/2024 1:45:11 PM	Backup	✓ Healthy	S3 Standard	Europe (London)
06/12/2024 1:45:11 PM	Archive	—	S3 Glacier Flexible ...	Europe (London)


At the bottom of the "Choose restore point" panel, there are "Apply" and "Cancel" buttons.

Step 3. Specify Data Retrieval Settings

[This step applies only if you have selected to restore from the archived restore point]

At the **Data Retrieval** step of the wizard, choose a retrieval mode and specify a period for which you want to keep the data available. To do that:

1. In the **Retrieval mode** section, click the link.
 - a. In the **Retrieval settings** window, for each processed EC2 instance, do the following:
 - i. Select an EC2 instance and click **Edit**.
 - ii. In the **Edit Retrieval Mode** window, select the retrieval mode that Veeam Backup for AWS will use to retrieve the archived data, and click **Save**. For more information on data retrieval modes, see [Retrieving EC2 Data From Archive](#).
 - b. To save changes made to the data retrieval settings, click **Apply**.



1. In the **Availability period** section, click **Edit Availability Period**.
 - a. In the **Availability settings** window, specify the number of days for which you want to keep the data available for restore operations.

IMPORTANT

If the time period expires while a restore operation is still running, the restore operation will fail. To work around the issue, you can instruct Veeam Backup for AWS to send an email notification when data is about to expire, and [manually extend the availability period](#) if required. To send the notification, select the **Send email notification** check box and choose when you want to be notified (that is, the number of hours remaining until data expiration).

b. To save changes made to the availability period settings, click **Apply**.

The screenshot shows the Veeam Backup for AWS interface. At the top, the header includes the Veeam logo, the text "Veeam Backup for AWS", the server time "Nov 16, 2023 12:17 PM", and the user "administrator Portal Administrator". The main content area is titled "Volume Restore" and has a left-hand navigation menu with options: "Instances", "Data Retrieval" (selected), "Account", "Restore Mode", "Reason", and "Summary". The "Data Retrieval" section is expanded, showing "Configure data retrieval settings" with a sub-section "Availability period" that lists "Availability period: 1 day" and "Email notification: Disabled". An "Availability settings" dialog box is open on the right, containing the following text: "Specify a time period for which you want the retrieved data to be available. If the time period expires while a restore operation is still running, the period will be automatically extended to keep the retrieved data available for 1 more day. You can also manually extend this period later if required." The dialog has three input fields: "Keep data available for: 3 days", "Send email notification: 1 hour before data expires" (checked), and "Notify when data retrieval completes" (checked). At the bottom of the dialog are "Apply" and "Cancel" buttons.

Step 4. Specify Restore Settings

At the **Account** step of the wizard, choose whether you want to use an IAM role or one-time access keys of an IAM user to allow Veeam Backup for AWS to perform the restore operation, and whether you want Veeam Backup for AWS to deploy worker instances in the [production account](#). For information on the permissions that the IAM role or IAM user must have to perform the restore operation, see [EC2 Restore IAM Permissions](#).

IMPORTANT

Make sure that the specified IAM role or one-time access keys belong to an AWS account to which you plan to restore EBS volumes.

Specifying IAM Role

To specify an IAM role, select the **IAM role** option and choose the necessary IAM role from the list.

For an IAM role to be displayed in the **IAM Role** list, it must be added to Veeam Backup for AWS with the *Amazon EC2 Restore* operation selected as described in section [Adding IAM Roles](#). If you have not added the necessary IAM role to Veeam Backup for AWS beforehand, you can do it without closing the **Volume Restore** wizard. To do that, click **Add** and complete the [Add IAM Role](#) wizard.

IMPORTANT

It is recommended that you check whether the selected IAM role has all the permissions required to perform the operation. If some permissions of the IAM role are missing, the restore operation will fail to complete successfully. To run the IAM role permission check, click **Check Permissions** and follow the instructions provided in section [Checking IAM Role Permissions](#).

Specifying One-Time Access Keys

To specify one-time access keys, select the **Temporary access keys** option, and use the **Access key** and **Secret key** fields to provide the access key ID and the secret access key.

NOTE

Veeam Backup for AWS does not store one-time access keys in the configuration database.

Enabling Worker Deployment in Production Account

[This option applies only if you restore volumes from image-level backups and have selected the **IAM role** option]

By default, Veeam Backup for AWS deploys worker instances used to perform restore operations in the [backup account](#). However, you can instruct Veeam Backup for AWS to deploy worker instances in a production account – that is, an account to which the volumes will be restored. To do that, set the **Deploy workers in production account** toggle to *On*, and specify an IAM role that will be attached to the worker instances and used by Veeam Backup for AWS to communicate with these instances. The specified IAM role must belong to the same account to which the IAM role specified to perform the restore operation belongs, and must be assigned permissions listed in section [Worker Deployment Role Permissions in Production Accounts](#).

For an IAM role to be displayed in the **IAM role** list, it must be added to Veeam Backup for AWS with the *Production worker role* selected as described in section [Adding IAM Roles](#). If you have not added the necessary IAM role to Veeam Backup for AWS beforehand, you can do it without closing the **Add Policy** wizard. To do that, click **Add** and complete the **Add IAM Role** wizard.

IMPORTANT

If you instruct Veeam Backup for AWS to deploy worker instances in production accounts, you must assign additional permissions to the IAM role used to perform the restore operation. For more information on the required permissions, see [EC2 Restore IAM Permissions](#).

The screenshot shows the 'Volume Restore' wizard in Veeam Backup for AWS. The current step is 'Choose IAM role and specify worker deployment settings'. The interface includes a top navigation bar with the Veeam logo, server time (Jun 13, 2024 4:38 PM), and user information (administrator, Portal Administrator). A left sidebar contains navigation options: Instances, Data Retrieval, Account (selected), Restore Mode, Reason, and Summary. The main content area is divided into two sections: 'IAM role' and 'Worker deployment'. In the 'IAM role' section, the 'IAM role' radio button is selected, and a dropdown menu shows 'Default Backup Restore (Default Backup Restore)'. There are '+ Add' and 'Check Permissions' buttons. Below this, there are input fields for 'Access key:' and 'Secret key:', followed by an information icon and a note: 'The keys are used to perform this operation only. They are not saved or stored. To learn what permissions are required for performing the operation, see the User Guide.' The 'Worker deployment' section has a toggle switch for 'Deploy workers in production account:' which is turned 'On'. Below this is another information icon and note: 'To be able to restore instances with volumes encrypted using the default AWS managed key, it is required to deploy worker instances in the production account.' At the bottom of the wizard, there are 'Previous', 'Next', and 'Cancel' buttons.

Related Topics

- [Managing Worker Instances](#)
- [Managing Worker Configurations](#)

Step 5. Choose Restore Mode

At the **Restore Mode** step of the wizard, choose whether you want to restore the selected EBS volumes to the original or to a custom location. If you select the **Restore to new location, or with different settings** option, specify the AWS Region and Availability Zone to which Veeam Backup for AWS will place the restored EBS volumes.

IMPORTANT

For Veeam Backup for AWS to be able to perform restore to the original location, the IAM role specified at the [Account](#) step of the wizard must belong to the AWS account to which the source EC2 instance belongs.

If you have AWS Outposts in your infrastructure, you can restore EBS volumes to an AWS Outpost. To do that:

1. Select the **Restore to new location, or with different settings** option.
2. From the region drop-down list, select the AWS Region to which the AWS Outpost is connected.
3. From the **Availability zone** drop-down list, select the Availability Zone that the AWS Outpost is homed to.
4. Click the link to the right of **Select AWS Outpost**.
5. In the **Choose AWS Outpost** window, select the AWS Outpost where you want to restore EBS volumes of the selected instances.
6. Click **Apply**.

NOTE

Consider the following:

- All objects residing in an AWS Outpost are encrypted.
- An AWS Outpost supports a limited list of EBS volume types. If the type of the restored EBS volume is not supported in the selected AWS Outpost, the restore operation will fail.
- Before you select an AWS Outpost, check limitations and requirements described in section [Before You Begin](#).

The screenshot shows the 'Volume Restore' wizard in Veeam Backup for AWS. The top navigation bar includes the Veeam logo, 'Veeam Backup for AWS', server time 'Jun 13, 2024 4:38 PM', user 'administrator', and 'Portal Administrator'. The left sidebar lists steps: Instances, Data Retrieval, Account, Restore Mode (selected), Encryption, Settings, Reason, and Summary. The main content area is titled 'Choose restore mode' and contains two radio button options: 'Restore to original location' and 'Restore to new location, or with different settings' (selected). Below the selected option are dropdown menus for 'Region' (Asia Pacific (Mumbai)) and 'Availability zone' (ap-south-1b). A 'Select AWS Outpost' field is currently set to 'Not set...'. At the bottom, there are 'Previous', 'Next', and 'Cancel' buttons.

Step 6. Enable Encryption

[This step applies only if you have selected the **Restore to new location, or with different settings** option at the **Restore Mode** step of the wizard]

At the **Encryption** step of the wizard, choose whether the restored EBS volumes will be encrypted with AWS KMS keys:

- If you do not want to encrypt the EBS volumes or want to apply the existing encryption scheme, select the **Use original encryption scheme** option.
- If you want to encrypt the EBS volumes, select the **Restore as encrypted volumes** option and choose the necessary KMS key from the **Encryption key** list.

For a KMS key to be displayed in the list of available encryption keys, it must be stored in the AWS Region selected at [step 5](#) of the wizard and the IAM role or user specified for the restore operation at [step 4](#) of the wizard must have permissions to access the key. For more information on KMS keys, see [AWS Documentation](#).

TIP

If the necessary KMS key is not displayed in the list, or if you want to use a KMS key from an AWS account other than the AWS account to which the specified IAM role belongs, you can select *Add custom key ARN* from the **Encryption key** drop-down list, and specify the amazon resource name (ARN) of the key in the **Add Custom Key ARN** window.

For Veeam Backup for AWS to be able to encrypt the restored EBS volumes using the provided KMS key, either the IAM role or user specified for the restore operation, or the IAM role used to create the restore point selected at [step 2](#) of the wizard must have permissions to access the key.

The screenshot shows the 'Volume Restore' wizard in the Veeam Backup for AWS console. The 'Encryption' step is active, showing options to 'Configure encryption settings'. The 'Restore as encrypted volume' option is selected. The 'Encryption key' dropdown menu is set to 'aws/backup'. An information icon and a link to a Veeam KB article are visible below the dropdown. The bottom of the wizard has 'Previous', 'Next', and 'Cancel' buttons.

Veeam Backup for AWS

Server time: Jun 13, 2024 4:39 PM

administrator Portal Administrator

Configuration

Volume Restore

Instances

Data Retrieval

Account

Restore Mode

Encryption

Settings

Reason

Summary

Configure encryption settings
Choose whether you want to use the original encryption scheme or encrypt the restored instances with a new key.

Use original encryption scheme

Restore as encrypted volume

Encryption key:
aws/backup

i To learn how to work with AWS encryption keys, see [this Veeam KB article](#).

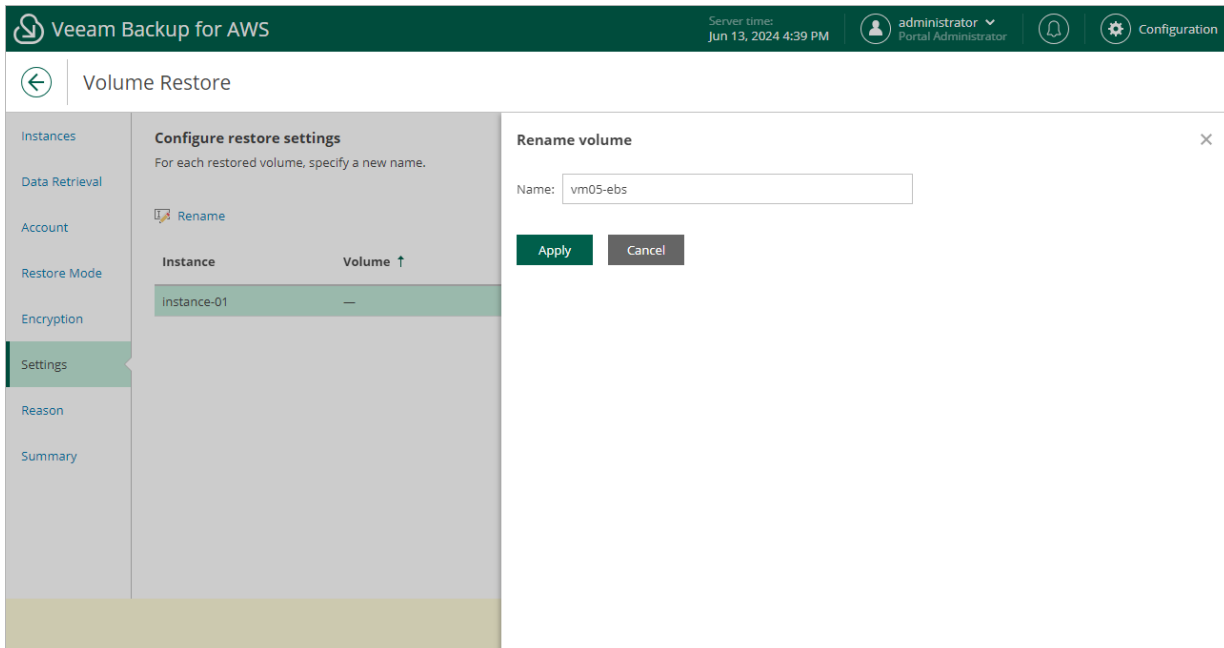
Previous Next Cancel

Step 7. Specify EBS Volume Name

[This step applies only if you have selected the **Restore to new location, or with different settings** option at the **Restore Mode** step of the wizard]

At the **Settings** step of the wizard, you can specify a name for each restored EBS volume:

1. Select the necessary EBS volume and click **Rename**.
2. In the **Rename volume** window, specify a name for the restored EBS volume and click **Apply**.



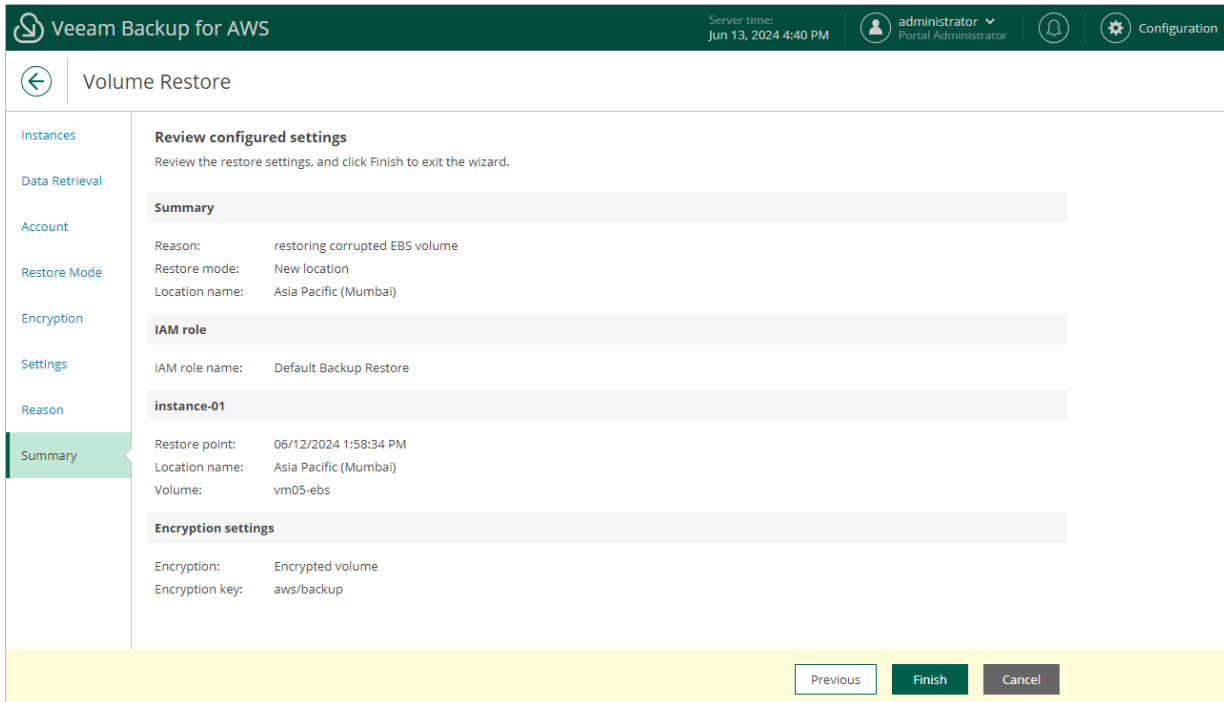
Step 8. Specify Restore Reason

At the **Reason** step of the wizard, you can specify a reason for restoring EBS volumes. The information you provide will be saved in the session history and you can reference it later.

The screenshot shows the Veeam Backup for AWS interface. The top navigation bar includes the Veeam logo, the text "Veeam Backup for AWS", the server time "Jun 13, 2024 4:40 PM", the user "administrator Portal Administrator", and a "Configuration" link. The main content area is titled "Volume Restore" and features a left-hand navigation menu with the following items: Instances, Data Retrieval, Account, Restore Mode, Encryption, Settings, Reason (highlighted in green), and Summary. The "Restore reason" section contains the instruction "Specify a reason for performing the restore operation." and a text input field with the value "restoring corrupted EBS volume". At the bottom of the wizard, there are three buttons: "Previous", "Next", and "Cancel".

Step 9. Finish Working with Wizard

At the **Summary** step of the wizard, review summary information and click **Finish**.



Performing File-Level Recovery

In case a disaster strikes, you can recover corrupted or missing files of an EC2 instance from a cloud-native snapshot or image-level backup.

IMPORTANT

Restore of files and folders is supported only for the following file systems: FAT, FAT32, NTFS, ext2, ext3, ext4, XFS, Btrfs. For EC2 instances running Microsoft Windows OSes, Veeam Backup for AWS supports file-level recovery only for basic volumes.

You can use the following options:

- Download the necessary files and folders to a local machine.
- Restore the files and folders of the source EC2 instance to the original location.

By default, Veeam Backup for AWS restores files and folders to a local machine. If you want to perform restore to the original location, you must enable the [Additional restore mode](#) in the restore settings.

IMPORTANT

Before you start the restore operation, check the prerequisites described in section [Before You Begin](#).

To learn how EC2 file-level recovery works, see [File-Level Recovery](#). To learn how to configure network settings that will be used to deploy workers during the restore process, see [Managing Worker Configurations](#).

How to Perform EC2 File-Level Recovery

To recover files and folders of a protected EC2 instance, do the following:

1. [Check prerequisites and limitations.](#)
2. [Launch the EC2 File-level Recovery wizard.](#)
3. [Select a restore point.](#)
4. [Specify restore settings.](#)
5. [Specify a restore reason.](#)
6. [Finish working with the wizard – start a recovery session.](#)
7. [Choose files and folders to recover.](#)
8. [Stop the recovery session.](#)

Before You Begin

Before you start file-level recovery, check the following limitations and prerequisites:

- To recover files and folders of an EC2 instance from a backup that is stored in an archive backup repository, you must retrieve the archived data manually before you begin the file-level recovery operation. For more information on data retrieval, see [Retrieving EC2 Data From Archive](#).
- The **443** port must be open on worker instances to allow inbound network access from the machine from which you plan to open the file-level recovery browser. To enable access for a worker instance, update the security group specified in [worker instance settings](#) to add an inbound rule. To learn how to add rules to security groups, see [AWS Documentation](#).

If you want worker instances to operate in a private network, enable the [private network deployment](#) functionality and configure specific VPC endpoints for all subnets to which the worker instances will be connected. Alternatively, configure VPC endpoints for all subnets as described in section [Appendix C. Configuring Endpoints in AWS](#).

TIP

It is recommended that you run a file-level recovery test before you start a file-level recovery operation in a specific AWS Region. For more information, see [Testing Configurations for FLR](#).

Restoring to Original Location

If you plan to perform file-level recovery to the original location, consider the following additional limitations and prerequisites:

- To perform restore to the original location, Veeam Backup for AWS deploys worker instances in the backup account. That is why you must specify network settings for worker instances beforehand as described in section [Adding Configurations for Backup Account](#).
- [For Linux-based EC2 instances] Python v2 or v3 with module 6 must be installed on the source instance.
- The source instance must be configured to communicate with AWS System Manager. To learn how to configure instance permissions for Systems Manager, see [AWS Documentation](#).
- SSM Agent must be installed on the source instance. To learn how to install SSM Agent, see [AWS Documentation](#).

- The IAM role attached to the source EC2 instance must meet the following requirements:
 - a. The IAM role must be included in the instance profile. For more information on instance profiles, see [AWS Documentation](#).
 - b. The Amazon EC2 service must be granted permissions to assume the IAM role.

To allow the Amazon EC2 service to assume the IAM role, configure trust relationships for the role and add the following statement to the trust policy.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "sts:AssumeRole",
      "Principal": {
        "Service": "ec2.amazonaws.com"
      }
    }
  ]
}
```

- c. During the file-level recovery session, Veeam Backup for AWS will create a temporary IAM role in the backup account to perform data transmission using [Amazon Kinesis Data Streams](#). That is why the IAM role attached to the source EC2 instance must have the permissions to assume the temporary role, as well as the permissions to work with Amazon Simple Queue Service (SQS) and Amazon Kinesis Data Streams:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "VisualEditor0",
      "Effect": "Allow",
      "Action": [
        "sqs:ListQueues",
        "sqs:GetQueueUrl",
        "kinesis:List*",
        "kinesis:Describe*",
        "kinesis:Get*",
        "sqs:GetQueueAttributes",
        "sqs:ListDeadLetterSourceQueues"
      ],
      "Resource": "*"
    },
    {
      "Sid": "VisualEditor1",
      "Effect": "Allow",
      "Action": "sts:AssumeRole",
      "Resource": "arn:aws:iam:: <service-account-id> :role/ veeam
_rto_<original-instance-id>"
    }
  ]
}
```

Where the `<service-account-id>` is an AWS ID of the trusted backup AWS account, and `<original-instance-id>` is an AWS ID of the source EC2 instance.

- If the source EC2 instance operates in a private network, you must create the following VPC endpoints for the subnet to which the instance is connected:
 - o `com.amazonaws.<region>.ec2messages`
 - o `com.amazonaws.<region>.ssm`
 - o `com.amazonaws.<region>.sqs`
 - o `com.amazonaws.<region>.kinesis-streams`
 - o `com.amazonaws.<region>.sts`

To learn how to create interface VPC endpoints, see [AWS Documentation](#).

Step 1. Launch EC2 File-level Recovery Wizard

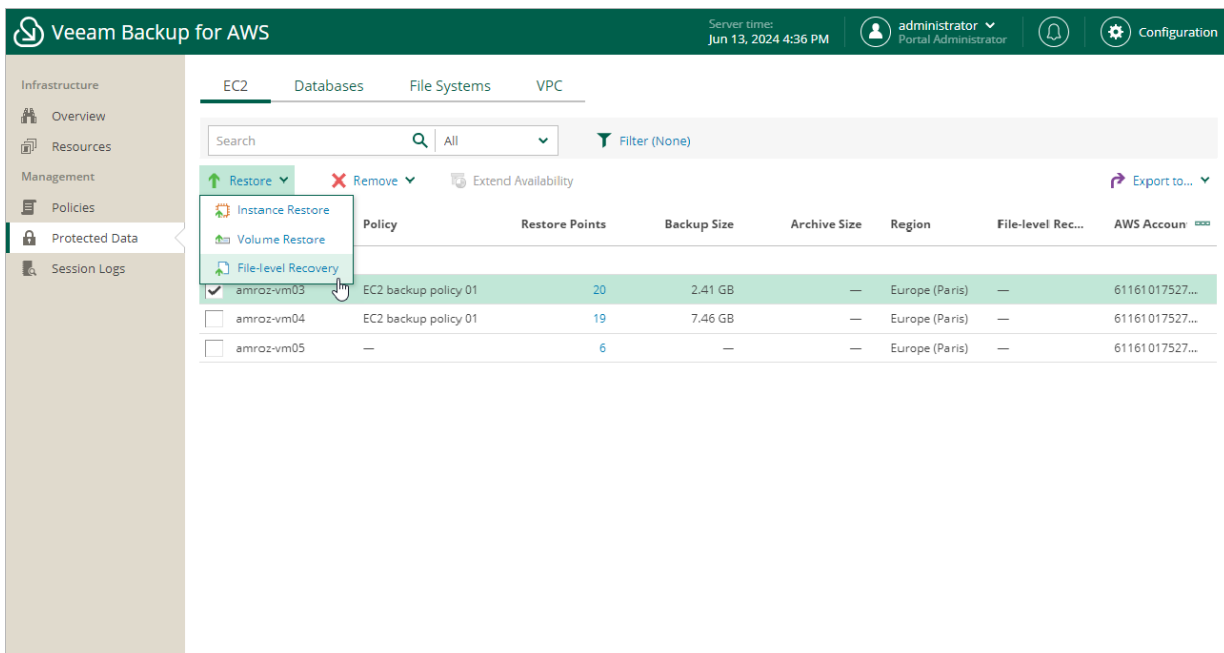
To launch the **EC2 File-level Recovery** wizard, do the following:

1. Navigate to **Protected Data > EC2**.
2. Select the EC2 instance whose files and folders you want to recover.
3. Click **Restore > File-level Recovery**.

Alternatively, click the link in the **Restore Points** column. Then, in the **Available Restore Points** window, select the necessary restore point and click **Restore > File-level Recovery**.

IMPORTANT

If you select multiple EC2 instances, you will not be able to proceed with the **EC2 File-level Recovery** wizard.



Step 2. Select Restore Point

At the **Instances** step of the wizard, select restore points to be used to perform the restore operation for added instance. By default, Veeam Backup for AWS uses the most recent valid restore point. However, you can restore files and folders of the backed-up EC2 instance to an earlier state.

To select a restore point:

1. Select the EC2 instance.
2. Click **Restore Point**.
3. In the **Choose restore point** window, select the necessary restore point and click **Apply**.

To help you choose a restore point, Veeam Backup for AWS provides the following information on each available restore point:

- **Date** – the date when the restore point was created.
- **Type** – the type of the restore point:
 - *Snapshot* – a cloud-native snapshot created by a backup policy.
 - *Replica* – a snapshot replica created by a backup policy.
 - *Manual Snapshot* – a cloud-native snapshot created manually.
 - *Backup* – an image-level backup created by a backup policy.
 - *Archive* – an archived backup created by a backup policy.
- **State** – the state of the restore point (for image-level backups):
 - *Healthy* – the restore point has been verified by the health check session and reported to be healthy.
 - *Incomplete* – the restore point has been verified by the health check session and reported to be corrupted or incomplete.
- **Storage Class** – a storage class of the backup repository where the restore point is stored (for image-level backups).
- **Restore Point Region** – an AWS Region where the restore point is stored (for cloud-native snapshots and snapshot replicas) or where the backup repository is located (for image-level backups).
- **IAM Role** – an IAM role used to create the restore point (for cloud-native snapshots and snapshot replicas) or an IAM role used to access the backup repository (for image-level backups).

IMPORTANT

To recover files and folders of an EC2 instance from a restore point that is stored in the archive backup repository of the S3 Glacier Flexible Retrieval or S3 Glacier Deep Archive storage class, you must retrieve the archived data manually before you begin the file-level recovery operation. For more information on data retrieval, see [Retrieving EC2 Data From Archive](#).

The screenshot shows the Veeam Backup for AWS interface. The main window is titled 'EC2 File-level Recovery'. On the left, there are tabs for 'Instances', 'Restore Settings', 'Reason', and 'Summary'. The 'Instances' tab is active, showing a table with columns for 'Instance', 'Type', and 'Restore Point'. The instance 'amroz-vm03' is selected, with a type of 'Snapshot' and a restore point of '10/19/2023'. A 'Choose restore point' dialog is open, displaying a table of restore points for this instance. The table has columns for 'Date', 'Type', 'State', 'Storage Class', and 'Restore Point R'. The selected restore point is from 10/19/2023 at 10:00:15 AM, with a state of 'Healthy' and storage class 'S3 Standard'. The dialog also has 'Apply' and 'Cancel' buttons at the bottom.

Date	Type ↑	State	Storage Class	Restore Point R
10/09/2023 4:01:19 PM	Manual snaps...	—	—	Europe (Pari
10/09/2023 4:02:24 PM	Manual snaps...	—	—	Europe (Pari
10/09/2023 10:00:18 AM	Backup	✓ Healthy	S3 Standard	Europe (Pari
07/24/2023 10:00:16 AM	Backup	✓ Healthy	S3 Standard	Europe (Pari
10/16/2023 10:00:13 AM	Backup	✓ Healthy	S3 Standard	Europe (Pari
08/28/2023 10:00:22 AM	Backup	✓ Healthy	S3 Standard	Europe (Pari
08/07/2023 10:00:16 AM	Backup	✓ Healthy	S3 Standard	Europe (Pari
08/21/2023 10:00:16 AM	Backup	✓ Healthy	S3 Standard	Europe (Pari
08/14/2023 10:00:23 AM	Backup	✓ Healthy	S3 Standard	Europe (Pari
07/31/2023 10:00:15 AM	Backup	✓ Healthy	S3 Standard	Europe (Pari
10/09/2023 10:00:18 AM	Snapshot	—	—	Europe (Pari
10/19/2023 10:00:15 AM	Snapshot	—	—	Europe (Pari
10/18/2023 10:00:16 AM	Snapshot	—	—	Europe (Pari
10/16/2023 10:00:13 AM	Snapshot	—	—	Europe (Pari
10/17/2023 10:08:55 AM	Snapshot	—	—	Europe (Pari
10/19/2023 10:00:15 AM	Replica	—	—	Europe (Mila

Step 3. Specify Restore Settings

At the **Restore Settings** step of the wizard, choose whether you want to restore files and folders to the original location, and to deploy worker instances in the production account.

Configuring Restore To Original Location

[This option applies only if you choose not to deploy worker instances in the production account]

To be able to restore files and folders to the original EC2 instance, set the **Additional restore mode** toggle to *On*.

To perform the restore operation, Veeam Backup for AWS will use the IAM role attached to the source instance. That is why before enabling the additional restore mode, assign all the required permissions to the IAM role. For more information on the required permissions, see [Before You Begin](#).

IMPORTANT

Consider the following limitations:

- For EC2 instances running Linux OS, restore of files and folders to the original location is supported only for systemd-based distributions.
- For EC2 instances running Windows OS, restore of files and folders to the original location is supported only if Windows Management Framework (WMF) version 5.1 is installed on the processed instances.

To restore files and folders to the source EC2 instance, Veeam Backup for AWS uses Amazon Kinesis Data Streams. Kinesis Data Streams are charged on a per-shard basis. By default, Veeam Backup for AWS uses streams that are composed of 1 shard with a fixed data transfer rate of 1 MB per second. However, you can change the number of shards in the streams by moving the **Restore rate** slider. For more information on Kinesis Data Streams, see [AWS Documentation](#).

Enabling Worker Deployment in Production Account

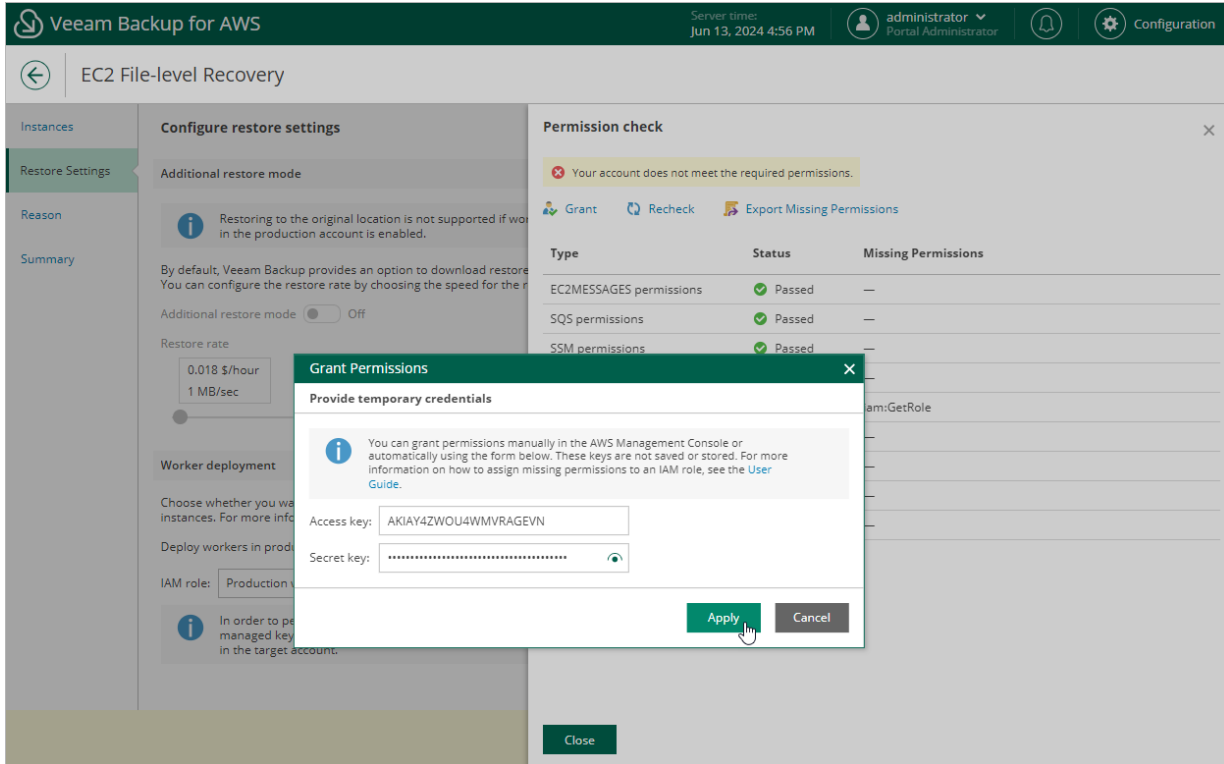
[This option applies only if you have selected a restore point of the **Snapshot, Replica** or **Manual Snapshot** type at the **Restore Point** step of the wizard]

By default, Veeam Backup for AWS deploys worker instances used to perform restore operations in the [backup account](#). However, you can instruct Veeam Backup for AWS to deploy worker instances in a production account – that is, an account in which the snapshot that is used to restore files and folders of the source EC2 instance resides. To do that, set the **Deploy workers in production account** toggle to *On*, and specify an IAM role that will be used to deploy worker instances, and further attached to these instances and used by Veeam Backup for AWS to communicate with them. The role must be assigned permissions listed in section [FLR Worker IAM Role Permissions](#).

For an IAM role to be displayed in the **IAM role** list, it must belong to the AWS account in which the snapshot resides, and must be added to Veeam Backup for AWS as described in section [Adding IAM Roles](#). If you have not added the necessary IAM role to Veeam Backup for AWS beforehand, you can do it without closing the **EC2 File-level Recovery** wizard. To do that, click **Add** and complete the [Add IAM Role](#) wizard.

IMPORTANT

It is recommended that you check whether the selected IAM role has all the permissions required to perform the operation. If some permissions of the IAM role are missing, the restore operation will fail to complete successfully. To run the IAM role permission check, click **Check Permissions** and follow the instructions provided in section [Checking IAM Role Permissions](#).



Step 4. Specify Restore Reason

At the **Reason** step of the wizard, you can specify a reason for recovering files and folders. This information will be saved to the session history and you will be able to reference it later.

The screenshot shows the Veeam Backup for AWS interface. At the top, the header includes the Veeam logo, the text 'Veeam Backup for AWS', the server time 'Jun 13, 2024 4:56 PM', and the user 'administrator Portal Administrator'. Below the header, the breadcrumb 'EC2 File-level Recovery' is visible. On the left, a navigation pane contains 'Instances', 'Restore Settings', 'Reason' (highlighted), and 'Summary'. The main area is titled 'Restore reason' and contains a text input field with the text 'restoring corrupted files'. At the bottom right, there are three buttons: 'Previous', 'Next', and 'Cancel'.

Step 5. Start Recovery Session

At the **Summary** step of the wizard, review summary information and click **Finish**.

As soon as you click **Finish**, Veeam Backup for AWS will close the **File-level Recovery** wizard, start a recovery session and display the **FLR Running Sessions** window. During the recovery session, Veeam Backup for AWS will deploy a worker instance and attach EBS volumes of the processed EC2 instance to it.

TIP

If you accidentally close the **FLR Running Sessions** window, navigate to **Protected Data > EC2** and click the link in the **File-Level Recovery URL** column to open the window again.

In the **FLR Running Sessions** window you can track the progress of the recovery session. In the **URL** column of the window, Veeam Backup for AWS will display a link to the file-level recovery browser. You can use the link in either of the following ways:


- Click the link to open the file-level recovery browser on your local machine while the recovery session is running.
- Copy the link, close the **FLR Running Sessions** window and open the file-level recovery browser on another machine.

IMPORTANT

When you click **Copy URL**, Veeam Backup for AWS copies the following information to the clipboard:

- A link to the file-level recovery browser includes a public DNS name of the worker instance hosting the browser and authentication information used to access the browser.
- A thumbprint of a TLS certificate installed on the worker instance hosting the file-level recovery browser.

To avoid a man-in-the-middle attack, before you start recovering files and folders, check that the certificate thumbprint displayed in the web browser from which you access the file-level recovery browser matches the provided certificate thumbprint.



Step 6. Choose Items to Recover

In the file-level recovery browser, you can find and recover items (files and folders) of the selected EC2 instance. All recovered items are either saved as a single .ZIP archive to the default download directory on a local machine from which you access the browser, or restored to the original EC2 instance.

To recover files and folders from a specific folder, follow the steps:

1. On the **Browse** tab, specify files and folders that you want to recover:
 - a. Navigate to the folder that contains the files and folders.
 - b. In the working area, select check boxes next to the necessary items and click **Add to Restore List**.
2. Switch to the **Restore List** tab, review the list of files and folders, select check boxes next to the items that you want to recover and do the following:
 - To download the selected files and folders to the local machine, click **Download**.
 - To download the selected files and folders to the source EC2 instance, click **Restore > Keep**.
Veeam Backup for AWS will save the files with the `restored-` prefix to the same directory where the source files are located.
 - To restore the selected files and folders to the source EC2 instance, click **Restore > Overwrite**.
Veeam Backup for AWS will overwrite the source files.

As soon as you click **Restore** or **Download**, Veeam Backup for AWS will recover the selected files. You can track the progress and view the results of the restore operation in the **Session Log** section of the **Restore List** tab.

Restore List: amroz-vm03

Restore Status: All ✓ ⚠ ✗

Download Stop Remove

<input type="checkbox"/>	Name ↑	Location	Size	Last Modified	Restore Point	Restore Date	Restore Status
Selected: 2 of 3							
<input checked="" type="checkbox"/>	dev	/		9/14/2022 9:03:44 PM	10/19/2023 10:00:15 AM	11/21/2023 2:42:33 PM	Restoring...
<input checked="" type="checkbox"/>	home	/		10/14/2022 2:59:01 PM	10/19/2023 10:00:15 AM	11/21/2023 2:42:33 PM	Queued
<input type="checkbox"/>	proc	/		9/14/2022 8:59:38 PM	10/19/2023 10:00:15 AM	—	—

Session Log

Status: All ✓ ⚠ ✗

Action	Status	Start Time	End Time	Duration
Select a single item to view sessions details				

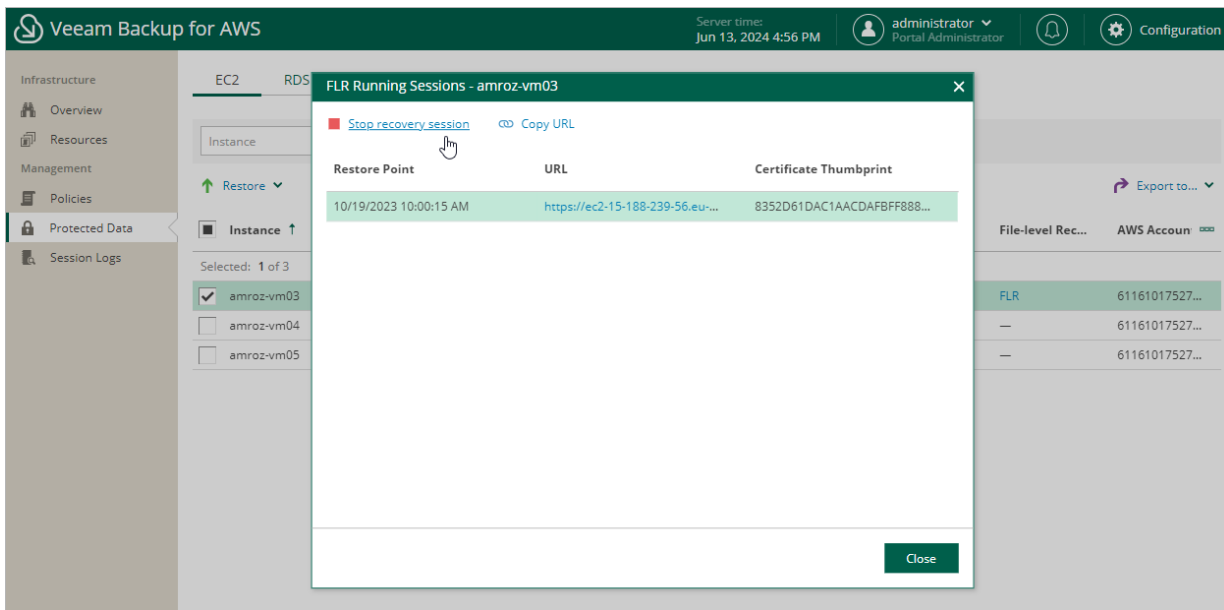
Step 7. Stop Recovery Session

After you finish working with the file-level recovery browser, it is recommended that you stop the recovery session so that Veeam Backup for AWS can unmount and detach EBS volumes of the processed EC2 instance from the worker instance and remove the worker instance from Amazon EC2.

To stop the recovery session, click **Stop recovery session** in the **FLR Running Sessions** window. If you do not perform any actions in the file-level recovery browser for 30 minutes, Veeam Backup for AWS will stop the recovery session automatically.

TIP

If you accidentally close the **FLR Running Sessions** window, navigate to **Protected Data > EC2** and click the link in the **File-Level Recovery URL** column to open the window again.



RDS Restore

The actions that you can perform with restore points of RDS resources depend on whether you access the restore points using the Veeam Backup & Replication console or the Veeam Backup for AWS Web UI.

RDS Restore Using Console

Veeam Backup & Replication offers the following restore operations:

- [DB instance restore](#) – start an entire DB instance from a restore point.
- [Aurora DB clusters restore](#) – start an entire Aurora DB cluster from a restore point.
- [Database restore](#) – restore specific databases of a PostgreSQL DB instance.

You can restore RDS resource data to the most recent state or to any available restore point.

Restoring DB Instances

To restore a DB instance, do the following:

1. [Launch the Restore to Amazon RDS wizard.](#)
2. [Select a restore point.](#)
3. [Choose a restore mode.](#)
4. [Select an AWS Region.](#)
5. [Specify instance type and enable encryption.](#)
6. [Specify parameter and option groups.](#)
7. [Specify a database identifier.](#)
8. [Configure network settings.](#)
9. [Specify a restore reason.](#)
10. [Finish working with the wizard.](#)

Step 1. Launch Restore to Amazon RDS Wizard

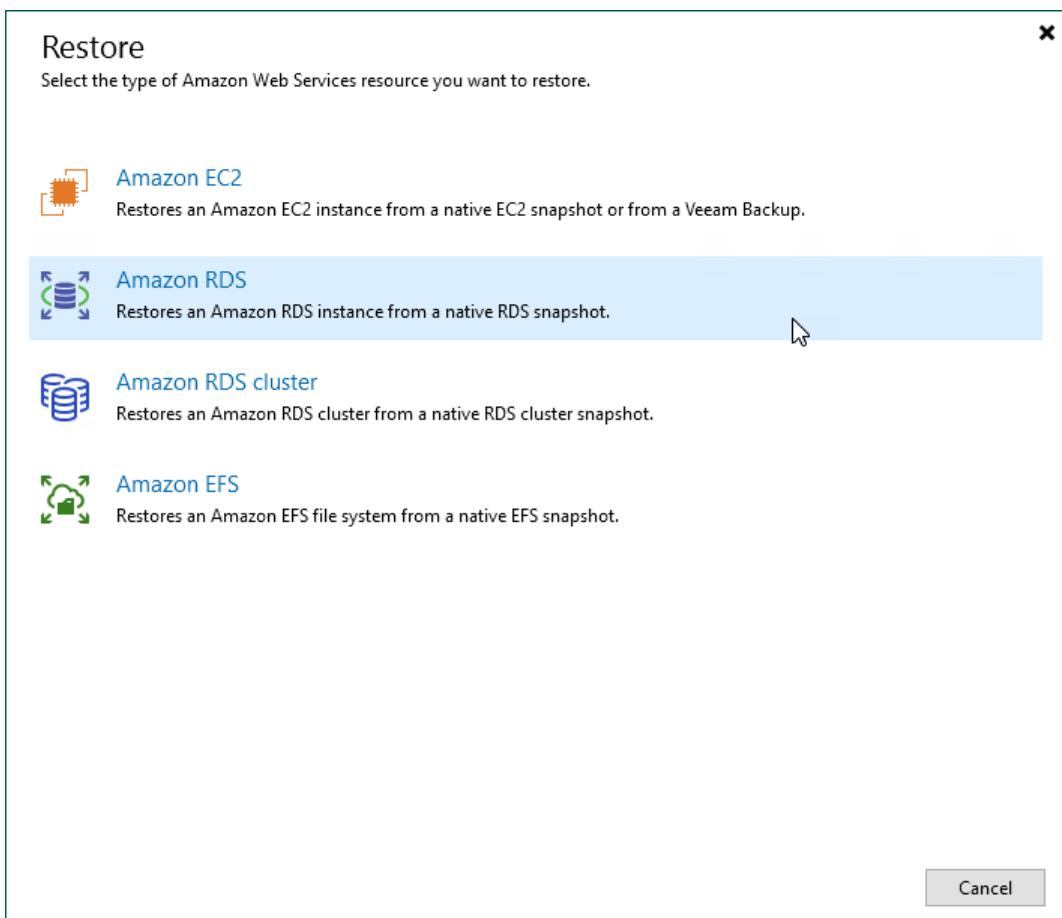
To launch the **Restore to Amazon RDS** wizard, do the following:

1. In the Veeam Backup & Replication console, open the **Home** view.
2. Navigate to **Backups > Snapshots**.
3. In the working area, expand the backup policy that protects a DB instance that you want to restore, select the necessary instance and click **Amazon RDS** on the ribbon.

Alternatively, you can right-click the instance and select **Amazon RDS**.

TIP

You can also launch the **Restore to Amazon RDS** wizard from the **Home** tab. To do that, click **Restore** and select **AWS**. Then, select **Amazon RDS** in the **Restore** window.



Step 2. Select Restore Point

At the **RDS Instance** step of the wizard, choose a restore point that will be used to restore the selected DB instance. By default, Veeam Backup & Replication uses the most recent valid restore point. However, you can restore the instance data to an earlier state.

To select a restore point, do the following:


1. In the **RDS instance** list, select the DB instance and click **Point**.
2. In the **Restore Points** window, expand the backup policy that protects the DB instance, select the necessary restore point and click **OK**.

To help you choose a restore point, Veeam Backup & Replication provides the following information on each available restore point:

- **Job** – the name of the backup policy that created the restore point and the date when the restore point was created.
- **Type** – the type of the restore point.
- **Location** – the AWS Region where the restore point is stored.

TIP

You can use the wizard to restore multiple instances at a time. To do that, click **Add**, select more DB instances to restore and choose a restore point for each of them.



Step 3. Choose Restore Mode

At the **Restore Mode** step of the wizard, do the following:

1. Choose whether you want to restore the selected DB instances to the original or to a new location.

NOTE

Restore to the original location is not supported if the IAM role that will be used to perform the restore operation belongs to an AWS account that differs from the AWS account to which the source resources belong.

2. Click **Pick an IAM identity to use** to select an IAM identity whose permissions will be used to perform the restore operation:
 - To specify an IAM role, select the **IAM role** option and choose the necessary IAM role from the **IAM role** drop-down list.

For an IAM role to be displayed in the list of available roles, it must be added to the backup appliance as described in section [Adding IAM Roles](#).
 - To specify access keys of an IAM user, select the **Temporary access key** option, and use the **Access key** and **Secret key** fields to provide the access key ID and the secret access key.

NOTE

By default, to perform the restore operation, Veeam Backup & Replication uses permissions of either the IAM role that was used to protect the source EC2 instance, or the IAM role used to update information on restore points that were created for the instance while rescanning AWS infrastructure.

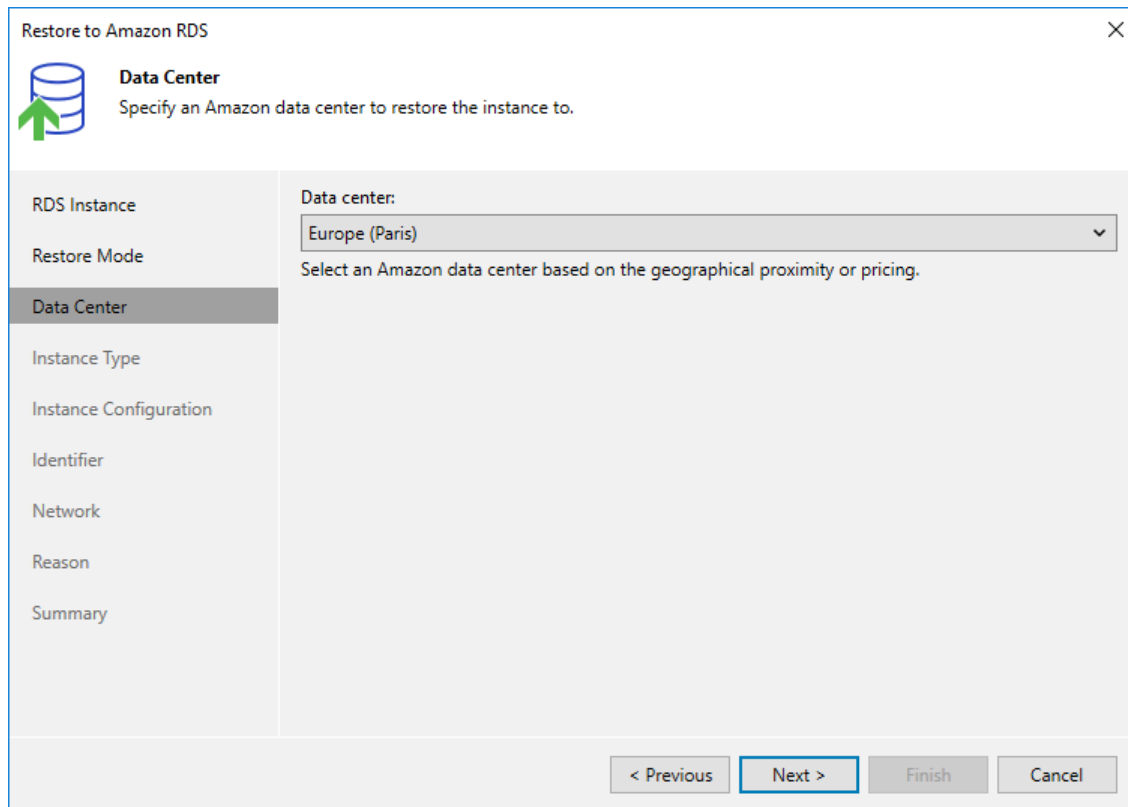
The screenshot shows a wizard window titled "Restore to Amazon RDS" with a close button (X) in the top right corner. On the left is a navigation pane with a tree view containing: RDS Instance, Restore Mode (highlighted), Data Center, Instance Type, Instance Configuration, Identifier, Network, Reason, and Summary. An icon of a database cylinder with an upward arrow is next to the "Restore Mode" header. Below the header, the text reads: "Specify whether selected RDS instances should be restored back to the original location, or to a new location or with different settings." The main area contains two radio button options: "Restore to the original location" (unselected) and "Restore to a new location, or with different settings" (selected). The selected option has a description: "Customize the restored RDS instance location, and change its settings. The wizard will automatically populate all controls with the original RDS instance settings as the defaults." Below this is a blue link: "Pick an IAM identity to use". At the bottom of the window are four buttons: "< Previous" (disabled), "Next >" (active/highlighted), "Finish" (disabled), and "Cancel" (disabled).

Step 4. Select Region

[This step applies only if you have selected the **Restore to a new location, or with different settings** option at the **Restore Mode** step of the wizard]

At the **Data Center** step of the wizard, select an AWS Region where the restored DB instance will operate.

If the selected location differs from the original location of the DB instance, Veeam Backup & Replication will raise a warning message notifying that the locations do not match. Click **Yes** to acknowledge the warning. Otherwise, you will not be able to proceed with the wizard.



The screenshot shows the 'Restore to Amazon RDS' wizard window. The title bar reads 'Restore to Amazon RDS' with a close button (X) on the right. Below the title bar is a header section with a database icon and a green arrow pointing up, followed by the text 'Data Center' and 'Specify an Amazon data center to restore the instance to.' Below this is a list of wizard steps on the left: 'RDS Instance', 'Restore Mode', 'Data Center' (highlighted), 'Instance Type', 'Instance Configuration', 'Identifier', 'Network', 'Reason', and 'Summary'. The main area contains a 'Data center:' label above a dropdown menu showing 'Europe (Paris)'. Below the dropdown is the instruction 'Select an Amazon data center based on the geographical proximity or pricing.' At the bottom of the window are four buttons: '< Previous', 'Next >' (highlighted with a blue border), 'Finish', and 'Cancel'.

Step 5. Specify Instance Type and Enable Encryption

[This step applies only if you have selected the **Restore to a new location, or with different settings** option at the **Restore Mode** step of the wizard]

At the **Instance Type** step of the wizard, you can configure settings for the restored DB instance. To do that, select the instance and do the following:

- If you want to specify a new machine type for the restored DB instance, click **Type** and select the necessary type in the **Instance Type** window. For the list of all existing RDS instance types, see [AWS Documentation](#).

You can also choose a new disk storage type for the restored DB instance. For more information on RDS storage types, see [AWS Documentation](#).

- If you want to change the encryption settings of the restored DB instance, click **Encryption** and do the following in the **Disk Encryption** window:
 - Select the **Preserve the original encryption settings** option if you do not want to encrypt the DB instance or want to apply the original encryption scheme of the source DB instance.

NOTE

You will not be able to select the **Preserve the original encryption settings** option if the AWS KMS key used to encrypt the source DB instance is not available in the region to which the DB instance will be restored.


- Select the **Use the following encryption key** option if you want to encrypt the DB instance with an AWS KMS key. Then, choose the necessary KMS key from the list.

For a KMS key to be displayed in the list of available encryption keys, it must be stored in the AWS Region selected at [step 4](#) of the wizard, and the IAM role specified for the restore operation must have permissions to access the key. For more information on KMS keys, see [AWS Documentation](#).

TIP

If the necessary KMS key is not displayed in the list, or if you want to use a KMS key from an AWS account other than the AWS account to which the specified IAM role belongs, you can specify the Amazon Resource Number (ARN) of the key in the **Use the following encryption key** field.

For Veeam Backup for AWS to be able to encrypt the restored DB instance using the provided KMS key, either the IAM role or user specified for the restore operation, or the IAM role used to create the restore point selected at [step 2](#) of the wizard must have permissions to access the key.



Step 6. Specify Parameter and Option Groups

[This step applies only if you have selected the **Restore to a new location, or with different settings** option at the **Restore Mode** step of the wizard]

At the **Instance Configuration** step of the wizard, you can choose the parameter and option groups with which the restored DB instance will be associated. To do that, select the instance and click **Edit**. In the **Group** window, do the following:

1. From the **Parameter group** drop-downlist, select the parameter group containing database engine configuration values that will be applied to the restored DB instance.


For a parameter group to be displayed in the list of available groups, the group must be created in AWS as described in [AWS Documentation](#), and the group settings must be compatible with the database engine and version of the original DB instance.

2. From the **Option group** drop-downlist, select the option group containing database configuration values and security settings that will be applied to the restored DB instance.

For an option group to be displayed in the list of available groups, it must be created in AWS as described in [AWS Documentation](#), and the group settings must be compatible with the database engine and version of the original DB instance.

NOTE

If Veeam Backup for AWS fails to find any option or parameter groups compatible with the database engine and version of the original DB instance, the **default** option will be selected automatically. In this case, Veeam Backup & Replication will create the necessary group during the restore session and associate the restored DB instance with the group.



Step 7. Specify Database Identifier

[This step applies only if you have selected the **Restore to a new location, or with different settings** option at the **Restore Mode** step of the wizard]

At the **Identifier** step of the wizard, you can specify a new identifier for the restored DB instance.

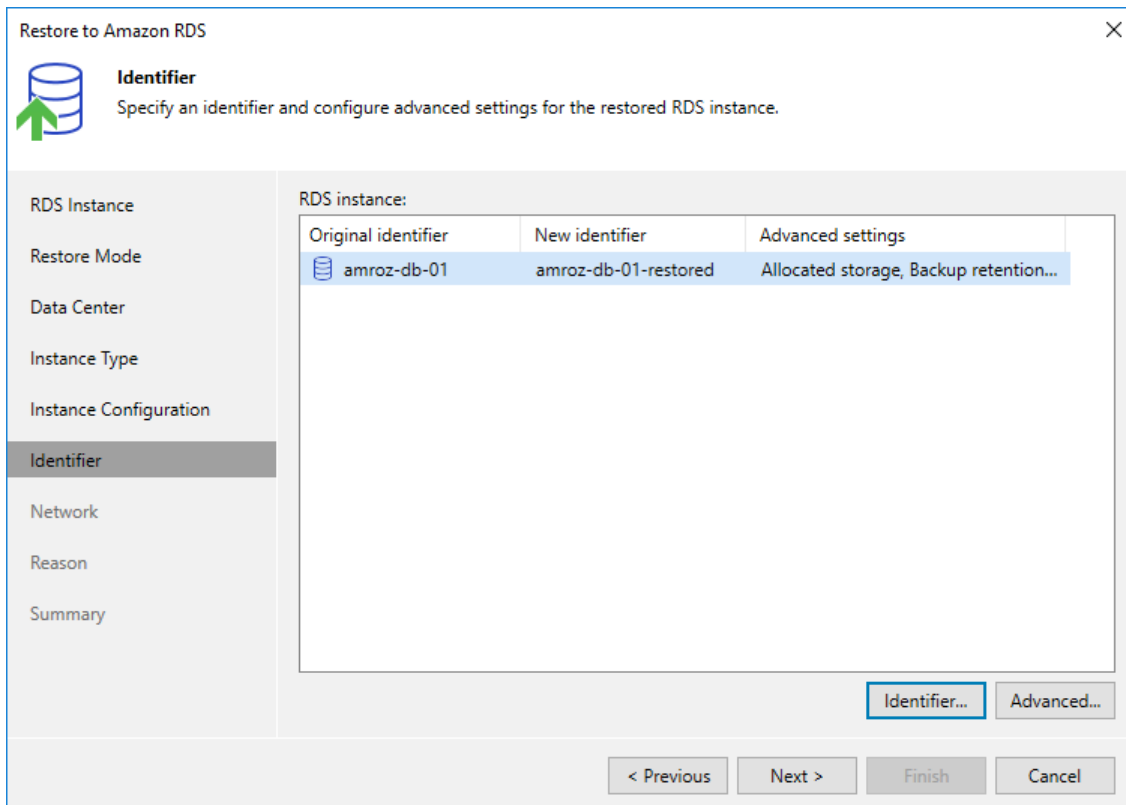
Consider the following limitations:

- The instance identifier must be unique for each AWS Region within one AWS Account.
- The instance identifier can contain only lowercase Latin letters and hyphens, but cannot contain two consecutive hyphens.
- The first character of the instance identifier must be a letter. The last character of the identifier must not be a hyphen.
- The maximum length of the instance identifier is 63 characters.

For more information on limitations for DB instance identifiers, see [AWS Documentation](#).

TIP

The **Identifier** step of the wizard contains preconfigured settings retrieved from the source DB instance. If you want to specify advanced configuration settings for the restored DB instance, click **Advanced** and edit the necessary settings in the **Advanced Settings** window. For more information on all available settings that can be specified for DB instances, see [AWS Documentation](#).



Step 8. Configure Network and Availability Settings

[This step applies only if you have selected the **Restore to a new location, or with different settings** option at the **Restore Mode** step of the wizard]

At the **Network** step of the wizard, you can configure specific network settings for the restored DB instance. To do that, select the instance and do the following:

1. Click **Customize**. Then, in the **Amazon VPC** window:
 - a. From the **Amazon VPC**, **Subnet group** and **Security group** drop-down lists, select an Amazon VPC to which the instance will be connected, a subnet group in which the instance will be launched, and a security group that will be associated with the instance. Note that the **Amazon VPC** list shows only VPCs that include one or more subnet groups.

For an Amazon VPC, subnet group and security group to be displayed in the list of available network specifications, they must be created in AWS in the AWS Region specified at [step 4](#) of the wizard, as described in [AWS Documentation](#).

- b. In the **Database port** field, specify the number of a port that will be used to access the DB instance. The port number must be within the following range: 1150–65535.

For SQL database engines, do not use the following port numbers: 1234, 1434, 3260, 3343, 3389, 47001 and 49152–49156.

2. Click **Availability**. Then, in the **Availability Settings** window:
 - a. From the **Public access** drop-down list, select *Enabled* if you want to make the restored DB instance accessible outside the selected Amazon VPC. Note that the DB instance must belong to a public subnet group to become publicly accessible.
 - b. From the **Availability type** drop-down list, select *Multiple zone* if you want to create a passive secondary replica (standby instance) of the restored DB instance. Note that Multi-AZ deployments are not supported for instances running MS SQL Server Express and MS SQL Server Web editions.

For more information on the Multi-AZ deployment, see [AWS Documentation](#).

- c. [This step applies only if you have selected the **Single zone** option] From the **Availability zone** drop-down list, select an Availability Zone where the restored DB instance will reside.

The screenshot shows the 'Restore to Amazon RDS' wizard in the 'Network' step. The main window has a sidebar with options: RDS Instance, Restore Mode, Data Center, Instance Type, Instance Configuration, Identifier, Network (selected), Reason, and Summary. The main area contains a 'Network' section with a description: 'Specify the virtual private cloud and additional network settings for the restored RDS instance.' An 'Amazon VPC' dialog box is open, with the following fields:

- Amazon VPC:** vpc-0d6107b77a93eb9a1
- Subnet group:** default-vpc-0d6107b77a93eb9a1
- Security group:** amroz-sec
- Database port:** 5432

The 'Availability zone' dropdown is set to 'eu-west-3c'. At the bottom of the dialog are 'OK' and 'Cancel' buttons. Below the dialog, there are 'Customize...' and 'Availability...' buttons. At the bottom of the main window are '< Previous', 'Next >', 'Finish', and 'Cancel' buttons.

Step 9. Specify Restore Reason

At the **Reason** step of the wizard, specify a reason for restoring the Amazon DB instance. The information you provide will be saved in the session history and you can reference it later.

Restore to Amazon RDS

Reason
Type in the reason for performing this restore operation. This information will be logged in the restore sessions history for later reference.

RDS Instance
Restore Mode
Data Center
Instance Type
Instance Configuration
Identifier
Network
Reason
Summary

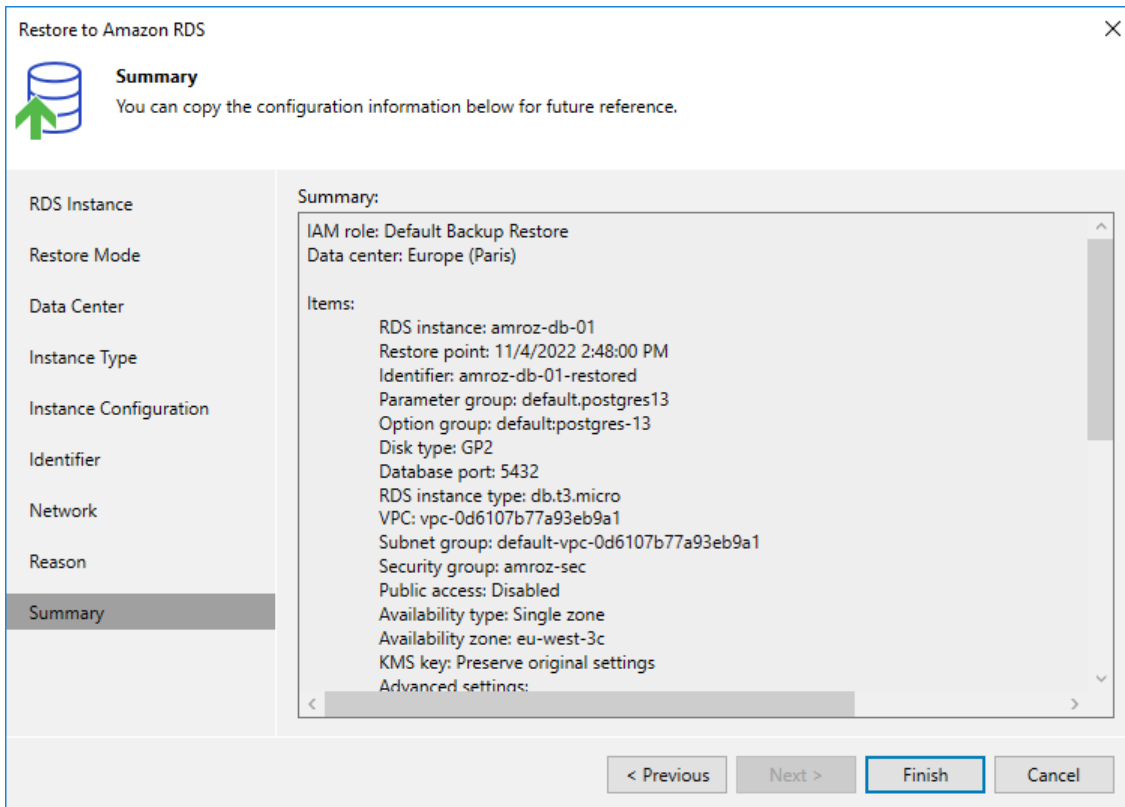
Restore reason:
Restore failed RDS instances

Do not show me this page again

< Previous Next > Finish Cancel

Step 10. Finish Working with Wizard

At the **Summary** step of the wizard, review summary information and click **Finish**.



Restoring Aurora DB Clusters

To restore a cluster, do the following:

1. [Check prerequisites and limitations.](#)
2. [Launch the Restore Amazon RDS Cluster wizard.](#)
3. [Select a restore point.](#)
4. [Choose a restore mode.](#)
5. [Select an AWS Region.](#)
6. [Choose capacity type and enable encryption.](#)
7. [Specify cluster and instance parameter groups.](#)
8. [Specify cluster and database identifiers.](#)
9. [Configure network and availability settings.](#)
10. [Specify a restore reason.](#)
11. [Finish working with the wizard.](#)

Limitations and Considerations

When restoring Aurora DB clusters, keep in mind the following limitations and considerations.

IAM Roles and Users

An IAM role and IAM user that you plan to use to perform the restore operation must have permissions described in section [RDS Restore IAM Permissions](#).

Restore Mode

Before you choose the restore mode, consider the following limitations:

- Restore of Aurora DB clusters to the original location is not supported if the [IAM role specified](#) for the restore operation belongs to an AWS account that differs from the AWS account to which the source cluster belongs.
- Restore of Aurora multi-master clusters is not supported if the source region differs from the target region specified for the restore operation. However, you can restore these clusters to the source region in the same or in the another AWS account. To specify an AWS account to which the cluster will be restored, select an IAM role that belongs to the necessary account at [step 3](#) of the **Restore Amazon RDS Cluster** wizard.
- When restoring to a new location, Veeam Backup & Replication creates only the primary DB instances in the restored clusters. Additional writer DB instances (for Aurora multi-master clusters) and Aurora Replicas (for Aurora DB clusters with single-master replication) must be added manually in the AWS after the restore operation completes.

To learn how to add DB instances to Amazon Aurora DB clusters, see [AWS Documentation](#).

- When restoring Aurora global databases, Veeam Backup & Replication restores only primary Aurora DB clusters in the primary AWS Regions; secondary clusters must be created manually in the AWS after the restore operation completes. If source clusters are still present in AWS, primary DB clusters will be restored with the `veeam-temp-<cluster_name>-<guid>` name pattern; the source clusters will not be removed automatically.

For more information on Amazon Aurora global databases, see [AWS Documentation](#).

Capacity Types

Before you choose a capacity type for the restored cluster, consider the following limitations:

- You can restore an Aurora Serverless DB cluster either as an Aurora Serverless DB cluster or as an Aurora provisioned DB cluster. However, you cannot restore an Aurora provisioned DB cluster as an Aurora Serverless DB cluster unless the source cluster is running MySQL 2.11.4.
- Aurora Serverless v1 is supported for a limited list of AWS Regions and only for MySQL 2.11.4. For more information, see [AWS Documentation](#).

Step 1. Launch Restore Amazon RDS Cluster Wizard


To launch the **Restore to Amazon RDS cluster** wizard, do the following:

1. In the Veeam Backup & Replication console, open the **Home** view.
2. Navigate to **Backups > Snapshots**.
3. In the working area, expand the backup policy that protects an Aurora DB cluster that you want to restore, select the necessary cluster and click **Amazon RDS cluster** on the ribbon.

Alternatively, you can right-click the instance and select **Restore to Amazon RDS cluster**.

TIP

You can also launch the **Restore to Amazon RDS cluster** wizard from the **Home** tab. To do that, click **Restore** and select **AWS**. Then, select **Amazon RDS cluster** in the **Restore** window.



Step 2. Select Restore Point

At the **RDS Cluster** step of the wizard, choose a restore point that will be used to restore the selected Aurora DB cluster. By default, Veeam Backup & Replication uses the most recent valid restore point. However, you can restore the cluster data to an earlier state.

To select a restore point, do the following:

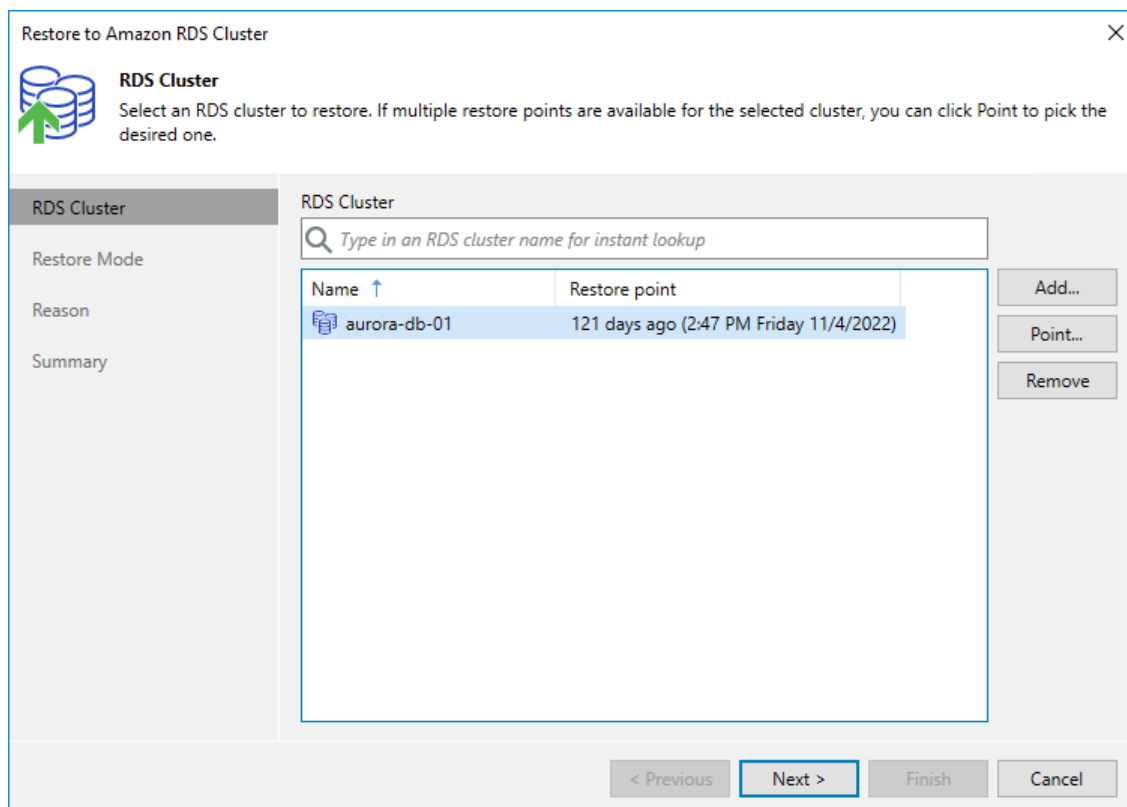
1. In the **RDS cluster** list, select the Aurora DB cluster and click **Point**.
2. In the **Restore Points** window, expand the backup policy that protects the cluster, select the necessary restore point and click **OK**.

To help you choose a restore point, Veeam Backup & Replication provides the following information on each available restore point:

- **Job** – the name of the backup policy that created the restore point and the date when the restore point was created.
- **Type** – the type of the restore point.
- **Location** – the AWS Region where the restore point is stored.

TIP

You can use the wizard to restore multiple clusters at a time. To do that, click **Add**, select more clusters to restore and choose a restore point for each of them.



Step 3. Choose Restore Mode

At the **Restore Mode** step of the wizard, do the following:

IMPORTANT

Before choosing a restore mode, check the limitations and prerequisites described in section [Limitations and Considerations](#).

1. Choose whether you want to restore the Aurora DB cluster to the original or to a new location.
2. Click **Pick an IAM identity to use** to select an IAM identity whose permissions will be used to perform the restore operation:
 - To specify an IAM role for the restore operation, select the **IAM role** option and choose the necessary IAM role from the **IAM role** drop-down list.

For an IAM role to be displayed in the list of available roles, it must be added to the backup appliance as described in section [Adding IAM Roles](#).
 - To specify access keys of an IAM user, select the **Temporary access key** option and use the **Access key** and **Secret key** fields to provide the access key ID and the secret access key.

NOTE

By default, to perform the restore operation, Veeam Backup & Replication uses permissions of either the IAM role that was used to protect the source EC2 instance, or the IAM role used to update information on restore points created for the instance while rescanning AWS infrastructure.

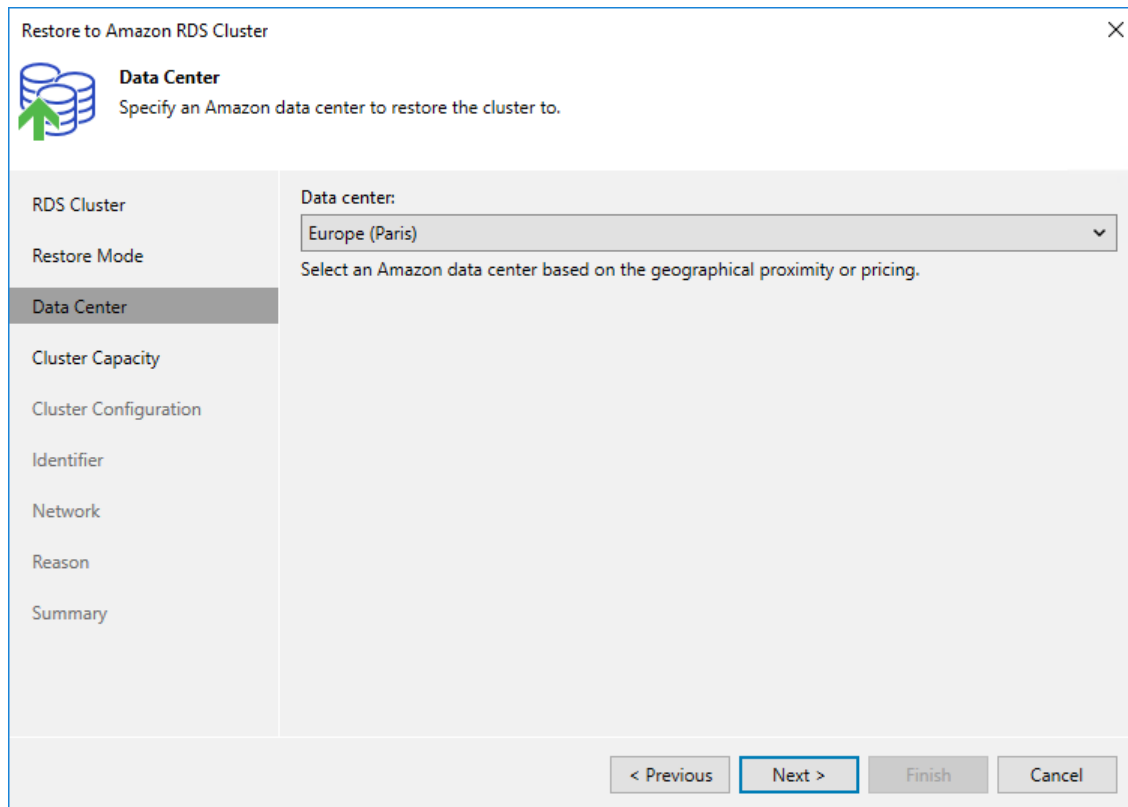
The screenshot shows a wizard window titled "Restore to Amazon RDS Cluster" with a close button (X) in the top right corner. The main content area is titled "Restore Mode" and includes a sub-header "Specify whether selected RDS clusters should be restored back to the original location, or to a new location or with different settings." Below this, there are two radio button options: "Restore to the original location" (unselected) and "Restore to a new location, or with different settings" (selected). The second option includes a link "Pick an IAM identity to use". On the left side, there is a vertical navigation pane with the following items: "RDS Cluster", "Restore Mode" (highlighted), "Data Center", "Cluster Capacity", "Cluster Configuration", "Identifier", "Network", "Reason", and "Summary". At the bottom of the window, there are four buttons: "< Previous", "Next >" (highlighted), "Finish", and "Cancel".

Step 4. Select Region

[This step applies only if you have selected the **Restore to a new location, or with different settings** option at the **Restore Mode** step of the wizard]

At the **Data Center** step of the wizard, select an AWS Region where the restored Aurora DB cluster will operate.

If the selected location differs from the original location of the Aurora DB cluster, Veeam Backup & Replication will raise a warning notifying that the locations do not match. Click **Yes** to acknowledge the warning. Otherwise, you will not be able to proceed with the wizard.



The screenshot shows a wizard window titled "Restore to Amazon RDS Cluster" with a close button (X) in the top right corner. The main heading is "Data Center" with a sub-heading "Specify an Amazon data center to restore the cluster to." and an icon of a database with an upward arrow. On the left is a vertical navigation pane with the following items: "RDS Cluster", "Restore Mode", "Data Center" (highlighted), "Cluster Capacity", "Cluster Configuration", "Identifier", "Network", "Reason", and "Summary". The main area contains a "Data center:" label above a dropdown menu showing "Europe (Paris)". Below the dropdown is the instruction "Select an Amazon data center based on the geographical proximity or pricing." At the bottom of the window are four buttons: "< Previous", "Next >" (highlighted with a blue border), "Finish", and "Cancel".

Step 5. Choose Capacity Type and Enable Encryption

[This step applies only if you have selected the **Restore to a new location, or with different settings** option at the **Restore Mode** step of the wizard]

At the **Cluster Capacity** step of the wizard, you can configure capacity and encryption settings for the restored Aurora DB cluster:

IMPORTANT

Before configuring capacity settings, check the limitations and prerequisites described in section [Limitations and Considerations](#).

1. Click **Capacity**. Then, in the **Capacity Settings** window:
 - Select the **Provision cluster of the specified instance type** option if you want perform to restore to Aurora provisioned. Then, choose the DB instance class that will be used to create the primary DB instance in the restored cluster.

For a DB instance class to be displayed in the list, it must be supported for the Aurora DB engine of the source Aurora DB cluster. For more information on supported DB instance classes, see [AWS Documentation](#).
 - Select the **Minimum and maximum amount of resources** option if you want to perform restore to Aurora Serverless. Then, specify a range of capacity units that will be used to create scaling rules for the restored cluster. These rules will define thresholds for CPU utilization, connections and available memory.

For more information on capacity units and scaling rules, see [AWS Documentation](#).

TIP

To restore the primary DB instance of a cluster as an Aurora Serverless v2 DB instance, select the **Provision cluster of the specified instance type** option, and choose the **db.serverless** instance type.

However, consider the following limitations:

- Aurora Serverless v2 is supported only for a limited list of DB engine versions. For more information, see [AWS Documentation](#).
 - You cannot specify a capacity range for the restored Aurora Serverless v2 DB instance. If the source DB instance had the same instance class as the restored instance, Veeam Backup & Replication will restore the instance with the backed-up capacity range. Otherwise, Veeam Backup & Replication will restore the Aurora Serverless v2 DB instance with the default capacity range – 8-64 Aurora Capacity Units.
- Select an Aurora database engine version for the restored cluster from the **Database engine version** drop-down list. The list shows only DB engine versions supported in the target AWS Region, and is filtered based on the DB engine type and DB engine version of the source Aurora DB cluster.

For more information on Amazon Aurora database engine versions, see [AWS Documentation](#).

IMPORTANT

Consider the following:

- When restoring Amazon Aurora global databases, make sure you select an Aurora database version that supports the global database feature. For the list of supported Aurora database versions, see [AWS Documentation](#).
- To be able to use the Aurora MySQL parallel query feature when restoring a cluster, make sure you select an Aurora database version that supports the parallel query feature. Keep in mind that to use this feature, you must also enable the `aurora_parallel_query` parameter in the DB cluster parameter group that you will specify at [step 6](#) of the wizard.

For more information on Aurora MySQL parallel query, see [AWS Documentation](#).

2. Click **Encryption**. Then, in the **Disk encryption** window:

- Select the **Preserve the original encryption settings** option if you do not want to encrypt the restored cluster or want to apply the original encryption scheme of the source cluster.

If you plan to restore an unencrypted provisioned DB cluster to Aurora Serverless and want to preserve the original encryption settings, note that Veeam Backup & Replication will encrypt the newly created Aurora Serverless DB cluster with the default KMS key in the target AWS Region. For more information on Aurora Serverless, see [AWS Documentation](#).

- Select the **Use the following encryption password** option if you want to encrypt the restored cluster with an AWS KMS key. Then, choose the necessary KMS key from the list.

For a KMS key to be displayed in the list of available encryption keys, it must be stored in AWS Region select at [step 4](#) of the wizard, and the IAM role specified for the restore operation must have permissions to access the key. For more information on KMS keys, see [AWS Documentation](#).

TIP

If the necessary KMS key is not displayed in the list, or if you want to use a KMS key from an AWS account other than the AWS account to which the specified IAM role belongs, you can specify the Amazon Resource Number (ARN) of the key in the **Use the following encryption key** field.

For Veeam Backup for AWS to be able to encrypt the restored Aurora DB cluster using the provided KMS key, either the IAM role or user specified for the restore operation, or the IAM role used to create the restore point selected at [step 2](#) of the wizard must have permissions to access the key.

The screenshot shows the 'Restore to Amazon RDS Cluster' wizard, specifically the 'Cluster Capacity' step. The wizard's main window has a sidebar with navigation options: RDS Cluster, Restore Mode, Data Center, Cluster Capacity (selected), Cluster Configuration, Identifier, Network, Reason, and Summary. The main area displays a table of RDS clusters:

RDS cluster:				
Name	Compute	Engine version	Encryption	
aurora-db-01	Provisioned	13.7	Preserve original settings	

A 'Capacity Settings' dialog box is open, allowing the user to select capacity options for the restored cluster. The dialog includes the following options:

- Provision cluster of the specified instance type:
 - db.r6g.2xlarge (8 cores, 64.0 GB memory)
- Minimum and maximum amount of resources:
 - Use between 2 (4 GB RAM) and 64 (122 GB RAM) capacity units
- Database engine version: 13.7

The dialog has 'OK' and 'Cancel' buttons. Below the dialog, the wizard's navigation buttons are visible: '< Previous', 'Next >', 'Finish', and 'Cancel'. An 'Encryption...' button is also present near the bottom right of the dialog.

Step 6. Specify Cluster and Instance Parameter Groups

[This step applies only if you have selected the **Restore to a new location, or with different settings** option at the **Restore Mode** step of the wizard]

At the **Instance Configuration** step of the wizard, you can choose the cluster parameter group that will be associated with the restored cluster, and the parameter group that will be associated with the primary DB instance. To do that, select the cluster and click **Edit**. In the **Group** window, do the following:

1. From the **Cluster parameter group** drop-down list, select the parameter group containing database engine configuration values that will be applied to each DB instance launched in the restored cluster.


For a DB cluster parameter group to be displayed in the list of available groups, the group must be created in AWS as described in [AWS Documentation](#).

2. [This step applies only to provisioned Aurora DB clusters and Aurora Serverless v2 DB clusters] From the **Parameter group** drop-down list, select the DB parameter group containing database engine configuration values that will be applied to the primary DB instance in the restored cluster.

For a DB parameter group to be displayed in the list of available groups, the group must be created in AWS as described in [AWS Documentation](#).

NOTE

If Veeam Backup for AWS fails to find any parameter groups in the target AWS Region, the **default** option will be selected automatically. In this case, Veeam Backup & Replication will create the necessary group during the restore session and associate the restored DB cluster and primary DB instance with the group.



Step 7. Specify Cluster and Database Identifiers

[This step applies only if you have selected the **Restore to a new location, or with different settings** option at the **Restore Mode** step of the wizard]

At the **Identifier** step of the wizard, you can specify a new identifier for the restored Aurora DB cluster and for the primary DB instance.

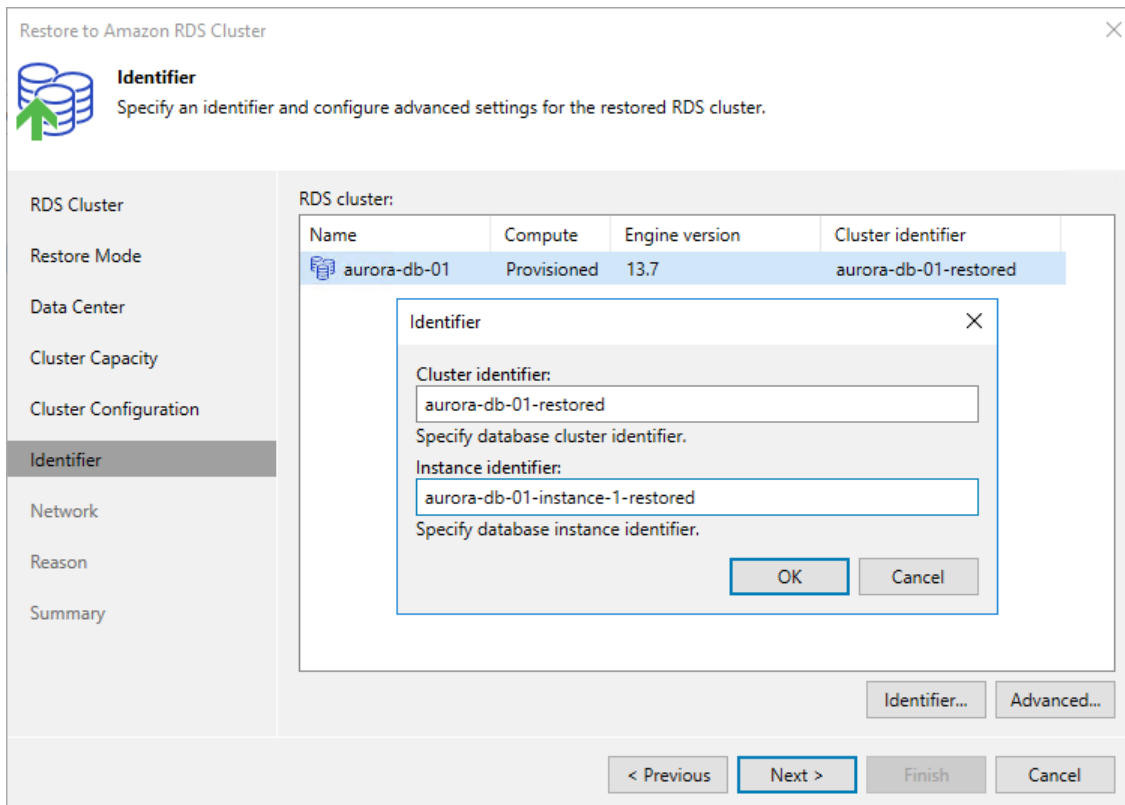
Consider the following limitations:

- The identifier must be unique for each AWS Region within one AWS Account.
- The identifier can contain only lowercase Latin letters and hyphens, but cannot contain two consecutive hyphens.
- The first character of the identifier must be a letter. The last character of the identifier must not be a hyphen.
- The maximum length of the identifier is 63 characters.

For more information on limitations for DB instance identifiers, see [AWS Documentation](#). For more information on limitations for Aurora DB cluster identifiers, see [AWS Documentation](#).

TIP

The **Identifier** step of the wizard contains preconfigured settings retrieved from the source primary DB instance. If you want to specify advanced configuration settings for the restored primary DB instance, click **Advanced** and edit the necessary settings in the **Advanced Settings** window. For more information on all available settings that can be specified for DB instances, see [AWS Documentation](#).



Step 8. Configure Network and Availability Settings

[This step applies only if you have selected the **Restore to a new location, or with different settings** option at the **Restore Mode** step of the wizard]

At the **Network** step of the wizard, you can configure specific network settings for the restored Aurora DB cluster. To do that, select the cluster and do the following:

1. Click **Customize**. Then, in the **Amazon VPC** window:
 - a. From the **Amazon VPC**, **Subnet group** and **Security group** drop-down lists, select an Amazon VPC to which the cluster will be restored, a subnet group in which the cluster will be launched, and a security group that will control access to the restored cluster. Note that the subnet group must include at least 2 subnets created in 2 different Availability Zones of the AWS Region specified at [step 4](#) of the wizard.

For an Amazon VPC, subnet group, security group to be displayed in the list of available network specifications, they must be created in the AWS Region specified at [step 4](#) of the wizard as described in [AWS Documentation](#).

- b. In the **Database port** field, specify the number of a port that will be used to access the primary DB instance.

The port number must be within the following range: 1150–65535.

2. Click **Availability**. Then, in the **Availability Settings** window:

- a. From the **Public access** drop-down list, select *Enabled* if you want to make the restored cluster accessible outside the selected Amazon VPC. Note that the cluster must belong to a public subnet group to become publicly accessible.
- b. From the **Availability type** drop-down list, select an Availability Zone where the primary DB instance will reside.

The screenshot shows the 'Restore to Amazon RDS Cluster' wizard at the 'Network' step. The main window has a sidebar with options: RDS Cluster, Restore Mode, Data Center, Cluster Capacity, Cluster Configuration, Identifier, Network (selected), Reason, and Summary. The main area displays 'Specify the virtual private cloud and additional network settings for the restored RDS cluster.' An 'Amazon VPC' dialog box is open, containing the following fields:

- Amazon VPC:** vpc-0d6107b77a93eb9a1
- Subnet group:** default-vpc-0d6107b77a93eb9a1
- Security group:** amroz-sec
- Database port:** 5432

Buttons for 'OK', 'Cancel', 'Customize...', and 'Availability...' are visible at the bottom of the dialog. The main window also has '< Previous', 'Next >', 'Finish', and 'Cancel' buttons at the bottom.

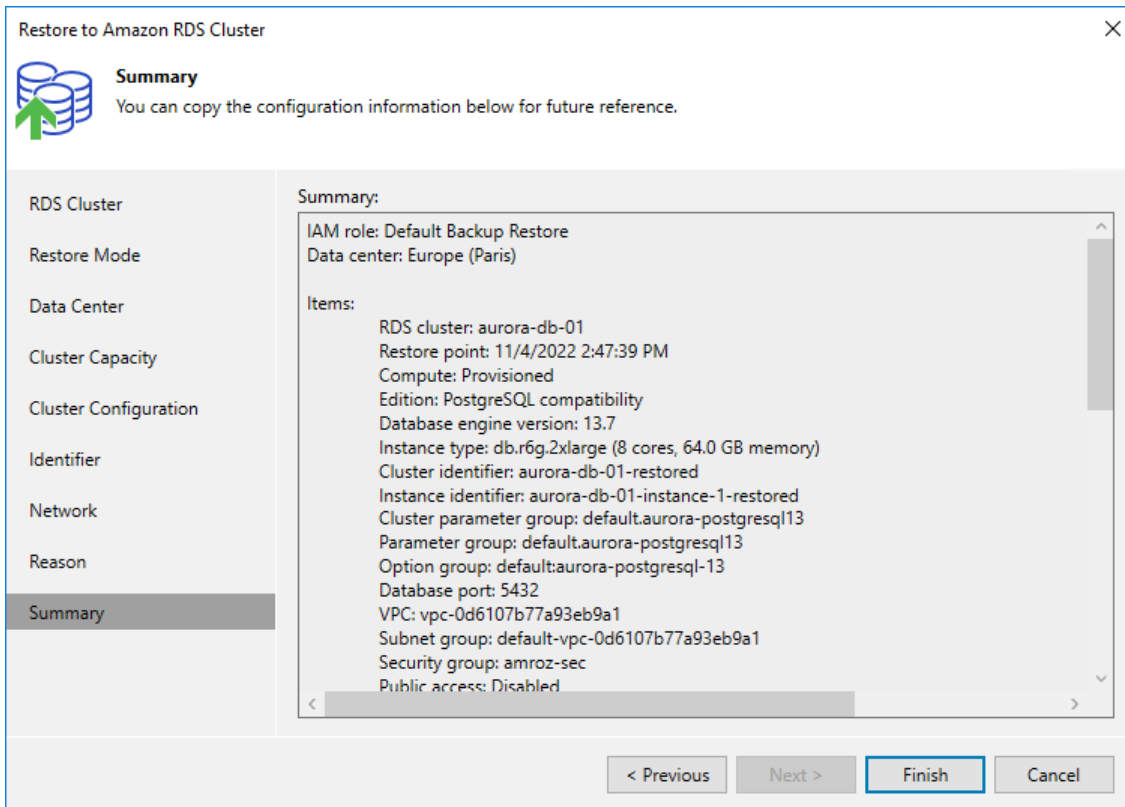
Step 9. Specify Restore Reason

At the **Reason** step of the wizard, specify a reason for restoring the Aurora DB cluster. The information you provide will be saved in the session history and you can reference it later.

The screenshot shows a wizard window titled "Restore to Amazon RDS Cluster" with a close button (X) in the top right corner. The window is divided into a left sidebar and a main content area. The sidebar contains a list of steps: RDS Cluster, Restore Mode, Data Center, Cluster Capacity, Cluster Configuration, Identifier, Network, Reason (highlighted), and Summary. Above the sidebar, there is a database icon with a green arrow pointing up, followed by the heading "Reason" and the instruction: "Type in the reason for performing this restore operation. This information will be logged in the restore sessions history for later reference." The main content area has a text input field labeled "Restore reason:" containing the text "Restore failed clusters". Below the input field is a checkbox labeled "Do not show me this page again". At the bottom of the window, there are four buttons: "< Previous", "Next >", "Finish", and "Cancel".

Step 10. Finish Working with Wizard

At the **Summary** step of the wizard, review summary information and click **Finish**.



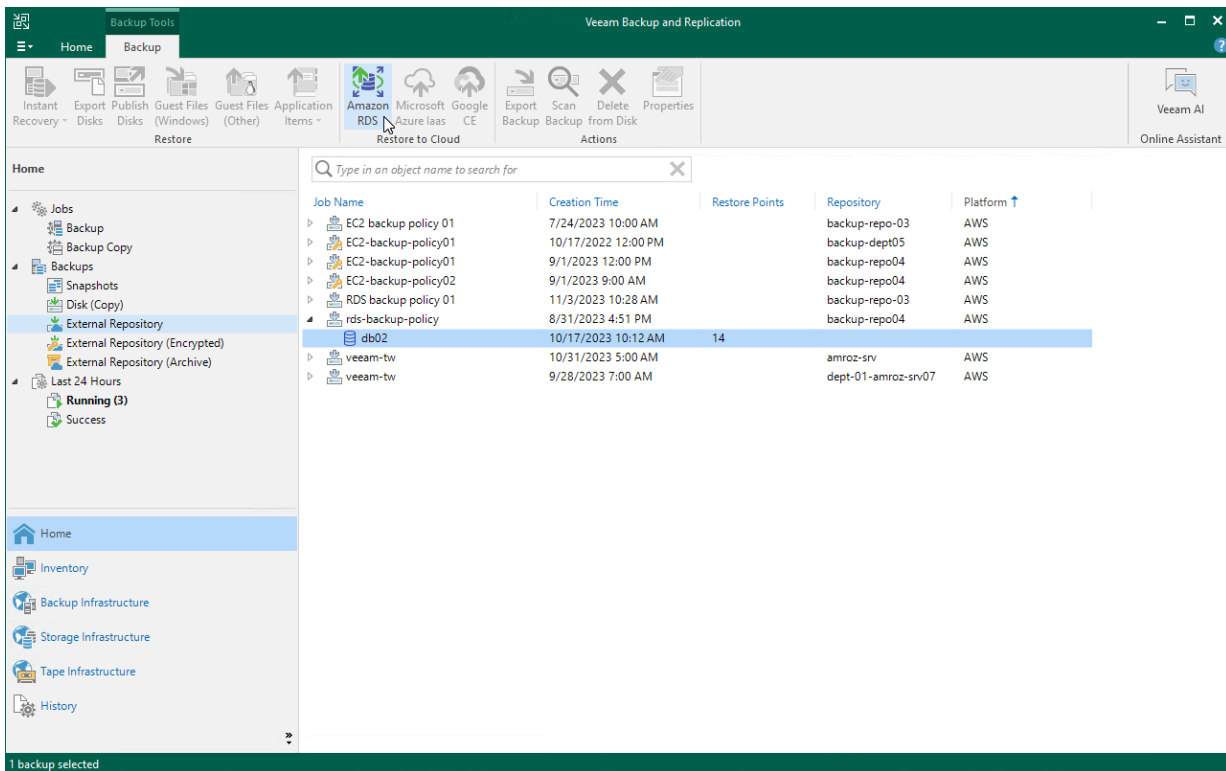
Restoring RDS Databases

You can recover corrupted databases of a DB instance running the PostgreSQL database engine from an image-level backup in the Veeam Backup for AWS Web UI only. However, you can launch the **RDS Database Restore** wizard directly from the Veeam Backup & Replication console to start the restore operation:

1. In the Veeam Backup & Replication console, open the **Home** view.
2. Navigate to **Backups > External Repository**.
3. Expand the backup policy that protects the database you want to recover, select the necessary database and click **Amazon RDS** on the ribbon.

Alternatively, you can right-click the selected database and click **Restore to Amazon RDS**.

Veeam Backup & Replication will open the **RDS Database Restore** wizard in a web browser. Complete the wizard as described in section [Performing Database Restore](#).



RDS Restore Using Web UI

Veeam Backup for AWS offers the following restore options:

- [RDS instance restore](#) – restores an entire DB instance or an Aurora DB cluster from a restore point.
- [Database restore](#) – restores specific databases of a PostgreSQL DB instance.

You can restore RDS resource data to the most recent state or to any available restore point.

Performing RDS Instance Restore

In case of a disaster, you can restore a DB instance or an Aurora DB cluster from a cloud-native snapshot or snapshot replica. Veeam Backup for AWS allows you to restore one or more RDS resources at a time, to the original location or to a new location.

How to Perform RDS Restore

To restore a protected RDS resource, do the following:

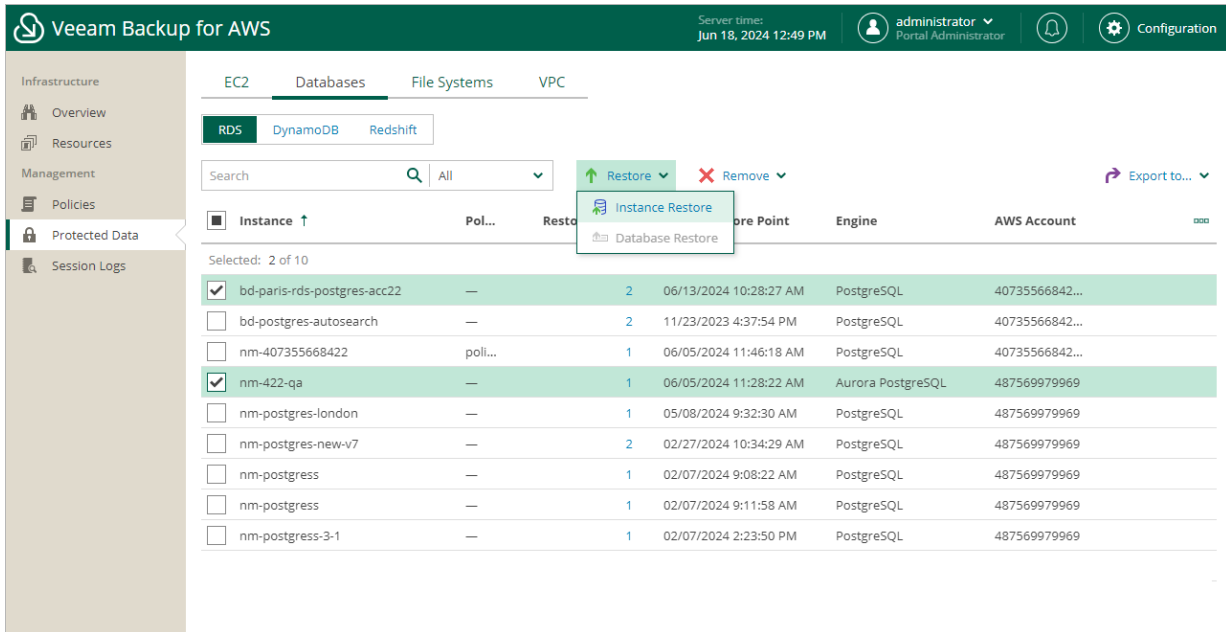
1. [Launch the RDS Restore wizard.](#)
2. [Select a restore point.](#)
3. [Specify an IAM identity for restore.](#)
4. [Choose a restore mode.](#)
5. [Enable encryption.](#)
6. [Configure RDS instance settings.](#)
7. [Configure network settings.](#)
8. [Specify a restore reason.](#)
9. [Finish working with the wizard.](#)

Step 1. Launch RDS Restore Wizard

To launch the **RDS Restore** wizard, do the following:

1. Navigate to **Protected Data > Databases > RDS**.
2. Select the RDS resource you want to restore.
3. Click **Restore > Instance Restore**.

Alternatively, click the link in the **Restore Points** column. Then, in the **Available Restore Points** window, select the necessary restore point and click **Restore**.



Step 2. Select Restore Point

At the **Instances** step of the wizard, you can add DB instances and Aurora DB clusters to the restore session and select restore points to be used to perform restore for each added RDS resource. By default, Veeam Backup for AWS uses the most recent valid restore point. However, you can restore an RDS resource to an earlier state.

To select a restore point, do the following:

1. Select the DB instance or Aurora DB cluster, and click **Restore Point**.
2. In the **Choose restore point** window, select the necessary restore point and click Apply.

To help you choose a restore point, Veeam Backup for AWS provides the following information on each available restore point:

- **Date** – the date when the restore point was created.
- **Type** – the type of the restore point:
 - *Snapshot* – a cloud-native snapshot created by a backup policy.
 - *Replica* – a snapshot replica created by a backup policy.
 - *Manual Snapshot* – a cloud-native snapshot created manually.
- **Restore Point Region** – the AWS Region where the restore point is stored.

The screenshot shows the Veeam Backup for AWS interface. At the top, the title bar reads 'Veeam Backup for AWS' with a server time of 'Jun 18, 2024 12:50 PM' and a user profile for 'administrator'. The main window is titled 'RDS Restore'. On the left, there is a sidebar with 'Instances' selected. The main area is divided into two panes. The left pane, 'Specify instances to restore', contains a search bar and a table with columns 'Instance', 'Engine', and 'Type'. The right pane, 'Choose restore point', is a modal dialog with a close button (X) and a table with columns 'Date', 'Type', and 'Restore Point Region'. The table lists two restore points: '06/13/2024 10:28:27 AM Manual snaps... Europe (Paris)' and '06/13/2024 12:03:37 AM Snapshot Europe (Paris)'. The second row is highlighted in green. At the bottom of the dialog are 'Apply' and 'Cancel' buttons.

Date ↓	Type	Restore Point Region
06/13/2024 10:28:27 AM	Manual snaps...	Europe (Paris)
06/13/2024 12:03:37 AM	Snapshot	Europe (Paris)

Step 3. Specify IAM Identity

At the **Account** step of the wizard, choose whether you want to [use an IAM role](#) or [one-time access keys of an IAM user](#) to allow Veeam Backup for AWS to perform the restore operation. For information on the permissions that the IAM role or IAM user must have to perform the restore operation, see [RDS Restore IAM Permissions](#).

IMPORTANT

Make sure that the specified IAM role or one-time access keys belong to an AWS account in which you plan to restore the selected RDS resources.

Specifying IAM Role

To specify an IAM role for restore, select the **IAM role** option and choose the necessary IAM role from the list.

For an IAM role to be displayed in the **IAM Role** list, it must be added to Veeam Backup for AWS with the *Amazon RDS Restore* operation selected as described in section [Adding IAM Roles](#). If you have not added the necessary IAM role to Veeam Backup for AWS beforehand, you can do it without closing the **RDS Restore** wizard. To do that, click **Add** and complete the [Add IAM Role](#) wizard.

IMPORTANT

It is recommended that you check whether the selected IAM role has all the permissions required to perform the operation. If some permissions of the IAM role are missing, the restore operation will fail to complete successfully. To run the IAM role permission check, click **Check Permissions** and follow the instructions provided in section [Checking IAM Role Permissions](#).

The screenshot shows the Veeam Backup for AWS interface during the RDS Restore process. The top navigation bar includes the Veeam logo, the text 'Veeam Backup for AWS', the server time 'Jun 18, 2024 12:50 PM', and user information 'administrator Portal Administrator'. The main content area is titled 'RDS Restore' and is currently on the 'Account' step. The 'Choose IAM role' section is active, with the instruction: 'Specify an IAM role that will be used to access resources for the restore operation or provide temporary access keys.' The 'IAM role' option is selected with a radio button. A dropdown menu is open, showing three roles: 'Default Backup Restore (Default Backup Restore)', 'rds-backup (Created by nm at 11/15/2023 9:21 AM)', and 'admin-942676447840 (Created by nm at 10/12/2023 12:19 PM)'. To the right of the dropdown are '+ Add' and 'Check Permissions' buttons. Below the dropdown is a 'Secret key:' field. An information icon and a note state: 'The keys are used to perform this operation only. They are not saved or stored. To learn what permissions are required for performing the operation, see the User Guide.' The bottom of the screen features 'Previous', 'Next', and 'Cancel' buttons.

Specifying One-Time Access Keys

To specify one-time access keys for restore, select the **Temporary access keys** option and use the **Access key** and **Secret key** fields to provide the access key ID and the secret access key.

NOTE

Veeam Backup for AWS does not store one-time access keys in the configuration database.

Veeam Backup for AWS Server time: Oct 30, 2023 12:40 PM administrator Portal Administrator Configuration

RDS Restore

Instances

Account

Restore Mode

Encryption

Settings

Network

Reason

Summary

Choose IAM role

Specify an IAM role that will be used to access resources for the restore operation or provide temporary access keys.

IAM role

Default Backup Restore (Default Backup Restore) + Add Check Permissions

Temporary access keys

Access key: DFRT6TGGDJKLP

Secret key:

i The keys are used to perform this operation only. They are not saved or stored. To learn what permissions are required for performing the operation, see the [User Guide](#).

Previous Next Cancel

Step 4. Choose Restore Mode

At the **Restore Mode** step of the wizard, choose whether you want to restore the selected RDS resources to the original or to a custom location. If you select the **Restore to new location, or with different settings** option, specify the target AWS Region where the restored DB instances and Aurora DB clusters will operate.

Limitations and Requirements

Before you choose the restore mode, consider the following limitations:

- Restore of RDS resources to the original location is not supported if the IAM role specified for the restore operation belongs to an AWS account that differs from the AWS account to which the source resources belong.
- Restore of RDS resources to the original location is not supported if deletion protection is enabled for the source resource.
- When restoring Aurora global databases, Veeam Backup for AWS restores only primary Aurora DB clusters in the primary AWS Regions; secondary clusters must be created manually in the AWS Management Console after the restore operation completes.

For more information on Amazon Aurora global databases, see [AWS Documentation](#).

- While restoring to a new location, Veeam Backup for AWS creates only primary DB instances in the restored clusters. Aurora Replicas for Aurora DB clusters with single-master replication must be added manually in the AWS Management Console after the restore operation completes.

To learn how to add DB instances to Amazon Aurora DB clusters, see [AWS Documentation](#).

The screenshot shows the 'RDS Restore' wizard in Veeam Backup for AWS. The 'Restore Mode' step is active, showing two options: 'Restore to original location' and 'Restore to new location, or with different settings'. The second option is selected, and a dropdown menu shows 'Europe (Ireland)' as the chosen region. A navigation bar at the bottom includes 'Previous', 'Next', and 'Cancel' buttons.

Server time: Jun 18, 2024 12:54 PM

administrator Portal Administrator Configuration

RDS Restore

Instances

Account

Restore Mode

Encryption

Settings

Network

Reason

Summary

Choose restore mode

Specify whether you want to restore the instance to the original location or to a new one, or with different settings.

i To restore to the original location, the IAM role specified for the operation must be from the same AWS account to which the backed-up instance belongs.

Restore to original location
Quickly restore the selected instance to the original location, with the same name and settings as the source instance.

Restore to new location, or with different settings
Perform additional configuration steps to restore the selected instance to a new location or to use settings that differ from the source settings.

Europe (Ireland)

Previous Next Cancel

Step 5. Enable Encryption

[This step applies only if you have selected the **Restore to new location, or with different settings** option at the **Restore Mode** step of the wizard]

At the **Encryption** step of the wizard, choose whether the restored RDS resources will be encrypted with AWS KMS keys:

- If you do not want to encrypt the RDS resources or want to apply the existing encryption scheme, select the **Use original encryption scheme** option.

IMPORTANT

If you plan to restore an unencrypted Aurora provisioned DB cluster to an Aurora Serverless DB cluster, and you select the **Use original encryption scheme** option, note that Veeam Backup for AWS will encrypt the newly created Aurora Serverless DB cluster with the default KMS key in the target AWS Region. For more information on Aurora Serverless, see [AWS Documentation](#).


- If you want to encrypt the RDS resources, select the **Restore as encrypted instance** option and choose the necessary KMS key from the **Encryption key** list.

For a KMS key to be displayed in the list of available encryption keys, it must be stored in the AWS Region selected at [step 4](#) of the wizard and the IAM role or user specified for the restore operation at [step 3](#) of the wizard must have permissions to access the key. For more information on KMS keys, see [AWS Documentation](#).

TIP

If the necessary KMS key is not displayed in the list, or if you want to use a KMS key from an AWS account other than the AWS account to which the specified IAM role belongs, you can select *Add custom key ARN* from the **Encryption key** drop-down list, and specify the amazon resource name (ARN) of the key in the **Add Custom Key ARN** window.

For Veeam Backup for AWS to be able to encrypt the restored RDS resource using the provided KMS key, either the IAM role or user specified for the restore operation, or the IAM role used to create the restore point selected at [step 2](#) of the wizard must have permissions to access the key.



Step 6. Configure Restore Settings

[This step applies only if you have selected the **Restore to new location, or with different settings** option at the **Restore Mode** step of the wizard]

At the **Settings** step of the wizard, specify settings for the restored RDS resources. To do that, follow the instructions provided in sections [Configuring Settings for DB Instances](#) and [Configuring Settings for Aurora DB Clusters](#).

TIP

The **Settings** step also contains some preconfigured settings retrieved from the source RDS resources. If you want to specify advanced configuration settings for a restored DB instance or Aurora DB cluster, select the necessary resource and click **Advanced Options**. For more information on all available settings that can be specified for RDS resources, see the [Amazon RDS User Guide](#) and [Amazon Aurora User Guide](#).

Configuring Settings for DB Instances

To configure settings for a restored DB instance, at the **Settings** step of the wizard, select the necessary instance and click **Edit**. In the opened window, do the following:

1. In the **Instance identifier** section, specify an identifier for the restored DB instance. Consider the following limitations:
 - The instance identifier must be unique for each AWS Region within one AWS Account.
 - The instance identifier can contain only lowercase Latin letters and hyphens, but cannot contain two consecutive hyphens.
 - The first character of the instance identifier must be a letter. The last character of the identifier must not be a hyphen.
 - The maximum length of the instance identifier is 63 characters.

For more information on limitations for DB instance identifiers, see [AWS Documentation](#).

2. In the **Instance specifications** section, choose a DB instance class and storage type for the restored instance. If you choose the *Provisioned IOPS (SSD)* storage type, you must also specify an IOPS rate.

For the list of all supported DB instance classes and available storage types, see [AWS Documentation](#).

3. In the **Instance options** section, specify a parameter group and an option group that will be associated with the restored instance:
 - a. From the **Parameter group** drop-down list, select the parameter group containing database engine configuration values that will be applied to the restored DB instance.

For a parameter group to be displayed in the list of available groups, the group must be created beforehand as described in [AWS Documentation](#).

- b. [This step does not apply to DB instances running the PostgreSQL database engine] From the **Option group** drop-down list, select the option group containing database configuration values and security settings that will be applied to the restored DB instance.

For an option group to be displayed in the list of available groups, the group must be created beforehand as described in [AWS Documentation](#).

NOTE

If you select the **Use default group** option, Veeam Backup for AWS will associate the restored DB instance with the default parameter group and the default option group automatically created by AWS during the restore operation.

4. Click **Apply**.

The screenshot shows the Veeam Backup for AWS interface for RDS Restore. The top bar includes the Veeam logo, server name 'administrator', and date 'Jun 18, 2024 12:56 PM'. The main content area is titled 'RDS Restore' and features a left-hand navigation menu with options like 'Instances', 'Account', 'Restore Mode', 'Encryption', 'Settings', 'Network', 'Reason', and 'Summary'. The 'Settings' section is active, displaying a table of instances:

Name	Engine	Instance Class
pi-aurora	Aurora P...	—
pi-maria	MariaDB	—
pi-external-irela...	PostgreS...	db.t3.micro

Below the table, the configuration panel for the selected instance 'pi-maria' is shown. It includes the following sections:

- Instance identifier:** A text input field containing 'pi-maria-restored'.
- Instance specifications:** A section for configuring settings for the restored instance, including:
 - Instance class:** A dropdown menu set to 'db.m6g.4xlarge (16 cores, 64GB)'.
 - Storage type:** A dropdown menu set to 'Provisioned IOPS (SSD)'.
 - Provisioned IOPS:** A numeric input field set to '1000'.
- Instance options:** A section for specifying a parameter and an option group, including:
 - Parameter group:** A dropdown menu set to 'default:mariadb10.6'.
 - Option group:** A dropdown menu set to 'default:mariadb-10-6'.

At the bottom of the configuration panel, there are 'Apply' and 'Cancel' buttons. A warning message is displayed: 'Changing the storage type can impact instance performance.'

Configuring Settings for Aurora DB Clusters

A number of settings that you can configure for a restored cluster depends on the capacity type that you plan to choose for the cluster. AWS supports Aurora DB clusters of 2 different capacity types:

- **Aurora provisioned DB cluster** – a cluster whose capacity is managed manually by creating DB instances: a single primary DB instance (writer) and multiple Aurora Replicas (readers) in Aurora DB clusters. For more information on provisioned DB clusters, see [AWS Documentation](#).
- **Aurora Serverless v2** – a clusters whose capacity is scaled automatically according to the specified minimum and maximum capacity values. For more information on Aurora Serverless, see [AWS Documentation](#).

NOTE

You cannot change replication settings for restored Aurora DB clusters. Veeam Backup for AWS restores the clusters with the same replication settings configured for the source clusters.

Configuring Restore Settings

To specify settings for a restored Aurora DB cluster, at the **Settings** step of the wizard, select the necessary cluster and click **Edit**. In the opened window, do the following:

1. In the **Instance specifications** section, specify configuration settings for the restored Aurora DB cluster:
 - a. Set the **Use global database** toggle to *On* if you plan that the restored cluster will have secondary DB clusters in a number of AWS Regions. In this case, the **Version** list will be filtered to show only Aurora database versions that support this feature. However, Veeam Backup for AWS will still create only a primary cluster in the AWS Region selected at [step 4](#) of the wizard; secondary clusters must be created manually in the AWS Management Console after the restore operation completes.

For more information on Amazon Aurora global databases, see [AWS Documentation](#).

- b. [This step applies only to Aurora MySQL DB clusters] Set the **Use parallel query** toggle to *On* if you plan to use the Aurora MySQL parallel query feature to improve I/O performance and to reduce network traffic in the restored cluster. In this case, the **Version** list will be filtered to show only Aurora database versions that support this feature. Keep in mind that to be able to use the feature, you must enable the `aurora_parallel_query` parameter in the DB cluster parameter group that you will specify in the **Instance options** section.

For more information on Aurora MySQL parallel query, see [AWS Documentation](#).

- c. From the **Version** drop-down list, select an Aurora database engine version for the restored cluster. The list shows only DB engine versions supported in the target AWS Region, and is filtered based on the DB engine type and DB engine version of the source Aurora DB cluster. The number of versions displayed in the list also depends on the source cluster replication settings and options that you have selected at steps 1b and 1c.

For more information on Amazon Aurora database engine versions, see [AWS Documentation](#).

NOTE

If you restore Aurora PostgreSQL DB clusters and plan to use the **Babelfish** feature to allow the restored clusters to accept database connections from Microsoft SQL Server clients, note that this feature is supported only for Aurora PostgreSQL 13.4 and later engine versions.

- d. In the **Cluster identifier** field, specify an identifier for the restored cluster. Consider the following limitations:
 - The cluster identifier must be unique for each AWS Region within one AWS Account.
 - The cluster identifier can contain only lowercase Latin letters and hyphens, but cannot contain 2 consecutive hyphens.
 - The first character of the cluster identifier must be a letter. The last character of the identifier must not be a hyphen.
 - The maximum length of the cluster identifier is 63 characters.

For more information on limitations for Aurora DB cluster identifiers, see [AWS Documentation](#).

- e. From the **Instance class** drop-down list, select a DB instance class that Veeam Backup for AWS will use to create the primary DB instance in the restored cluster. For the list of all supported DB instance classes, see [AWS Documentation](#).

If you want to restore the primary DB instance of the provisioned cluster as an Aurora Serverless v2 DB instance, select `db.serverless` from the **Instance class** drop-down list. Consider that Aurora Serverless v2 is supported only for a limited list of DB engine versions. For more information, see [AWS Documentation](#).

- f. [This step applies only to Aurora Serverless v2] Use the **Minimum capacity** and **Maximum capacity** fields to specify a range of capacity units that will be used to create scaling rules for the restored cluster. These rules define thresholds for CPU utilization, connections and available memory.

For more information on capacity units and scaling rules, see [AWS Documentation](#).

- g. In the **Instance identifier** field, specify an identifier for the primary DB instance in the restored cluster. Consider the following limitations:

- The instance identifier must be unique for each AWS Region within one AWS Account.
- The instance identifier can contain only lowercase Latin letters and hyphens, but cannot contain two consecutive hyphens.
- The first character of the instance identifier must be a letter. The last character of the identifier must not be a hyphen.
- The maximum length of the instance identifier is 63 characters.

For more information on limitations for DB instance identifiers, see [AWS Documentation](#).

2. In the **Instance options** section, specify a DB cluster parameter group that will be associated with the restored cluster and a DB parameter group that will be associated with the primary DB instance:

- a. From the **Cluster parameter group** drop-down list, select the DB cluster parameter group containing database engine configuration values that will be applied to every DB instance launched in the restored cluster.

For a DB cluster parameter group to be displayed in the list, the group must be created beforehand as described in [AWS Documentation](#).

- b. From the **Parameter group** drop-down list, select the DB parameter group containing database engine configuration values that will be applied to the primary DB instance in the restored Aurora provisioned DB cluster or to the restored Aurora Serverless v2 cluster.

For a DB parameter group to be displayed in the list, the group must be created beforehand as described in [AWS Documentation](#).

NOTE

If Veeam Backup for AWS cannot find any parameter groups in the target AWS Region, the **Use default group option** will be displayed. Use this option to associate the restored DB cluster and the primary DB instance with the default parameter groups that will be automatically created by AWS during the restore operation.

3. Click Apply.

Veeam Backup for AWS Server time: Jun 18, 2024 12:56 PM administrator Portal Administrator Configuration

RDS Restore

Configure restore settings
Specify settings for the restored instances.

[Edit](#) [Advanced Options](#)

Name	Engine ↑	Instance Class
pi-maria-restored	MariaDB	db.m6g.4xlarge
pi-external-irela...	PostgreS...	db.t3.micro
pi-aurora	Aurora P...	—

Instance specifications
Specify configuration settings for the restored instance such as the capacity type, engine version, cluster identifier and others.

Capacity type:

Use global database: On

Version:

Cluster identifier:

Instance class:

Instance identifier:

Instance options
Specify a parameter and an option group that will be associated with the restored instance.

[Rescan](#)

Cluster parameter group:

Parameter group:

Option group:

Step 7. Configure Network Settings

[This step applies only if you have selected the **Restore to a new location, or with different settings** option at the **Restore Mode** step of the wizard]

At the **Network** step of the wizard, configure network and security settings for the restored DB instances and Aurora DB clusters. To do that, select the necessary RDS resource and click **Edit**. In the opened window, do the following:

1. In the **Network settings** section, specify network settings for the restored RDS resource:
 - For a restored DB instance, choose an Amazon VPC to which the instance will be connected, a subnet group that will be assigned to the instance, an Availability Zone where the instance will reside, and a port that will be used to access the DB instance. Note that the **VPC** list shows only Amazon VPCs that include one or more subnet groups.

For a VPC and a subnet group to be displayed in the lists of available network specifications, they must be created in the AWS Region specified at [step 4](#) of the wizard as described in [AWS Documentation](#).

TIP

If you want to create a passive secondary replica (standby instance) of the restored DB instance, set the **Multi-AZ deployment** toggle to *On*. Keep in mind that Multi-AZ deployments are not supported for instances running MS SQL Server Express and MS SQL Server Web editions. For more information on Multi-AZ deployments, see [AWS Documentation](#).

- For a restored Aurora DB cluster, choose an Amazon VPC to which the cluster will be restored, a subnet group that includes at least two subnets created in two different Availability Zones of the AWS Region specified at [step 4](#) of the wizard, an Availability Zone where the primary DB instance will reside, and a port that will be used to access the primary DB instance.
2. In the **Security settings** section, specify security settings to control what IP addresses will be able to connect to databases of the restored RDS resource.
 - a. To make the RDS resource accessible outside the selected Amazon VPC, set the **Public accessible** toggle to *On*. Note that the RDS resource must belong to a public subnet group to become publicly accessible.
 - b. To specify security groups that will control access to the RDS resource, click the link in the **Security group** field and then select the necessary groups in the **Select Security Group** window.

3. Click Apply.

Veeam Backup for AWS Server time: Jun 18, 2024 12:56 PM administrator Portal Administrator Configuration

RDS Restore

Configure network settings
Specify network settings for the restored instances.

Instance	VPC ↑	Subnet
pi-maria-restored	vpc-6d9e8c0b	—
pi-aurora	vpc-6d9e8c0b	—
pi-external-ireland	vpc-3748a14e	default

Network settings X

Specify network settings for the restored instance such as a VPC, subnet group, port, and choose whether you want to use Multi-AZ deployment or a preferred availability zone.

VPC: vpc-3748a14e
Subnet group: default
Multi-AZ deployment: No
Availability zone: eu-west-1b
Port: 3306

Security settings

Specify security settings for the restored instance such as instance public accessibility and security groups.

Public accessible: On
Security group: 1 security group selected.

Apply **Cancel**


Step 8. Specify Restore Reason

At the **Reason** step of the wizard, you can specify a reason for restoring the RDS resources. This information will be saved to the session history and you will be able to reference it later.

The screenshot shows the Veeam Backup for AWS interface during the RDS Restore process. The top navigation bar includes the Veeam logo, the text "Veeam Backup for AWS", the server time "Jun 18, 2024 12:57 PM", and user information "administrator Portal Administrator". A "Configuration" icon is also present. Below the navigation bar, the "RDS Restore" wizard is displayed. On the left, a sidebar menu lists the steps: "Instances", "Account", "Restore Mode", "Encryption", "Settings", "Network", "Reason" (which is highlighted in green), and "Summary". The main content area is titled "Specify restore reason" and contains a text input field with the text "Restoring RDS resources to Ireland". At the bottom of the wizard, there are three buttons: "Previous", "Next" (highlighted in green), and "Cancel".

Step 9. Finish Working with Wizard

At the **Summary** step of the wizard, review summary information and click **Finish**.



Performing RDS Database Restore

In case of a disaster, you can restore corrupted databases of a PostgreSQL DB instance from an image-level backup. Veeam Backup for AWS allows you to restore one or more databases of a PostgreSQL DB instance at a time, to the original location or to a new location.

How to Perform Database Restore

To restore databases of a protected DB instance, do the following:

1. [Launch the Database Restore wizard.](#)
2. [Select databases.](#)
3. [Specify an IAM identity for restore.](#)
4. [Specify data retrieval settings for archived backups.](#)
5. [Configure target instance settings.](#)
6. [Specify a restore reason.](#)
7. [Finish working with the wizard.](#)

Step 1. Launch RDS Database Restore Wizard

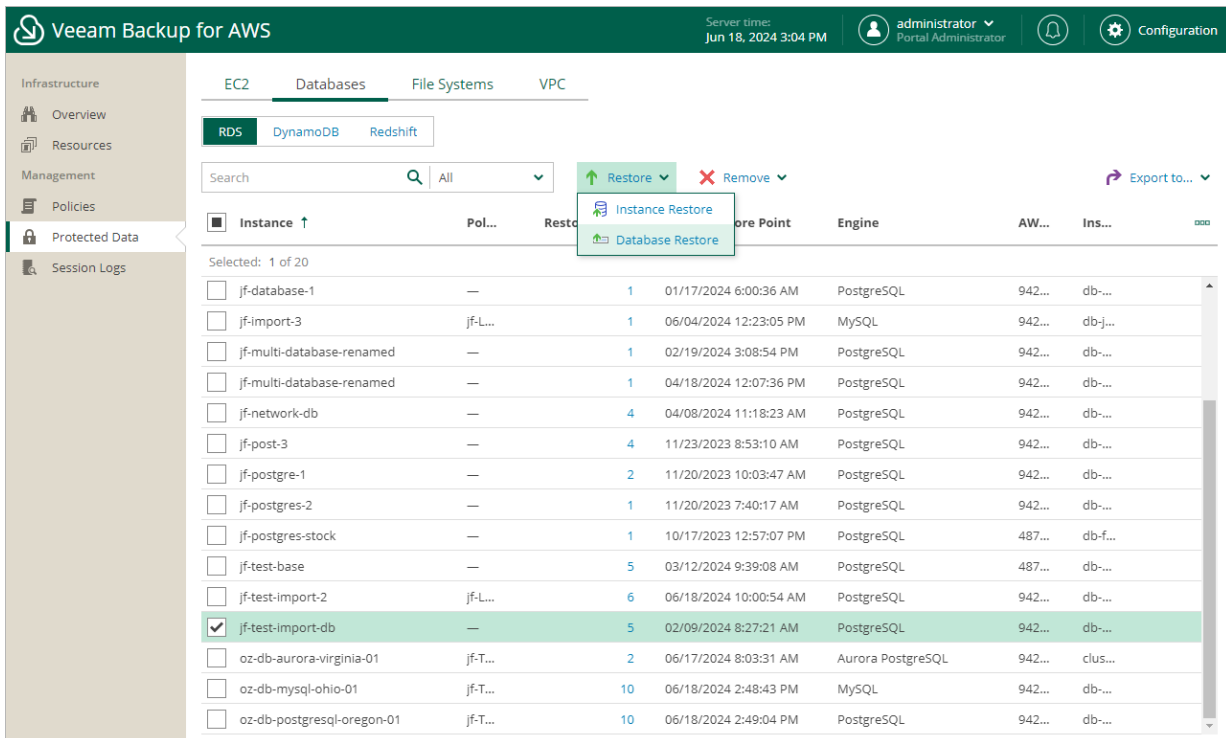
To launch the **RDS Database Restore** wizard, do the following.

1. Navigate to **Protected Data > Databases > RDS**.
2. Select the DB instance whose databases you want to restore, and click **Restore > Database Restore**.

Alternatively, click the link in the **Restore Points** column. Then, in the **Available Restore Points** window, select the necessary restore point and click **Restore > Database Restore**.

IMPORTANT

If you select multiple DB instances, you will not be able to proceed with the **RDS Database Restore** wizard.

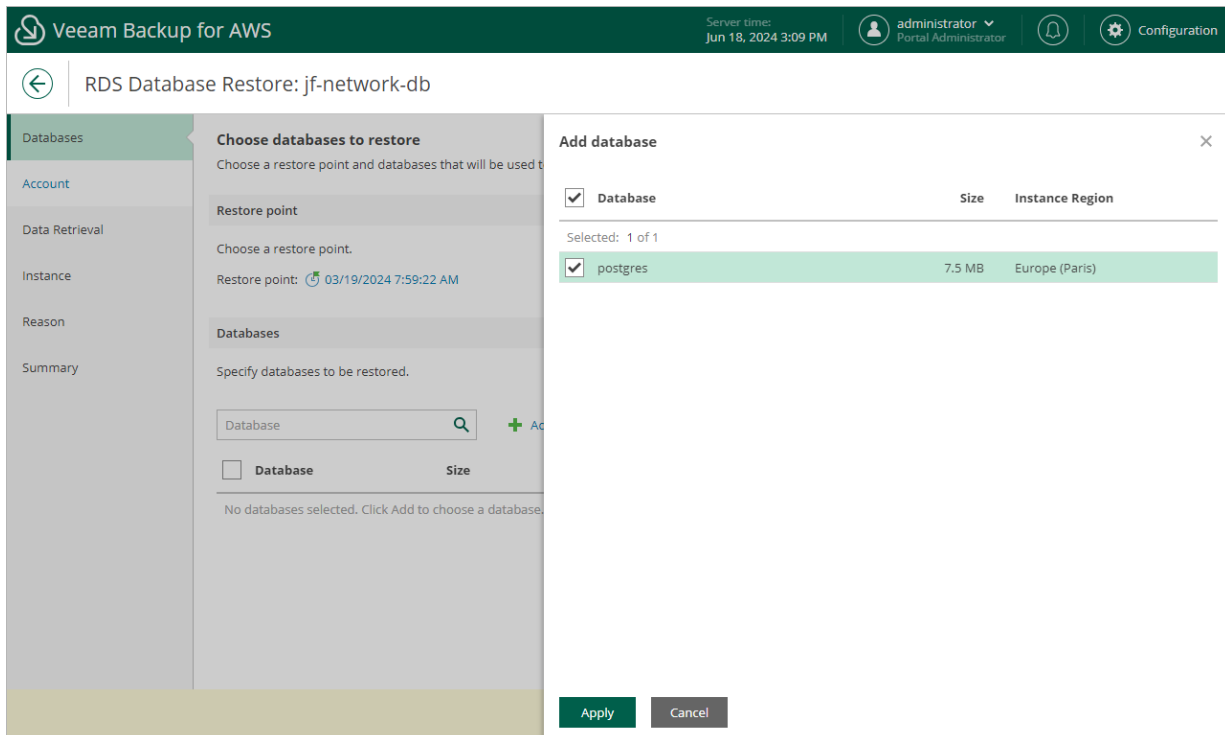


Step 2. Select Databases

At the **Databases** step of the wizard, select a restore point that will be used to perform the restore operation for each database, and then click **Add** to select databases to restore. By default, Veeam Backup for AWS uses the most recent valid restore point. However, you can restore the database data to an earlier state.

To help you choose a restore point, Veeam Backup for AWS provides the following information on each available restore point:

- **Date** – the date when the restore point was created.
- **Type** – the type of the restore point:
 - *Backup* – an image-level backup created by a backup policy.
 - *Archive* – an archived backup created by a backup policy.
- **State** – the state of the restore point stored in the standard backup repository:
 - *Healthy* – the restore point has been verified by the health check session and reported to be healthy.
 - *Incomplete* – the restore point has been verified by the health check session and reported to be corrupted or incomplete.
- **Storage Class** – the storage class of the backup repository where the restore point is stored.
- **Restore Point Region** – the AWS Region where the restore point is stored.
- **AWS Account** – the AWS account to which the DB instance belongs.



The screenshot shows the Veeam Backup for AWS interface during the 'RDS Database Restore' wizard. The main window is titled 'RDS Database Restore: jf-network-db'. On the left, there is a navigation pane with 'Databases' selected. The main content area is divided into two sections: 'Choose databases to restore' and 'Add database'.

The 'Choose databases to restore' section shows a restore point of '03/19/2024 7:59:22 AM' and a search bar for databases. Below it, there is a table with one row: 'Database' and 'Size'. The text below the table says 'No databases selected. Click Add to choose a database.'

The 'Add database' modal window is open, showing a table with one selected database:

<input checked="" type="checkbox"/>	Database	Size	Instance Region
<input checked="" type="checkbox"/>	postgres	7.5 MB	Europe (Paris)

Below the table, it says 'Selected: 1 of 1'. At the bottom of the modal, there are 'Apply' and 'Cancel' buttons.

Step 3. Specify IAM Identity

At the **Account** step of the wizard, specify IAM roles that Veeam Backup for AWS will use to perform the restore operation.

IMPORTANT

Make sure that the specified IAM roles belong to an AWS account in which you plan to restore the selected databases.

Configuring Worker Settings

At the **Account** step of the wizard, do the following:

1. In the **IAM role** section, specify an IAM role to allow Veeam Backup for AWS to perform the restore operation. For information on the permissions that the IAM role must have to perform the restore operation, see [RDS Database Restore IAM Permissions](#).

For an IAM role to be displayed in the **IAM role** list, it must be added to Veeam Backup for AWS with the *Amazon RDS Restore* operation selected as described in section [Adding IAM Roles](#). If you have not added the necessary IAM role to Veeam Backup for AWS beforehand, you can do it without closing the **RDS Database Restore** wizard. To do that, click **Add** and complete the **Add IAM Role** wizard.

2. In the **Worker deployment** section, specify an IAM role that will be attached to the worker instances and used by Veeam Backup for AWS to communicate with these instances. For information on the permissions that the IAM role must have to perform the restore operation, see [Worker Deployment Role Permissions in Production Accounts](#).

For an IAM role to be displayed in the **IAM role** list, it must be added to Veeam Backup for AWS with the *Production worker role* selected as described in section [Adding IAM Roles](#). If you have not added the necessary IAM role to Veeam Backup for AWS beforehand, you can do it without closing the **RDS Database Restore** wizard. To do that, click **Add** and complete the **Add IAM Role** wizard.

IMPORTANT

It is recommended that you check whether the selected IAM roles have all the permissions required to perform the operation. If some permissions of the IAM role permissions are missing, the restore operation will fail to complete successfully. To run the IAM role permission check, click **Check Permissions** and follow the instructions provided in section [Checking IAM Role Permissions](#).

Worker Instance Requirements

To restore DB instance databases from image-level backups, Veeam Backup for AWS deploys worker instances in an AWS Region where DB instance that will host the restored databases resides in an AWS account to which the instance belongs. By default, Veeam Backup for AWS uses the most appropriate network settings of AWS Regions to deploy worker instances. However, you can add [specific worker configurations](#) that will be used to deploy worker instances used for database restore operations.

If no specific [worker configurations](#) are added to Veeam Backup for AWS, the most appropriate network settings of AWS Regions are used to deploy worker instances for the database restore operation. For Veeam Backup for AWS to be able to deploy a worker instance used to perform the restore operation:

- The DNS resolution option must be enabled for the VPC. For more information, see [AWS Documentation](#).

- As Veeam Backup for AWS uses public access to communicate with worker instances, the [public IPv4 addressing](#) attribute must be enabled at least for one subnet in the Availability Zone where the DB instance resides and the VPC to which the subnet belongs must have an [internet gateway attached](#). VPC and subnet route tables must have routes that direct internet-bound traffic to this internet gateway.

If you want worker instances to operate in a private network, enable the [private network deployment](#) functionality and configure [specific VPC endpoints](#) for the subnet to let Veeam Backup for AWS use private IPv4 addresses. Alternatively, configure VPC interface endpoints as described in section [Appendix C. Configuring Endpoints in AWS](#).

NOTE

During RDS image-level backup operations, Veeam Backup for AWS creates 2 additional security groups that are further associated with the source DB instances and worker instances to allow direct network traffic between them. To learn how DB instance database restore works, see [Database Restore](#).

The screenshot shows the Veeam Backup for AWS interface. The main window is titled "RDS Database Restore: jf-network-db". On the left, there is a navigation menu with options: Databases, Account, Data Retrieval, Instance, Reason, and Summary. The "Account" section is currently selected. The main content area is divided into two panels. The left panel is titled "Choose IAM role" and contains instructions: "Specify the pre-created IAM role to use to deploy on required permissions see the User Guide." Below this, there are sections for "IAM role" and "Worker deployment", both with an "IAM role:" label and a dropdown menu set to "Default Backup Restore (Default Back...". The right panel is titled "Permission check" and contains a green checkmark with the text "Your account meets the required permissions." Below this, there are buttons for "Grant", "Recheck", and "Export Missing Permissions". At the bottom of the "Permission check" panel, there is a table with the following data:

Type	Status	Missing Permissions
EC2 permissions	Passed	—
SERVICEQUOTAS permissions	Passed	—
IAM permissions	Passed	—
RDS permissions	Passed	—
SQS permissions	Passed	—
SSM permissions	Passed	—

At the bottom of the "Permission check" panel, there is a "Close" button.

Step 4. Specify Data Retrieval Settings

[This step applies only if you have selected to restore from the archived restore point]

At the **Data Retrieval** step of the wizard, choose a retrieval mode and specify a period for which you want to keep the data available. To do that:

1. In the **Retrieval Mode** section, click the link.
 - a. In the **Choose retrieval mode** window, choose the retrieval mode that Veeam Backup for AWS will use to retrieve the archived data:
 - **Expedited** – the most expensive option. The retrieved data is available within 1-5 minutes.
Amazon does not support this option for data stored in the S3 Glacier Deep Archive storage class. For more information, see [AWS Documentation](#).
 - **Standard** – the recommended option. The retrieved data is available within 3-5 hours for data stored in the S3 Glacier Flexible Retrieval storage class and within 12 hours for data stored in the S3 Glacier Deep Archive storage class.
 - **Bulk** – the least expensive option. The retrieved data is available within 5-12 hours for data stored in the S3 Glacier Flexible Retrieval storage class and within 48 hours for data stored in the S3 Glacier Deep Archive storage class.
 - **Standard accelerated** – the option that is less expensive than the **Expedited** option. The retrieved data is available within 15-30 minutes for data stored in the S3 Glacier Flexible Retrieval storage class.

With this option enabled, Veeam Backup for AWS leverages the [S3 Batch Operations functionality](#) to retrieve the archived data.

TIP

Before you enable the **Standard accelerated** option, it is recommended that you check whether the IAM role specified to access the archive backup repository has all the permissions required to perform data retrieval operations using the S3 Batch Operations functionality, as described in section [Checking IAM Role Permissions](#).

If some of the IAM role permissions required to perform data retrieval operations using the S3 Batch Operations functionality are missing, Veeam Backup for AWS will use the **Standard** option to retrieve data.

For more information on archive retrieval options, see [AWS Documentation](#).

- b. To save changes made to the data retrieval settings, click **Apply**.
2. In the **Availability Period** section, click **Edit Availability Period**.
 - a. In the **Availability settings** window, specify the number of days for which you want to keep the data available for restore operations.

If you want to receive an email notification when the data is about to expire, select the **Send email notifications** check box and specify the number of hours before the expiration time when the notification will be sent.

b. To save changes made to the availability period settings, click **Apply**.

The screenshot shows the Veeam Backup for AWS interface. At the top, the header includes the Veeam logo, the text "Veeam Backup for AWS", the server time "Jun 18, 2024 3:09 PM", and the user "administrator Portal Administrator". The main content area is titled "RDS Database Restore: jf-network-db". On the left, a navigation menu lists "Databases", "Account", "Data Retrieval", "Instance", "Reason", and "Summary". The "Data Retrieval" section is expanded, showing "Configure data retrieval settings" with a sub-header "Based on your time and cost requirements, choose a retrieval mode you want the retrieved data to be available." Below this, it shows "Retrieval Mode: Standard" and "Availability Period" settings: "Data available for: 1 day" and "Notification email: Disabled". A modal window titled "Choose retrieval mode" is open, displaying four options: "Expedited", "Standard accelerated", "Standard" (selected), and "Bulk". Each option has a brief description of its characteristics and typical completion times. At the bottom of the modal, there are "Apply" and "Cancel" buttons.

Step 5. Configure Target Instance Settings

At the **Instance** step of the wizard, specify the target AWS Region where a DB instance will host the restored databases, and choose the target DB instance.

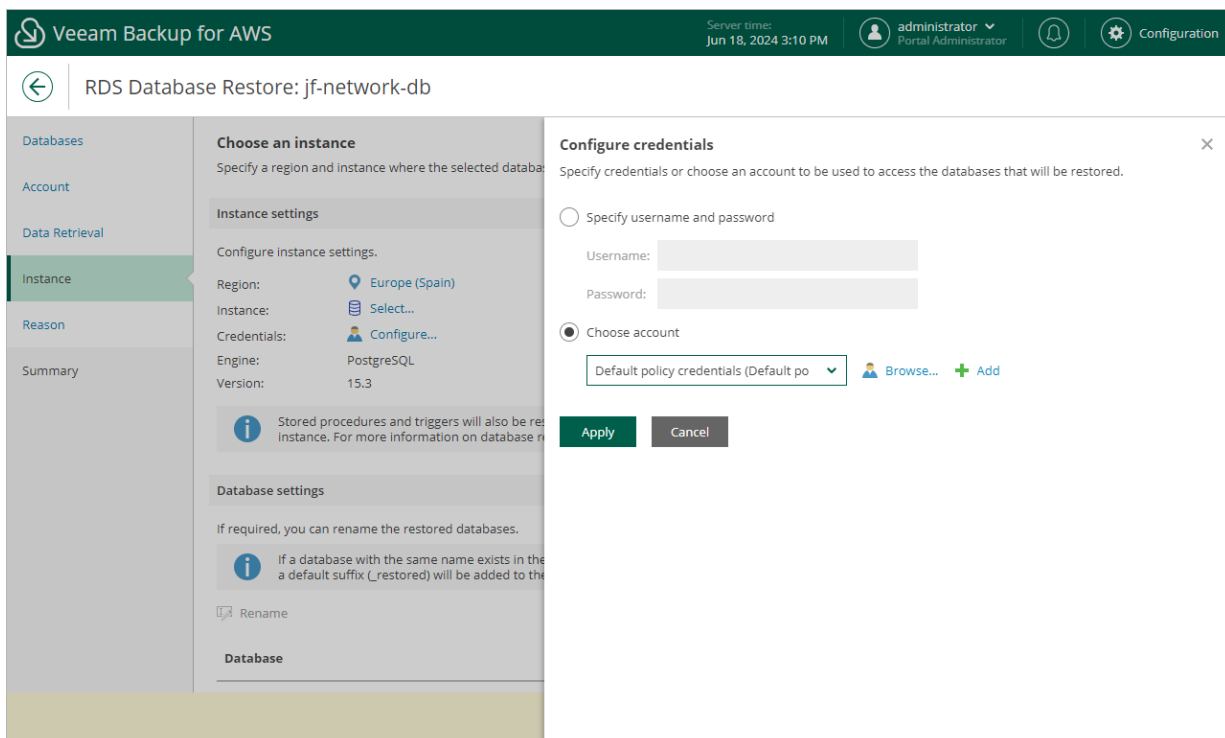
NOTE

If the specified target region is the same region where the source DB instance resided and if the specified DB instance is the same instance that hosted the source databases, the next run of the backup policy protecting the source DB instance will take more time to complete in case the restore operation completes successfully. This is the expected behavior caused by [AWS technical limitations](#).

You must also specify credentials of a database account that Veeam Backup for AWS will use to connect to the target DB instance. To do that, click the link in the **Credentials** field. In the **Configure credentials** window, choose whether you want to provide temporary credentials or select an existing database account. For an account to be displayed in the list of available accounts, it must be added to Veeam Backup for AWS as described in section [Adding Database Accounts](#). If you have not added an account beforehand, click **Add** and complete the **Add Account** wizard.

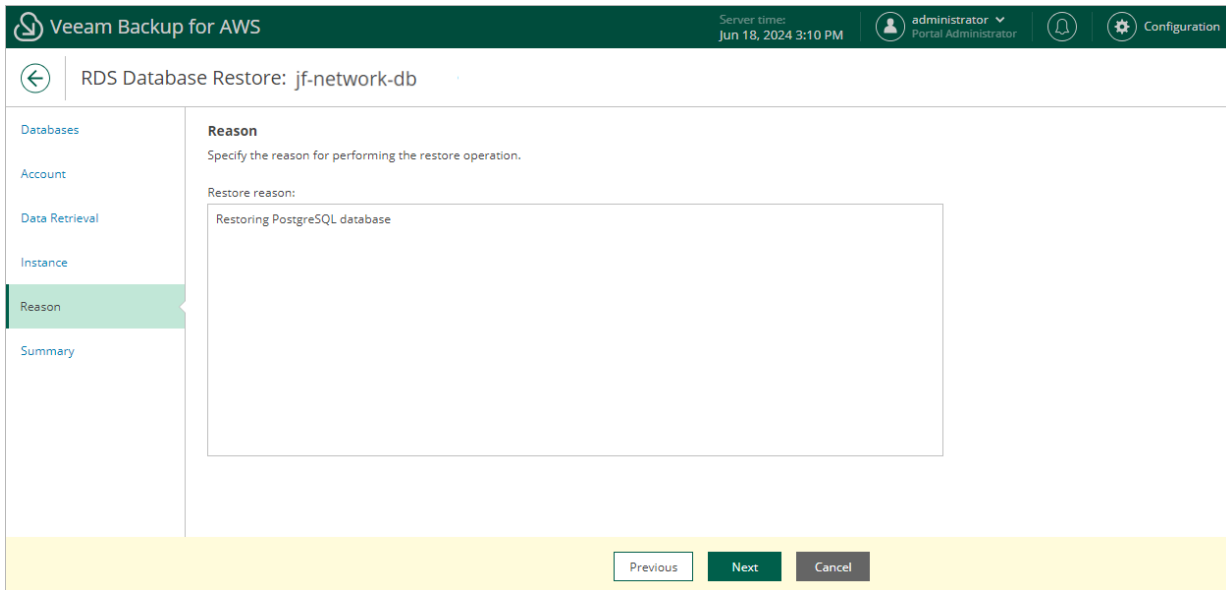
TIP

By default, Veeam Backup for AWS restores each database with the same name as the original database. If a database with this name already exists in the production infrastructure, a default suffix is added to the database name automatically. You can also specify a new name for each restored database manually – to do that, select the database in the **Database settings** section and click **Rename**.



Step 6. Specify Restore Reason

At the **Reason** step of the wizard, you can specify a reason for restoring the databases. This information will be saved to the session history, and you will be able to reference it later.



The screenshot displays the Veeam Backup for AWS interface during the 'RDS Database Restore' process. The top navigation bar includes the Veeam logo, the text 'Veeam Backup for AWS', the server time 'Jun 18, 2024 3:10 PM', and user information 'administrator Portal Administrator'. A left sidebar contains navigation links for 'Databases', 'Account', 'Data Retrieval', 'Instance', 'Reason' (highlighted), and 'Summary'. The main content area is titled 'RDS Database Restore: jf-network-db' and features a 'Reason' section with the instruction 'Specify the reason for performing the restore operation.'. Below this is a text input field containing 'Restoring PostgreSQL database'. At the bottom of the interface, there are three buttons: 'Previous', 'Next', and 'Cancel'.

Step 7. Finish Working with Wizard

At the **Summary** step of the wizard, review summary information and click **Finish**.

The screenshot shows the 'Summary' step of the 'RDS Database Restore: jf-network-db' wizard in the Veeam Backup for AWS interface. The interface includes a top navigation bar with the Veeam logo, server time (Jun 18, 2024 3:10 PM), and user information (administrator, Portal Administrator). A left sidebar lists navigation options: Databases, Account, Data Retrieval, Instance, Reason, and Summary (which is highlighted). The main content area displays the following summary information:

- Summary**: Review the restore settings, and click Finish to exit the wizard.
- Reason**: Restoring PostgreSQL database
- Instance**:
 - Region: Asia Pacific (Mumbai)
 - Instance: postgres-database
 - Credentials: donna_ortiz
 - Engine: PostgreSQL
 - Version: 15.3
- Account**: RDS Backup and Restore Role
- Databases**: 1 database will be restored

At the bottom of the wizard, there are three buttons: 'Previous', 'Finish', and 'Cancel'.

DynamoDB Restore

The actions that you can perform with restore points of DynamoDB tables depend on whether you access the restore points using the Veeam Backup & Replication console or the Veeam Backup for AWS Web UI.

DynamoDB Restore Using Console

You can recover corrupted DynamoDB tables in the Veeam Backup for AWS Web UI only. However, you can launch the **DynamoDB Table Restore** wizard directly from the Veeam Backup & Replication console to start the restore operation:


1. In the Veeam Backup & Replication console, open the **Home** view.
2. Navigate to **Backups > Snapshots**.
3. Expand the backup policy that protects the DynamoDB tables that you want to recover, select the necessary table and click **Amazon DynamoDB** on the ribbon.

Alternatively, you can right-click the selected table and click **Restore to Amazon DynamoDB**.

IMPORTANT

You cannot restore multiple DynamoDB tables from the Veeam Backup & Replication console.

Veeam Backup & Replication will open the **DynamoDB Table Restore** wizard in a web browser. Complete the wizard as described in section [DynamoDB Restore Using Web UI](#).



DynamoDB Restore Using Web UI

In case of a disaster, you can restore a DynamoDB table from a DynamoDB backup or backup copy. Veeam Backup for AWS allows you to restore one or more DynamoDB tables at a time, to the original location or to a new location. To learn how DynamoDB restore works, see [DynamoDB Restore](#).

IMPORTANT

Consider the following:

- Veeam Backup for AWS supports restore of DynamoDB tables only to the same AWS account where the source tables reside.
- Veeam Backup for AWS supports restore of only those DynamoDB table properties that are described in section [Protecting DynamoDB Tables](#).

How to Perform DynamoDB Restore

To restore a protected DynamoDB table, do the following:

1. [Launch the DynamoDB Restore wizard](#).
2. [Select a restore point](#).
3. [Specify an IAM identity for restore](#).
4. [Choose a restore mode](#).
5. [Enable encryption for the restored table](#).
6. [Specify configuration settings](#).
7. [Choose capacity mode for the restored table](#).
8. [Specify a restore reason](#).
9. [Finish working with the wizard](#).

Step 1. Launch DynamoDB Restore Wizard

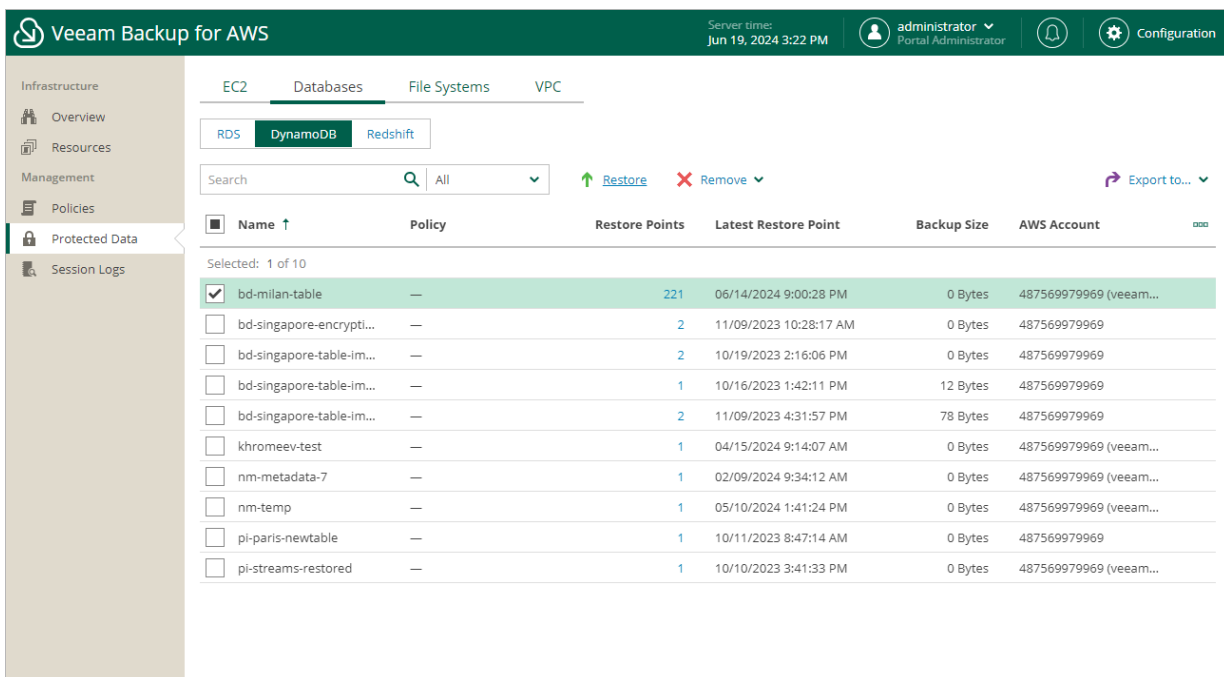
To launch the **DynamoDB Restore** wizard, do the following:

1. Navigate to **Protected Data > Databases > DynamoDB**.
2. Select the DynamoDB table that you want to restore.
3. Click **Restore**.

Alternatively, click the link in the **Restore Points** column. Then, in the **Available Restore Points** window, select the necessary restore point and click **Restore**.

NOTE

You can restore multiple DynamoDB tables if they belong to same AWS account only.



The screenshot shows the Veeam Backup for AWS interface. The top navigation bar includes the Veeam logo, server time (Jun 19, 2024 3:22 PM), user profile (administrator), and configuration options. The left sidebar shows navigation menus for Infrastructure, Management, and Protected Data. The main content area is titled 'Databases' and shows a list of DynamoDB tables. The table 'bd-milan-table' is selected, and its details are shown in the table below.

Name	Policy	Restore Points	Latest Restore Point	Backup Size	AWS Account
<input checked="" type="checkbox"/> bd-milan-table	—	221	06/14/2024 9:00:28 PM	0 Bytes	487569979969 (veeam...)
<input type="checkbox"/> bd-singapore-encrypti...	—	2	11/09/2023 10:28:17 AM	0 Bytes	487569979969
<input type="checkbox"/> bd-singapore-table-im...	—	2	10/19/2023 2:16:06 PM	0 Bytes	487569979969
<input type="checkbox"/> bd-singapore-table-im...	—	1	10/16/2023 1:42:11 PM	12 Bytes	487569979969
<input type="checkbox"/> bd-singapore-table-im...	—	2	11/09/2023 4:31:57 PM	78 Bytes	487569979969
<input type="checkbox"/> khromeev-test	—	1	04/15/2024 9:14:07 AM	0 Bytes	487569979969 (veeam...)
<input type="checkbox"/> nm-metadata-7	—	1	02/09/2024 9:34:12 AM	0 Bytes	487569979969 (veeam...)
<input type="checkbox"/> nm-temp	—	1	05/10/2024 1:41:24 PM	0 Bytes	487569979969 (veeam...)
<input type="checkbox"/> pi-paris-newtable	—	1	10/11/2023 8:47:14 AM	0 Bytes	487569979969
<input type="checkbox"/> pi-streams-restored	—	1	10/10/2023 3:41:33 PM	0 Bytes	487569979969 (veeam...)

Step 2. Select Restore Point

At the **Tables** step of the wizard, you can add tables to the restore session and select restore points to be used to perform the restore operation for each added DynamoDB table. By default, Veeam Backup for AWS uses the most recent valid restore point. However, you can restore a table to an earlier state.

To select a restore point, do the following:

1. Select the table and click **Restore Point**.
2. In the **Choose restore point** window, select the necessary restore point and click **Apply**.

To help you choose a restore point, Veeam Backup for AWS provides the following information on each available restore point:

- **Date** – the date when the restore point was created.
- **Size** – the size of the restore point.
- **Type** – the type of the restore point:
 - *DynamoDB backup* – an DynamoDB backup created by a backup policy.
 - *DynamoDB backup copy* – a backup copy created by a backup policy.
 - *Manual backup* – a DynamoDB backup created manually.
- **Storage Class** – the storage class of the restore point.
- **Restore Point Region** – the AWS Region where the restore point is stored.

IMPORTANT

Keep in mind that once stored in a cold storage tier in an AWS Region, backups cannot be copied to other AWS Regions. This means that you will only be able to use the backups to restore tables to the same AWS Region in which these backups reside after being moved from a warm storage tier. That is why if the selected restore points are stored in a cold storage tier in an AWS Region that differs from the AWS Region in which the backed-up tables reside, some of the [restore options](#) may not be available. To work around the issue, you can do either of the following:

- If you plan to perform restore to the original location, select restore points that are stored in a cold storage tier in the same AWS Region in which the backed-up tables reside.
- If you plan to perform restore either to a new location or to the original location but with different settings, select restore points that are stored in the target location.

The screenshot shows the Veeam Backup for AWS interface. At the top, the server time is Jun 19, 2024 3:23 PM, and the user is administrator (Portal Administrator). The main window is titled 'DynamoDB Table Restore'. On the left, there are tabs for 'Tables', 'Account', 'Restore Mode', 'Reason', and 'Summary'. The 'Specify tables to restore' section has a search box for 'Name' and a list of results, including 'bd-milan-ta...'. The 'Choose restore point' table is as follows:

Date ↓	Size	Type	Storage Class	Restore Point Region
06/02/2024 11:00:15 PM	24 Bytes	DynamoDB backup	Cold	Europe (Milan)
06/05/2024 10:00:19 PM	27 Bytes	DynamoDB backup	Cold	Europe (Milan)
06/05/2024 9:00:15 PM	27 Bytes	DynamoDB backup	Cold	Europe (Milan)
06/05/2024 8:00:19 PM	27 Bytes	DynamoDB backup	Cold	Europe (Milan)
06/05/2024 7:00:11 PM	27 Bytes	DynamoDB backup	Cold	Europe (Milan)
06/05/2024 6:00:11 PM	27 Bytes	DynamoDB backup	Cold	Europe (Milan)
06/05/2024 5:00:17 PM	24 Bytes	DynamoDB backup	Cold	Europe (Milan)
06/05/2024 4:00:10 PM	24 Bytes	DynamoDB backup	Cold	Europe (Milan)
06/05/2024 3:00:17 PM	24 Bytes	DynamoDB backup	Cold	Europe (Milan)
06/05/2024 2:00:12 PM	24 Bytes	DynamoDB backup	Cold	Europe (Milan)
05/27/2024 10:32:04 AM	32 Bytes	DynamoDB backup co...	Warm	Europe (Paris)
05/09/2024 3:07:00 PM	20 Bytes	DynamoDB backup co...	Warm	Europe (Paris)
04/16/2024 2:34:01 PM	24 Bytes	Manual backup	Warm	Europe (Milan)

At the bottom of the table, it shows 'Page 1 of 2' and 'Apply' and 'Cancel' buttons.

Step 3. Specify IAM Identity

At the **Account** step of the wizard, choose whether you want to [use an IAM role](#) or [one-time access keys of an IAM user](#) to allow Veeam Backup for AWS to perform the restore operation. For information on the permissions that the IAM role or IAM user must have to perform the restore operation, see [DynamoDB Restore IAM Permissions](#).

IMPORTANT

Make sure that the specified IAM role or one-time access keys belong to an AWS account where the source table resides.

Specifying IAM Role

To specify an IAM role for restore, select the **IAM role** option and choose the necessary IAM role from the list.

For an IAM role to be displayed in the **IAM role** list, it must be added to Veeam Backup for AWS with the *Amazon DynamoDB Restore* operation selected as described in section [Adding IAM Roles](#). If you have not added the necessary IAM role to Veeam Backup for AWS beforehand, you can do it without closing the **DynamoDB Table Restore** wizard. To do that, click **Add** and complete the [Add IAM Role](#) wizard.

IMPORTANT

It is recommended that you check whether the selected IAM role has all the permissions required to perform the operation. If some permissions of the IAM role are missing, the restore operation will fail to complete successfully. To run the IAM role permission check, click **Check Permissions** and follow the instructions provided in section [Checking IAM Role Permissions](#).

The screenshot shows the 'DynamoDB Table Restore' wizard in the 'Account' step. The interface includes a top navigation bar with the Veeam logo, server time (Jun 19, 2024 3:28 PM), user profile (administrator), and configuration icon. The main content area is titled 'DynamoDB Table Restore' and features a left-hand navigation menu with options: Tables, Account (selected), Restore Mode, Reason, and Summary. The 'Choose IAM role' section contains the instruction: 'Specify an IAM role that will be used to access resources for the restore operation, or provide temporary access keys.' There are two radio button options: 'IAM role' (selected) and 'Temporary access keys'. Under 'IAM role', a dropdown menu shows 'Default Backup Restore (Default Backup Restore)'. To the right of the dropdown are '+ Add' and 'Check Permissions' buttons. Under 'Temporary access keys', there are input fields for 'Access key:' and 'Secret key:'. A blue information icon is present next to a note: 'The keys are used to perform this operation only. They are not saved or stored. To learn what permissions are required for performing the operation, see the User Guide.' At the bottom of the wizard, there are three buttons: 'Previous', 'Next', and 'Cancel'.

Specifying One-Time Access Keys

To specify one-time access keys for restore, select the **Temporary access keys** option and use the **Access key** and **Secret key** fields to provide the access key ID and the secret access key.

NOTE

Veeam Backup for AWS does not store one-time access keys in the configuration database.

Veeam Backup for AWS Server time: Nov 3, 2023 9:19 AM administrator Portal Administrator Configuration

DynamoDB Table Restore

Tables

Account

Restore Mode

Reason

Summary

Choose IAM role

Specify an IAM role that will be used to access resources for the restore operation, or provide temporary access keys.

IAM role

Default Backup Restore (Default Backup Restore) + Add Check Permissions

Temporary access keys

Access key: AKIAY4ZWOU4WMVRAGEVN

Secret key:

i The keys are used to perform this operation only. They are not saved or stored. To learn what permissions are required for performing the operation, see the [User Guide](#).

Previous Next Cancel

Step 4. Choose Restore Mode

At the **Restore Mode** step of the wizard, choose whether you want to restore the selected DynamoDB table to the original or to a custom location. If you select the **Restore to a new location, or with different settings** option, specify the target AWS Region where the restored table will reside.

IMPORTANT

If any of the restore options are not available, make sure that the selected restore points meet all the requirements listed at [step 2](#).

Veeam Backup for AWS does not support restoring of provisioned throughput capacity values adjusted by Amazon DynamoDB auto scaling for tables and global secondary indexes (GSI). This means that if you add to the restore session a table with auto scaling enabled or a GSI-associated table with auto scaling enabled, the restore mode will affect the number of capacity units provisioned to the restored table or to the GSI:

- If you select the **Restore to original location** option, the restored table or GSI will be provisioned with the same numbers of capacity units that were used by the source table during the backup session. In this case, it is recommended to check values of capacity units for the restored table after the restore session completes to avoid unexpected charges.
- If you select the **Restore to new location, or with different settings** option, you will be able to specify the number of capacity units for the restored table at [step 7](#), which will apply both to the table and to the GSI.

TIP

If some of the selected tables still exist in AWS, the wizard will display a notification message and restore to the original location will not be available. To work around the issues, you can do either of the following:

- Remove the source tables from AWS.
- Use the **Restore to new location, or with different settings** option. In this case, you will also have to specify new names for the restored tables at [step 6](#).

The screenshot shows the 'DynamoDB Table Restore' wizard in Veeam Backup for AWS. The interface is in a dark theme. At the top, the header shows 'Veeam Backup for AWS', 'Server time: Jun 19, 2024 3:28 PM', and user information 'administrator Portal Administrator'. A navigation menu on the left includes 'Tables', 'Account', 'Restore Mode' (highlighted), 'Encryption', 'Settings', 'Capacity', 'Reason', and 'Summary'. The main content area is titled 'Restore Mode' and contains the instruction: 'Specify whether you want to restore the tables to the original location or to a new one, or with different settings.' There is an information icon and a message: 'Restore of Auto Scaling settings is not supported. To avoid unexpected expenses, we recommend to verify your capacity units configuration for tables and indexes which had Auto Scaling enabled. For more information see the User Guide.' Two radio button options are present: 'Restore to original location' (unselected) and 'Restore to new location, or with different settings' (selected). Below the selected option is a dropdown menu showing 'Asia Pacific (Sydney)'. At the bottom, there are three buttons: 'Previous', 'Next', and 'Cancel'.

Step 5. Enable Encryption

[This step applies only if you have selected the **Restore to new location, or with different settings** option at the **Restore Mode** step of the wizard]

At the **Encryption** step of the wizard, configure encryption settings:

- If you want to apply the existing encryption scheme, select the **Use original encryption scheme** option.
- If you want to change the key that is used for server-side encryption, select the **Change server-side encryption** option and choose the necessary key from the **Encryption key** drop-down list.

For a KMS key to be displayed in the list of available encryption keys, it must be stored in the AWS Region selected at [step 4](#) of the wizard, and the IAM role or user specified for the restore operation at [step 3](#) of the wizard must have permissions to access the key. For more information on KMS keys, see [AWS Documentation](#).

TIP

If the necessary KMS key is not displayed in the list, or if you want to use a KMS key from an AWS account other than the AWS account to which the specified IAM role belongs, you can select *Add custom key ARN* from the **Encryption key** drop-down list, and specify the Amazon resource name (ARN) of the key in the **Add Custom Key ARN** window.

For Veeam Backup for AWS to be able to encrypt the restored table using the provided KMS key, either the IAM role or user specified for the restore operation, or the IAM role used to create the restore point selected at [step 2](#) of the wizard must have permissions to access the key.

The screenshot shows the 'DynamoDB Table Restore' wizard in Veeam Backup for AWS. The interface is divided into a left sidebar and a main content area. The sidebar contains navigation options: Tables, Account, Restore Mode, Encryption (highlighted), Settings, Capacity, Reason, and Summary. The main content area is titled 'Configure encryption settings' and includes the instruction: 'Choose whether you want to use the original encryption scheme or encrypt the restored tables with a new key.' There are two radio button options: 'Use original encryption scheme' (unselected) and 'Change server-side encryption' (selected). Below the radio buttons is an 'Encryption key:' label and a dropdown menu currently showing 'Managed by DynamoDB'. An information icon (i) is present next to a link: 'To learn how to work with AWS encryption keys, see this Veeam KB article.' At the bottom of the wizard, there are three buttons: 'Previous' (disabled), 'Next' (active), and 'Cancel' (disabled).

Step 6. Configure General Settings

[This step applies only if you have selected the **Restore to new location, or with different settings** option at the **Restore Mode** step of the wizard]

At the **Settings** step of the wizard, you can specify a new name for the restored table. To do that, select the table and click **Rename**.

You can also choose a class for the restored table, decide whether you want to protect the table from accidental deletion, and enable point-in-time recovery to prevent accidental writes and to ensure restore to any point in time during the last 35 days. To specify the configuration settings, select the table and click **Edit**.

NOTE

By default, the AWS Backup service restores tables associated with the Standard table class only. To restore a table associated with the Standard-IA table class, Veeam Backup for AWS updates the table class of the restored table. Keep in mind that you can change table classes no more than two times during a 30-day period.

For more information on considerations and limitations when choosing a table class, see [AWS Documentation](#).

For more information on deletion protection, see [AWS Documentation](#). For more information on point-in-time recovery, see [AWS Documentation](#).

The screenshot shows the Veeam Backup for AWS interface. At the top, there is a dark green header with the Veeam logo, the text 'Veeam Backup for AWS', and the server time 'Jun 19, 2024 3:31 PM'. On the right side of the header, there is a user profile 'administrator Portal Administrator', a notification bell, and a 'Configuration' gear icon. Below the header, the main area is titled 'DynamoDB Table Restore'. On the left, there is a navigation menu with options: Tables, Account, Restore Mode, Encryption, Settings (highlighted), Capacity, Reason, and Summary. The main content area is divided into two panels. The left panel is titled 'Configure table settings' and contains a table with columns 'Name' and 'Table Class'. The table has one row with 'bd-milan-table' and 'Standard'. Above the table are 'Edit' and 'Rename' icons. The right panel is a modal dialog titled 'Configure table settings' with a close button. It contains a 'Table class' dropdown menu set to 'StandardIA', and two toggle switches for 'Deletion protection' and 'PITR backup', both of which are turned on. At the bottom of the modal are 'Apply' and 'Cancel' buttons.

Step 7. Choose Capacity Mode

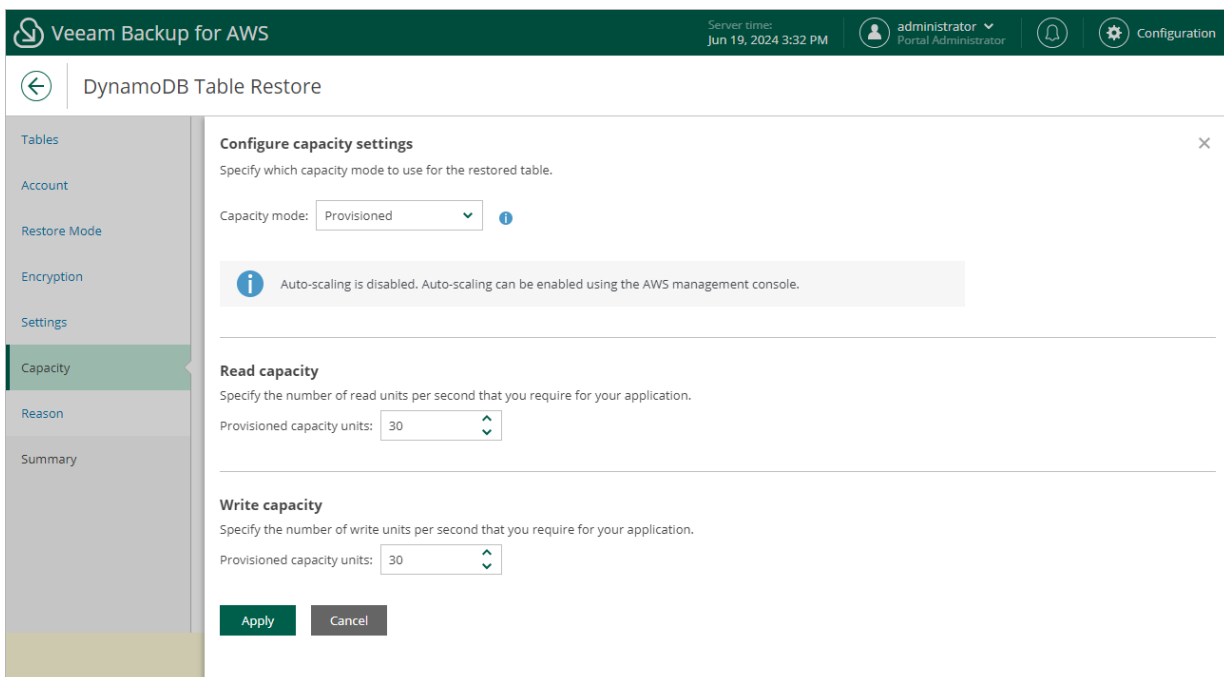
[This step applies only if you have selected the **Restore to new location, or with different settings** option at the **Restore Mode** step of the wizard]

At the **Capacity** step of the wizard, you can change the capacity mode and configure capacity settings for the restored table. To do that, select the table and click **Edit**.

NOTE

You can change the capacity mode only once within 24 hours. For more information on table capacity modes, see [AWS Documentation](#).

If you have selected the **Provisioned** capacity mode option, specify the value of the capacity units in the **Read capacity** and **Write capacity** fields. For more information on considerations and limitations when decreasing throughput for provisioned tables, see [AWS Documentation](#).



The screenshot shows the Veeam Backup for AWS interface. The top bar includes the Veeam logo, the text 'Veeam Backup for AWS', the server time 'Jun 19, 2024 3:32 PM', and user information 'administrator Portal Administrator'. The main window is titled 'DynamoDB Table Restore' and features a left-hand navigation menu with options: Tables, Account, Restore Mode, Encryption, Settings, Capacity (highlighted), Reason, and Summary. The main content area is titled 'Configure capacity settings' and contains the following elements:

- A heading 'Configure capacity settings' with a close button (X).
- Text: 'Specify which capacity mode to use for the restored table.'
- A dropdown menu for 'Capacity mode' set to 'Provisioned' with an information icon (i).
- An information message: 'Auto-scaling is disabled. Auto-scaling can be enabled using the AWS management console.'
- A heading 'Read capacity' with the text: 'Specify the number of read units per second that you require for your application.'
- A spinner input for 'Provisioned capacity units' set to '30'.
- A heading 'Write capacity' with the text: 'Specify the number of write units per second that you require for your application.'
- A spinner input for 'Provisioned capacity units' set to '30'.
- 'Apply' and 'Cancel' buttons at the bottom.

Step 8. Specify Restore Reason

At the **Reason** step of the wizard, you can specify a reason for restoring the DynamoDB table. The information you provide will be saved in the session history, and you can reference it later.

The screenshot shows the Veeam Backup for AWS interface. The top navigation bar includes the Veeam logo, the text "Veeam Backup for AWS", the server time "Jun 19, 2024 3:32 PM", the user "administrator Portal Administrator", and a "Configuration" link. The main content area is titled "DynamoDB Table Restore" and features a left-hand navigation menu with options: Tables, Account, Restore Mode, Encryption, Settings, Capacity, Reason (highlighted), and Summary. The "Reason" step is active, displaying the instruction "Specify the reason for performing the restore operation." and a text input field labeled "Restore reason:" containing the text "Restoring the table to Asia Pacific". At the bottom of the wizard, there are three buttons: "Previous", "Next" (with a mouse cursor over it), and "Cancel".

Step 9. Finish Working with Wizard

At the **Summary** step of the wizard, review summary information and click **Finish**.

The screenshot shows the Veeam Backup for AWS console interface. At the top, there is a dark green header with the Veeam logo and 'Veeam Backup for AWS' on the left. On the right, it displays 'Server time: Jun 19, 2024 3:33 PM', a user profile for 'administrator' (Portal Administrator), a notification bell, and a 'Configuration' gear icon. Below the header, a breadcrumb trail shows a back arrow and 'DynamoDB Table Restore'. A left-hand navigation menu lists several steps: 'Tables', 'Account', 'Restore Mode', 'Encryption', 'Settings', 'Capacity', 'Reason', and 'Summary'. The 'Summary' step is highlighted with a green bar. The main content area is titled 'Summary' and contains the following information:

- Summary**
Review the restore settings, and click Finish to exit the wizard.
- Reason**
Reason: Restoring the table to Asia Pacific
- General settings**
Restore mode: New location
Location Name: Asia Pacific (Sydney)
- IAM Role**
IAM Role name: Default Backup Restore
- Encryption settings**
Server-side encryption: Managed by DynamoDB

At the bottom of the wizard, there are three buttons: 'Previous' (disabled), 'Finish' (active), and 'Cancel' (disabled).

Redshift Restore

The actions that you can perform with restore points of Redshift clusters depend on whether you access the restore points using the Veeam Backup & Replication console or the Veeam Backup for AWS Web UI.

Redshift Restore Using Console

You can recover corrupted Redshift clusters in the Veeam Backup for AWS Web UI only. However, you can launch the **Redshift Cluster Restore** wizard directly from the Veeam Backup & Replication console to start the restore operation:


1. In the Veeam Backup & Replication console, open the **Home** view.
2. Navigate to **Backups > Snapshots**.
3. Expand the backup policy that protects the Redshift clusters that you want to recover, select the necessary cluster and click **Amazon Redshift** on the ribbon.

Alternatively, you can right-click the selected cluster and click **Restore to Amazon Redshift**.

IMPORTANT

You cannot restore multiple Redshift clusters from the Veeam Backup & Replication console.

Veeam Backup & Replication will open the **Redshift Cluster Restore** wizard in a web browser. Complete the wizard as described in section [Redshift Restore Using Web UI](#).



Redshift Restore Using Web UI

In case of a disaster, you can restore a Redshift cluster from a Redshift backup. Veeam Backup for AWS allows you to restore one or more Redshift clusters at a time to the original location, with the source or different settings. To learn how Redshift restore works, see [Redshift Restore](#).

IMPORTANT

Consider the following:

- Veeam Backup for AWS supports restore of Redshift clusters only to the same AWS accounts to which the source clusters belong and to the same AWS Region where the source cluster resides.
- Veeam Backup for AWS does not support restore of Redshift clusters with the Multi-AZ deployment. These clusters will be restored as clusters with the Single-AZ deployment.
- Veeam Backup for AWS does not support restore of Amazon Redshift Serverless.

How to Perform Redshift Restore

To restore a protected Redshift cluster, do the following:

1. [Launch the Redshift Restore wizard](#).
2. [Select a restore point](#).
3. [Specify an IAM identity for restore](#).
4. [Choose a restore mode](#).
5. [Enable encryption for the restored table](#).
6. [Configure Redshift cluster settings](#).
7. [Configure network settings](#).
8. [Specify a restore reason](#).
9. [Finish working with the wizard](#).

Step 1. Launch Redshift Restore Wizard

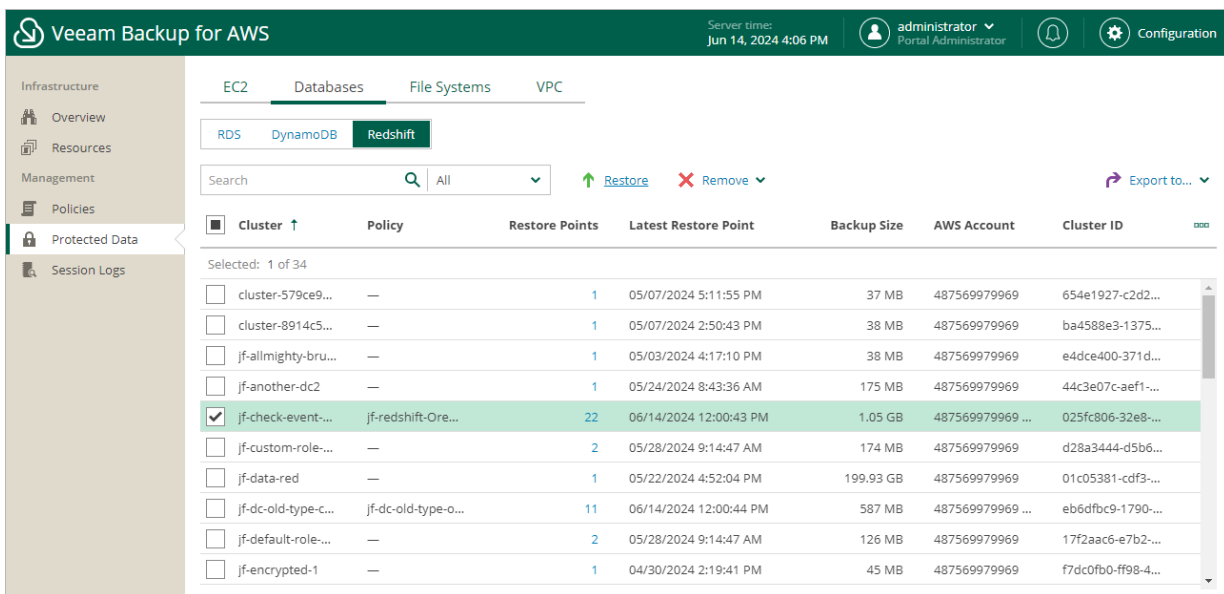
To launch the **Redshift Restore** wizard, do the following:

1. Navigate to **Protected Data > Databases > Redshift**.
2. Select the Redshift cluster that you want to restore.
3. Click **Restore**.

Alternatively, click the link in the **Restore Points** column. Then, in the **Available Restore Points** window, select the necessary restore point and click **Restore**.

NOTE

You can restore multiple Redshift clusters if they belong to same AWS account only.



The screenshot shows the Veeam Backup for AWS interface. The top navigation bar includes the Veeam logo, server time (Jun 14, 2024 4:06 PM), and user information (administrator, Portal Administrator). The left sidebar shows navigation options: Infrastructure, Overview, Resources, Management, Policies, Protected Data (selected), and Session Logs. The main content area is divided into tabs: EC2, Databases, File Systems, and VPC. Under Databases, there are sub-tabs for RDS, DynamoDB, and Redshift (selected). A search bar and a dropdown menu (All) are present. Below the search bar, there are buttons for Restore (up arrow), Remove (down arrow), and Export to... (right arrow). The main table displays a list of Redshift clusters with columns: Cluster, Policy, Restore Points, Latest Restore Point, Backup Size, AWS Account, and Cluster ID. One cluster, 'jf-check-event-...', is selected, indicated by a checkmark in the first column and a green highlight. The table shows 34 clusters in total, with 1 selected.

Cluster	Policy	Restore Points	Latest Restore Point	Backup Size	AWS Account	Cluster ID
<input type="checkbox"/> cluster-579ce9...	—	1	05/07/2024 5:11:55 PM	37 MB	487569979969	654e1927-c2d2...
<input type="checkbox"/> cluster-8914c5...	—	1	05/07/2024 2:50:43 PM	38 MB	487569979969	ba4588e3-1375...
<input type="checkbox"/> jf-allmighty-bru...	—	1	05/03/2024 4:17:10 PM	38 MB	487569979969	e4dce400-371d...
<input type="checkbox"/> jf-another-dc2	—	1	05/24/2024 8:43:36 AM	175 MB	487569979969	44c3e07c-aeff1...
<input checked="" type="checkbox"/> jf-check-event-...	jf-redshift-Ore...	22	06/14/2024 12:00:43 PM	1.05 GB	487569979969 ...	025fc806-32e8-...
<input type="checkbox"/> jf-custom-role-...	—	2	05/28/2024 9:14:47 AM	174 MB	487569979969	d28a3444-d5b6...
<input type="checkbox"/> jf-data-red	—	1	05/22/2024 4:52:04 PM	199.93 GB	487569979969	01c05381-cdf3...
<input type="checkbox"/> jf-dc-old-type-c...	jf-dc-old-type-o...	11	06/14/2024 12:00:44 PM	587 MB	487569979969 ...	eb6dfbc9-1790...
<input type="checkbox"/> jf-default-role-...	—	2	05/28/2024 9:14:47 AM	126 MB	487569979969	17f2aac6-e7b2...
<input type="checkbox"/> jf-encrypted-1	—	1	04/30/2024 2:19:41 PM	45 MB	487569979969	f7dc0fb0-ff98-4...

Step 2. Select Restore Point

At the **Clusters** step of the wizard, you can add clusters to the restore session and select restore points to be used to perform the restore operation for each added Redshift cluster. By default, Veeam Backup for AWS uses the most recent valid restore point. However, you can restore a cluster to an earlier state.

To select a restore point, do the following:

1. Select the cluster and click **Restore Point**.
2. In the **Choose restore point** window, select the necessary restore point and click **Apply**.

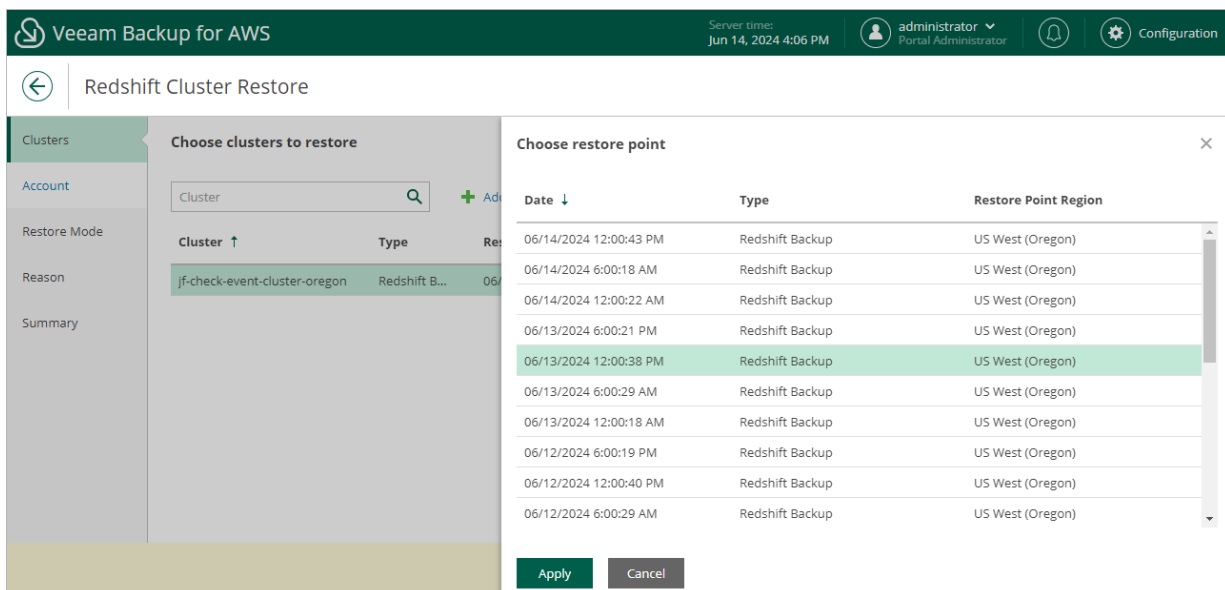
To help you choose a restore point, Veeam Backup for AWS provides the following information on each available restore point:

- **Date** – the date when the restore point was created.
- **Type** – the type of the restore point:
 - *Redshift backup* – a Redshift backup created by a backup policy.
 - *Manual backup* – a Redshift backup created manually.
- **Restore Point Region** – the AWS Region where the restore point is stored.

IMPORTANT

Consider the following:

- If you add Redshift clusters with the Multi-AZ deployment to the restore session, Veeam Backup for AWS will restore these clusters with the Single-AZ deployment.
-
- Since cross-region copying of Redshift backups is not supported for [Amazon Redshift](#), some of [restore options](#) may not be available. To work around the issue, it is recommended that you select restore points stored in the same AWS Region where the source clusters reside if you plan to perform restore to the original location but with different settings. Otherwise, Veeam Backup for AWS will be able to restore clusters belonging to different AWS Regions only to their original location with the source cluster settings.



The screenshot shows the Veeam Backup for AWS interface. The top bar includes the Veeam logo, the text 'Veeam Backup for AWS', the server time 'Jun 14, 2024 4:06 PM', and user information 'administrator Portal Administrator'. The main window is titled 'Redshift Cluster Restore' and has a sidebar with 'Clusters', 'Account', 'Restore Mode', 'Reason', and 'Summary'. The 'Clusters' section is active, showing 'Choose clusters to restore' with a search bar and a list of clusters. A 'Choose restore point' dialog box is open, displaying a table of restore points. The table has three columns: 'Date', 'Type', and 'Restore Point Region'. The selected restore point is highlighted in green.

Date ↓	Type	Restore Point Region
06/14/2024 12:00:43 PM	Redshift Backup	US West (Oregon)
06/14/2024 6:00:18 AM	Redshift Backup	US West (Oregon)
06/14/2024 12:00:22 AM	Redshift Backup	US West (Oregon)
06/13/2024 6:00:21 PM	Redshift Backup	US West (Oregon)
06/13/2024 12:00:38 PM	Redshift Backup	US West (Oregon)
06/13/2024 6:00:29 AM	Redshift Backup	US West (Oregon)
06/13/2024 12:00:18 AM	Redshift Backup	US West (Oregon)
06/12/2024 6:00:19 PM	Redshift Backup	US West (Oregon)
06/12/2024 12:00:40 PM	Redshift Backup	US West (Oregon)
06/12/2024 6:00:29 AM	Redshift Backup	US West (Oregon)

Step 3. Specify IAM Identity

At the **Account** step of the wizard, choose whether you want to [use an IAM role](#) or [one-time access keys of an IAM user](#) to allow Veeam Backup for AWS to perform the restore operation. For information on the permissions that the IAM role or IAM user must have to perform the restore operation, see [Redshift Restore IAM Permissions](#).

IMPORTANT

Make sure that the specified IAM role or one-time access keys belong to an AWS account where the source cluster resides.

Specifying IAM Role

To specify an IAM role for restore, select the **IAM role** option and choose the necessary IAM role from the list.

For an IAM role to be displayed in the **IAM role** list, it must be added to Veeam Backup for AWS with the *Amazon Redshift Restore* operation selected as described in section [Adding IAM Roles](#). If you have not added the necessary IAM role to Veeam Backup for AWS beforehand, you can do it without closing the **Redshift Cluster Restore** wizard. To do that, click **Add** and complete the [Add IAM Role](#) wizard.

IMPORTANT

It is recommended that you check whether the selected IAM role has all the permissions required to perform the operation. If some permissions of the IAM role are missing, the restore operation will fail to complete successfully. To run the IAM role permission check, click **Check Permissions** and follow the instructions provided in section [Checking IAM Role Permissions](#).

The screenshot shows the 'Redshift Cluster Restore' wizard in the 'Account' step. The interface includes a top navigation bar with 'Veeam Backup for AWS', a server time of 'Jun 14, 2024 4:12 PM', and user information 'administrator Portal Administrator'. The main content area is titled 'Choose IAM role' and contains the instruction: 'Specify an IAM role that will be used to access resources for restore operation or provide temporary access keys.' There are two radio button options: 'IAM role' (selected) and 'Temporary access keys'. Under 'IAM role', there is a dropdown menu showing 'Default Backup Restore (Default Backup Restore)', an '+ Add' button, and a 'Check Permissions' button. Under 'Temporary access keys', there are input fields for 'Access key:' and 'Secret key:'. An information icon and a note state: 'The keys are used to perform this operation only. They are not saved or stored. To learn what permissions are required for performing the operation, see the User Guide.' At the bottom, there are 'Previous', 'Next', and 'Cancel' buttons.

Specifying One-Time Access Keys

To specify one-time access keys for restore, select the **Temporary access keys** option and use the **Access key** and **Secret key** fields to provide the access key ID and the secret access key.

NOTE

Veeam Backup for AWS does not store one-time access keys in the configuration database.

Veeam Backup for AWS Server time: Jun 14, 2024 4:12 PM administrator Portal Administrator Configuration

Redshift Cluster Restore

Choose IAM role
Specify an IAM role that will be used to access resources for restore operation or provide temporary access keys.

IAM role
Default Backup Restore (Default Backup Restore) + Add Check Permissions

Temporary access keys

Access key: orhrt

Secret key:

i The keys are used to perform this operation only. They are not saved or stored. To learn what permissions are required for performing the operation, see the [User Guide](#).

Previous Next Cancel

Step 4. Choose Restore Mode

At the **Restore Mode** step of the wizard, choose whether you want to restore the selected Redshift cluster to the original location with the same or with different settings.

IMPORTANT

If any of the restore options are not available, make sure that the selected restore points meet all the requirements listed at [step 2](#).

The screenshot shows the 'Redshift Cluster Restore' wizard in Veeam Backup for AWS. The interface includes a top navigation bar with the Veeam logo, server time (Jun 14, 2024 4:13 PM), user information (administrator, Portal Administrator), and a Configuration icon. A left sidebar contains navigation links for Clusters, Account, Restore Mode (highlighted), Encryption, Settings, Network, Reason, and Summary. The main content area is titled 'Choose restore mode' and contains the following text: 'Specify whether you want to restore the cluster with the original or different settings'. There are two radio button options: 'Restore to original location with source settings' (unselected) and 'Restore to original location with different settings' (selected). Below the second option is a dropdown menu currently showing 'US West (Oregon)'. At the bottom of the wizard, there are three buttons: 'Previous', 'Next', and 'Cancel'.

Step 5. Enable Encryption

[This step applies only if you have selected the **Restore to original location with different settings** option at the **Restore Mode** step of the wizard]

At the **Encryption** step of the wizard, choose whether the restored Redshift cluster will be encrypted with an AWS KMS key:

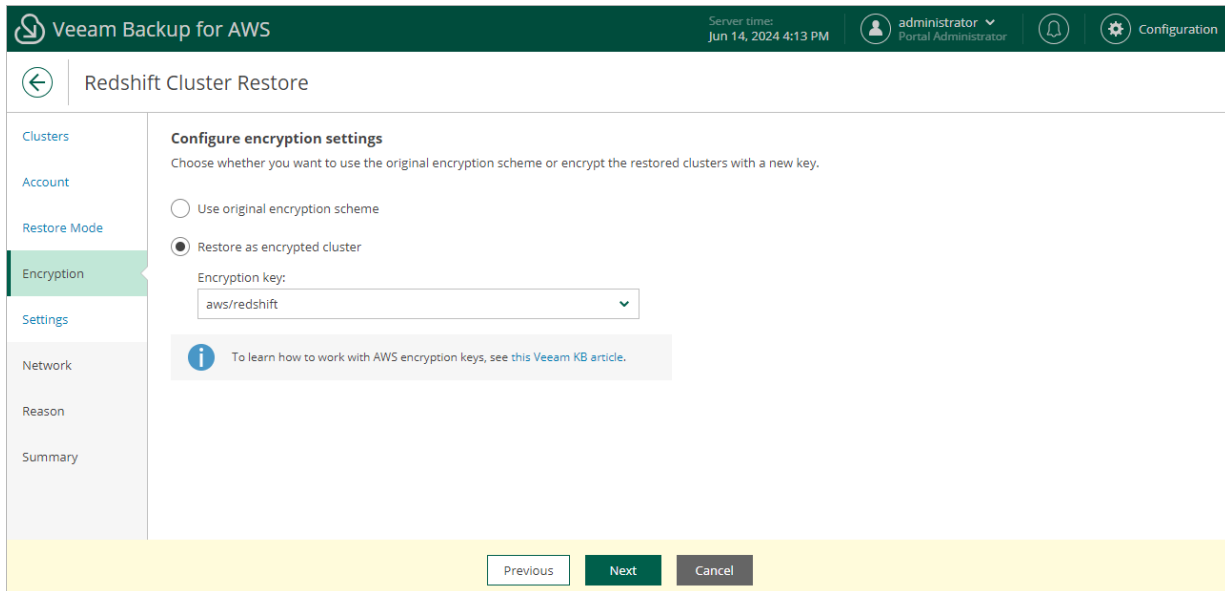
- If you want to apply the existing encryption scheme, select the **Use original encryption scheme** option.
- If you want to encrypt the Redshift cluster or want to change the key that is used for cluster encryption, select the **Restore as encrypted instance** option and choose the necessary key from the **Encryption key** drop-down list.

For a KMS key to be displayed in the list of available encryption keys, it must be stored in the AWS Region where the source Redshift cluster resides, and the IAM role or user specified for the restore operation at [step 3](#) of the wizard must have permissions to access the key. For more information on KMS keys, see [AWS Documentation](#).

TIP

If the necessary KMS key is not displayed in the list, or if you want to use a KMS key from an AWS account other than the AWS account to which the specified IAM role belongs, you can select *Add custom key ARN* from the **Encryption key** drop-down list, and specify the Amazon resource name (ARN) of the key in the **Add Custom Key ARN** window.

For Veeam Backup for AWS to be able to encrypt the restored cluster using the provided KMS key, either the IAM role (or user) specified for the restore operation or the IAM role used to create the restore point that was selected at [step 2](#) of the wizard must have permissions to access the key.



Step 6. Configure Restore Settings

[This step applies only if you have selected the **Restore to original location with different settings** option at the **Restore Mode** step of the wizard]

TIP

As soon as you proceed to the **Settings** step of the wizard, Veeam Backup for AWS will verify whether the original IAM roles associated with the Redshift cluster added to the restore session still exist in the AWS infrastructure. If the associated roles do not exist in the AWS infrastructure anymore, you will receive a warning in the **Associated IAM Roles** column. To work around the issue, select other IAM roles to be associated with the restored cluster.

You will also be able to proceed with the wizard and complete the restore operation without associating any new IAM roles. However, you will then need to associate the required roles with the cluster in the AWS Management Console as described in [AWS Documentation](#).

At the **Settings** step of the wizard, provide a new name and specify configuration settings for the restored Redshift cluster:

- To specify a new name, select the cluster and click **Rename**. Consider the following limitations:
 - The cluster identifier must be unique for each AWS Region within one AWS Account.
 - The cluster identifier can contain only lowercase Latin letters and hyphens, but cannot contain 2 consecutive hyphens.
 - The first character of the cluster identifier must be a letter. The last character of the identifier must not be a hyphen.
 - The maximum length of the cluster name is 63 characters.

For more information on limitations for Redshift identifiers, see [AWS Documentation](#).

- To specify configuration settings, select the cluster and click **Edit**. Then, in the **Cluster configuration** window, do the following:
 - a. To choose a node type for the restored cluster, use the **Node type** drop-down list. For more information on all existing cluster node types, see [AWS Documentation](#).
 - b. To choose a number of nodes for the restored cluster, use the **Number of nodes** drop-down list.

Note that the number of nodes that you can choose depends on the initial configuration of the source cluster. For more information on the node count dependency, see [AWS Documentation](#).
 - c. To choose a parameter group containing database parameter values that will be applied to the restored cluster, use the **Parameter group** drop-down list.

For a parameter group to be displayed in the list of available groups, the group must be created in the Amazon Redshift console as described in [AWS Documentation](#).
 - d. To associate IAM roles with the restored cluster or to replace the original IAM roles that are already associated with the cluster, set the **Configure IAM roles** toggle to *On*. Then, click the link in the **Associated IAM roles** field to select the necessary roles. If you set the toggle to *Off*, the cluster will be restored without any IAM role associated.

Note that the list shows all existing IAM roles from the same AWS account to which the restored cluster belongs.

- e. To set one of the selected IAM roles as the default one, use the **Default IAM role** drop-down list. For more information on default IAM roles in Amazon Redshift, see [AWS Documentation](#).

TIP

If the admin password that was used to access the source cluster is managed by AWS Secrets Manager, Veeam Backup for AWS will also require an AWS KMS key to encrypt the password for the restored cluster. You can either use the default KMS key or specify a new one – to use the default key, set the **Customize encryption settings** toggle to *Off*; to specify a new key, set the toggle to *On* and select the necessary key from the **Admin password** drop-down list.

For a KMS key to be displayed in the list of available encryption keys, it must be stored in the AWS Region where the source cluster resides, and the IAM role (or user) specified for the restore operation at [step 3](#) of the wizard must have permissions to access the key. If the necessary KMS key is not displayed in the list, or if you want to use a KMS key from an AWS account other than the AWS account to which the specified IAM role belongs, you can select *Add custom key ARN* from the **AWS KMS key** drop-down list, and specify the amazon resource number (ARN) of the key in the **Add Custom Key ARN** window. For more information on KMS keys, see [AWS Documentation](#).

The screenshot shows the Veeam Backup for AWS interface. At the top, there's a header with the Veeam logo, 'Veeam Backup for AWS', server time 'Jun 14, 2024 4:14 PM', and user 'administrator Portal Administrator'. The main content area is titled 'Redshift Cluster Restore'. On the left, a sidebar lists navigation options: Clusters, Account, Restore Mode, Encryption, Settings (highlighted), Network, Reason, and Summary. The central panel is titled 'Configure restore settings' and contains a table with columns 'Cluster' and 'Node Type'. The table has one row: 'jf-check-event-cluster-oregon' and 'dc2.large'. Above the table are 'Edit' and 'Rename' icons. The right panel is titled 'Cluster configuration' and contains the following settings: Node type: 'dc2.large'; Number of nodes: '1'; Parameter group: 'default.redshift-1.0'; Configure IAM roles: a toggle switch that is turned on; Associated IAM roles: '1 IAM role selected'; Default IAM role: 'AmazonRedshift-CommandsAccessRole-20231212...'; Admin password: 'Custom password'. At the bottom of the right panel are 'Apply' and 'Cancel' buttons.

Step 7. Configure Network Settings

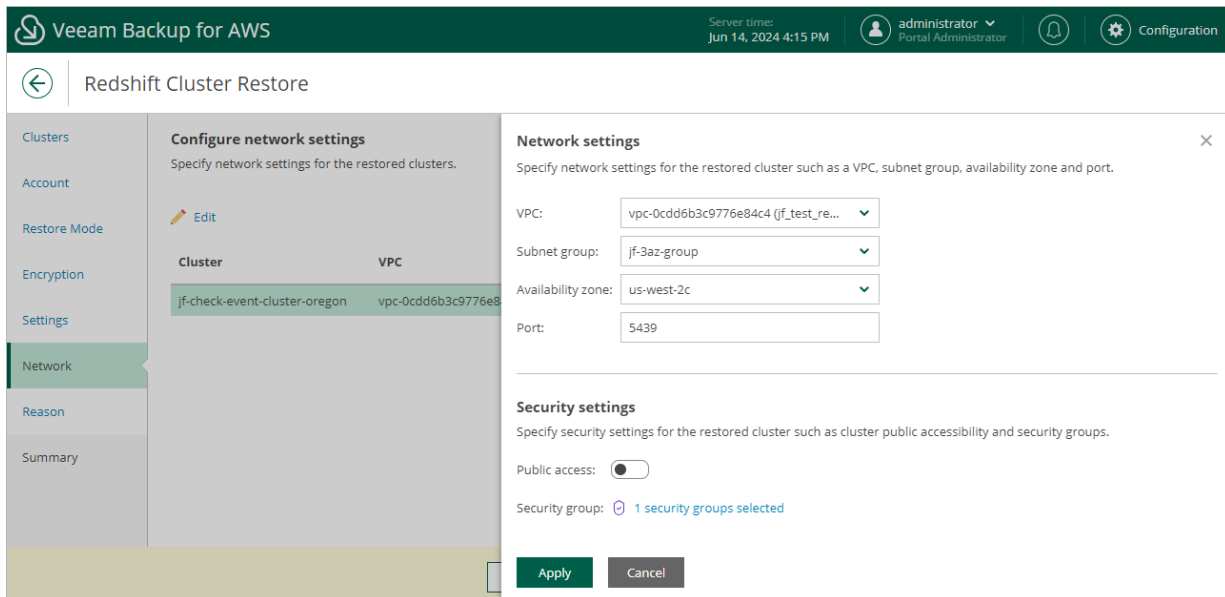
[This step applies only if you have selected the **Restore to original location with different settings** option at the **Restore Mode** step of the wizard]

At the **Network** step of the wizard, configure network and security settings for the restored Redshift cluster. To do that, select the necessary cluster and click **Edit**. In the opened window, do the following:

1. In the **Network settings** section, choose an Amazon VPC where the restored cluster will be deployed, a subnet group in which the cluster will be launched, an Availability Zone where the cluster will reside, and a port that will be used to access the cluster.

For a VPC and a subnet group to be displayed in the lists of available network specifications, they must be created in the AWS Region where the source cluster resides as described in [AWS Documentation](#).

2. In the **Security settings** section, specify security settings to control what IP addresses will be able to connect to databases in the restored cluster.
 - a. To make the cluster accessible outside the selected Amazon VPC, set the **Public accessible** toggle to *On*. Note that the cluster must belong to a public subnet group to become publicly accessible.
 - b. To specify security groups that will control access to the cluster, click the link in the **Security group** field and then select the necessary groups in the **Select Security Group** window.
3. To save changes made to the network and security settings, click **Apply**.



Step 8. Specify Restore Reason

At the **Reason** step of the wizard, you can specify a reason for restoring the Redshift cluster. The information you provide will be saved in the session history, and you can reference it later.

The screenshot shows the 'Redshift Cluster Restore' wizard in the Veeam Backup for AWS console. The interface is divided into a top navigation bar, a left sidebar, a main content area, and a bottom action bar.

- Top Bar:** Veeam Backup for AWS logo on the left. Server time: Jun 14, 2024 4:15 PM. User: administrator (Portal Administrator). Configuration icon on the right.
- Left Sidebar:** A list of steps: Clusters, Account, Restore Mode, Encryption, Settings, Network, Reason (highlighted in green), and Summary.
- Main Content Area:**
 - Restore reason:** Specify a reason for performing the restore operation.
 - Restore reason:** A text input field containing the text 'corrupted redshift cluster'.
- Bottom Bar:** Three buttons: 'Previous' (disabled), 'Next' (active), and 'Cancel' (disabled).

Step 9. Finish Working with Wizard

At the **Summary** step of the wizard, review summary information and click **Finish**.

The screenshot displays the 'Redshift Cluster Restore' wizard in the 'Summary' step. The top navigation bar includes the Veeam logo, the text 'Veeam Backup for AWS', the server time 'Jun 14, 2024 4:16 PM', the user 'administrator', and a 'Configuration' link. The left sidebar lists various settings categories, with 'Summary' highlighted. The main content area shows the following summary information:

- Summary**
Review the restore settings, and click Finish to exit the wizard.
- [Copy to Clipboard](#)
- Reason**
Reason: corrupted redshift cluster
- General settings**
Restore mode: Original location with different settings
Location name: US West (Oregon)
- IAM role**
IAM role: Default Backup Restore
- Encryption settings**
Encryption: Restore as an encrypted cluster

At the bottom of the wizard, there are three buttons: 'Previous', 'Finish', and 'Cancel'.

EFS Restore

The actions that you can perform with restore points of EFS file systems depend on whether you access the restore points using the Veeam Backup & Replication console or the Veeam Backup for AWS Web UI.

EFS Restore Using Console

Veeam Backup & Replication offers the following restore operations:

- [File system restore](#) – restore an entire Amazon EFS file system.
- [File-level recovery](#) – restore individual files and folders stored in a file system.

You can restore EFS file system data to the most recent state or to any available restore point.

IMPORTANT

You can restore an EFS file system only to the same AWS account to which the source file system belongs.

Performing Entire File System Restore

In case a disaster strikes, you can restore an entire Amazon EFS file system from an EFS backup or a backup copy. Veeam Backup & Replication allows you to restore one or more Amazon EFS file systems at a time, to the original location or to a new location. To learn how EFS restore works, see [EFS Restore](#).

How to Perform EFS File-Level Recovery

To restore a protected EFS file system, do the following:

1. [Launch the Restore to Amazon EFS wizard.](#)
2. [Select a restore point.](#)
3. [Choose a restore mode.](#)
4. [Select an AWS Region.](#)
5. [Configure restore settings.](#)
6. [Specify a new name for the file system.](#)
7. [Configure network and mount target settings.](#)
8. [Specify a restore reason.](#)
9. [Finish working with the wizard.](#)

Step 1. Launch Restore to Amazon EFS Wizard


To launch the **Restore to Amazon EFS** wizard, do the following:

1. In the Veeam Backup & Replication console, open the **Home** view.
2. Navigate to **Backups > Snapshots**.
3. In the working area, expand the backup policy that protects an EFS file system that you want to restore, select the necessary EFS file system and click **Amazon EFS** on the ribbon.

Alternatively, you can right-click the file system and select **Restore to Amazon EFS**.

TIP

You can also launch the **Restore to Amazon EFS** wizard from the **Home** tab. To do that, click **Restore** and select **AWS**. In the **Restore** window, select **Amazon EFS**.



Step 2. Select Restore Point

At the **EFS File System** step of the wizard, choose a restore point that will be used to restore the selected Amazon EFS file system. By default, Veeam Backup & Replication uses the most recent valid restore point. However, you can restore the EFS file system data to an earlier state.

To select a restore point, do the following:

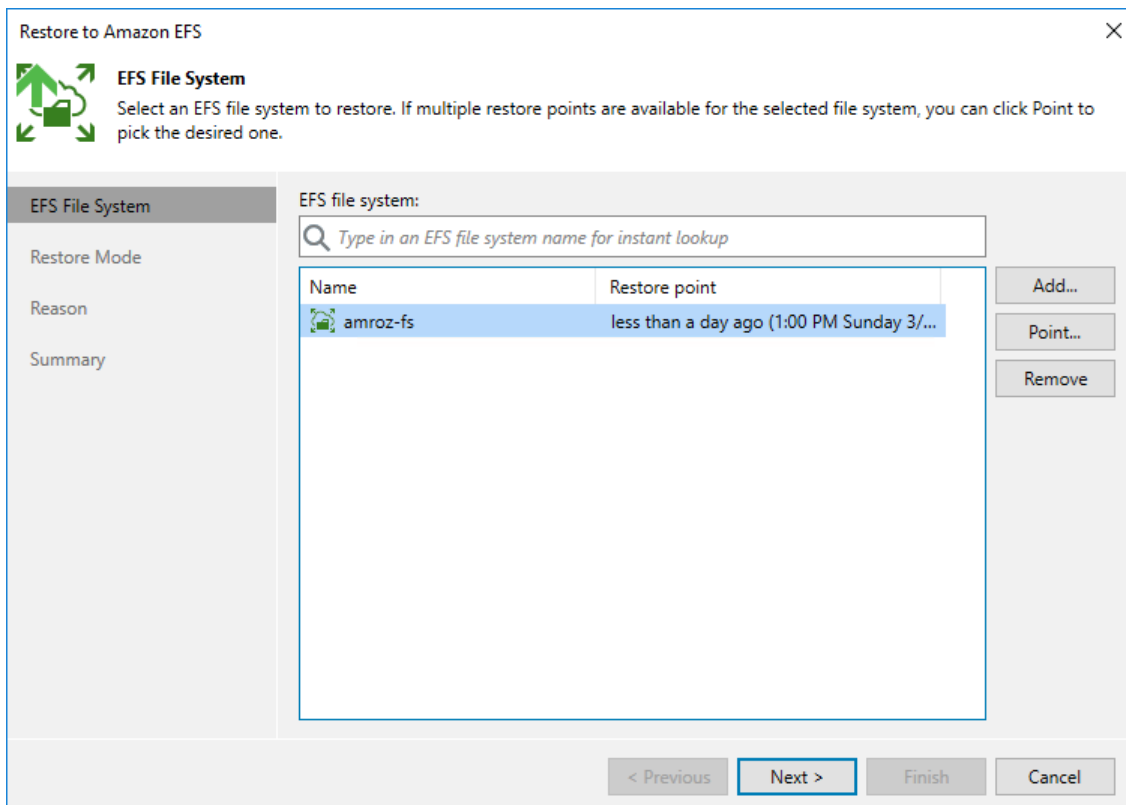
1. In the **EFS file system** list, select the EFS file system and click **Point**.
2. In the **Restore Points** window, expand the backup policy that protects the EFS file system, select the necessary restore point and click **OK**.

To help you choose a restore point, Veeam Backup & Replication provides the following information on each available restore point:

- **Job** – the name of the backup policy that created the restore point and the date when the restore point was created.
- **Type** – the type of the restore point.
- **Location** – the AWS Region or repository where the restore point is stored.

TIP

You can use the wizard to restore multiple file systems at a time. To do that, click **Add**, select more EFS file systems to restore and choose a restore point for each of them.



Step 3. Choose Restore Mode

At the **Restore Mode** step of the wizard, do the following:

1. Choose whether you want to restore the EFS file system to the original or to a new location.

NOTE

If you choose to restore to the original location, consider the following:

- The original EFS file system will be removed as soon as the restore process completes.
- The restored file system will not be mounted to any EC2 instances automatically. To mount the file system to an EC2 instance, you must do it manually in AWS as described in [AWS Documentation](#).

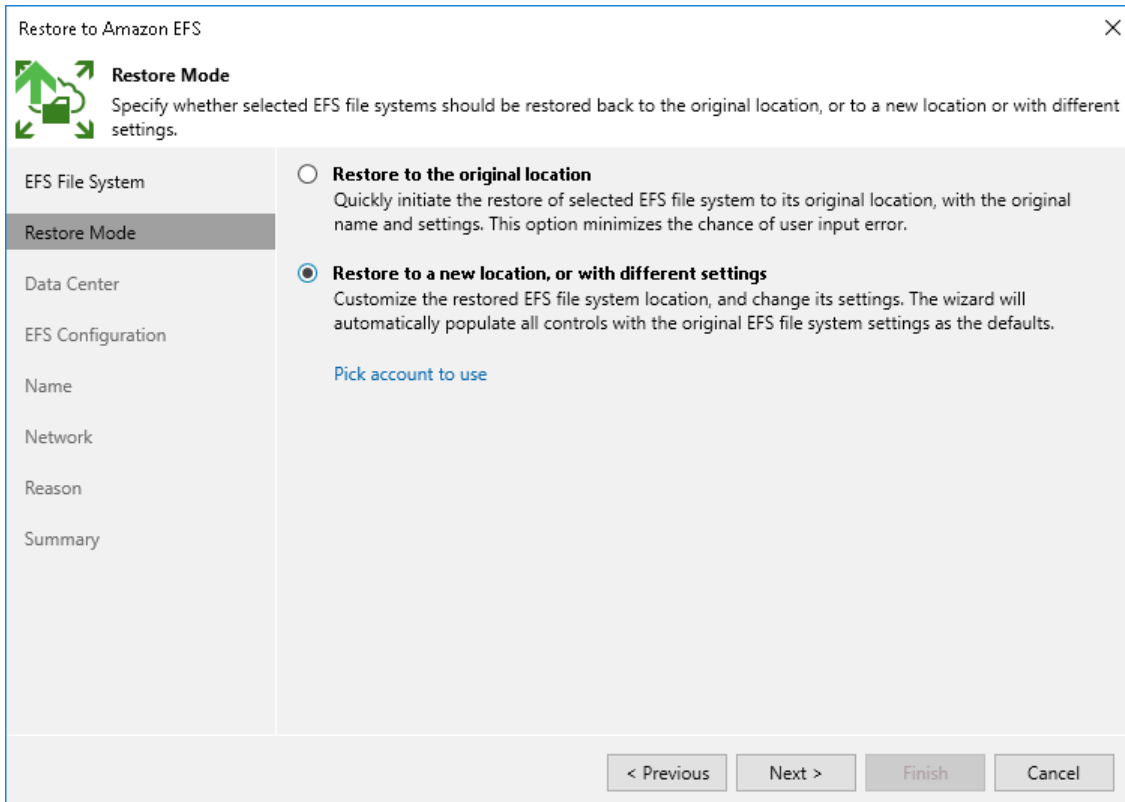
2. Click **Pick an IAM identity to use** to select an IAM identity whose permissions will be used to perform the restore operation:
 - To specify an IAM role for the restore operation, select the **IAM role** option and choose the necessary IAM role from the **IAM role** drop-down list.

For an IAM role to be displayed in the list of available roles, it must be added to the backup appliance as described in section [Adding IAM Roles](#).
 - To specify access keys of an IAM user, select the **Temporary access key** option, and use the **Access key** and **Secret key** fields to provide the access key ID and the secret access key.

NOTE

By default, to perform the restore operation, Veeam Backup & Replication uses permissions of either the *Default Backup Restore* IAM role, or the IAM role that has been used to protect the source EC2 instance, or the IAM role used to update information on restore points created for the instance while rescanning AWS infrastructure.

The *Default Backup Restore* IAM role is assigned all the permissions required to perform data protection and disaster recovery operations in the same AWS account where the backup appliance resides. For more information on the *Default Backup Restore* IAM role permissions, see [Full List of IAM Permissions](#).



The screenshot shows a wizard window titled "Restore to Amazon EFS" with a close button (X) in the top right corner. On the left is a navigation pane with the following items: "EFS File System", "Restore Mode" (highlighted), "Data Center", "EFS Configuration", "Name", "Network", "Reason", and "Summary".

The main content area is titled "Restore Mode" and includes a green icon of a folder with arrows. Below the icon is the text: "Specify whether selected EFS file systems should be restored back to the original location, or to a new location or with different settings."

There are two radio button options:

- Restore to the original location**
Quickly initiate the restore of selected EFS file system to its original location, with the original name and settings. This option minimizes the chance of user input error.
- Restore to a new location, or with different settings**
Customize the restored EFS file system location, and change its settings. The wizard will automatically populate all controls with the original EFS file system settings as the defaults.
[Pick account to use](#)

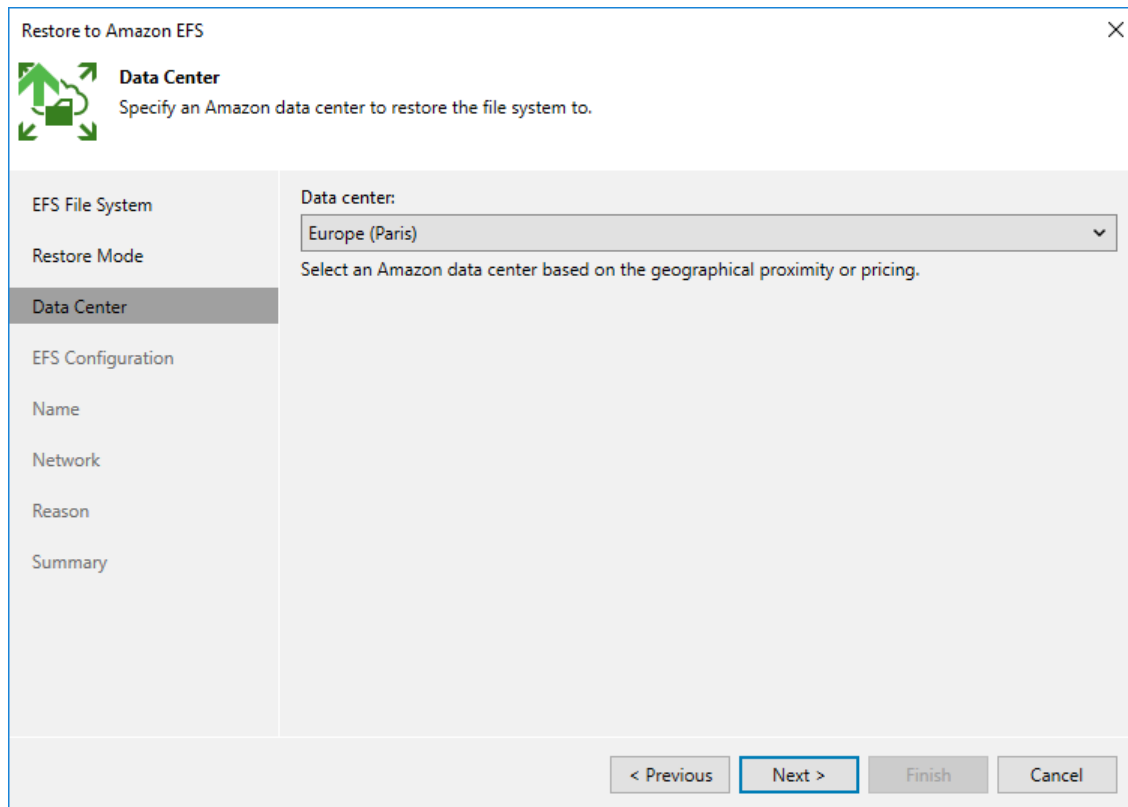
At the bottom of the window are four buttons: "< Previous", "Next >", "Finish", and "Cancel".

Step 4. Select Region

[This step applies only if you have selected the **Restore to a new location, or with different settings** option at the **Restore Mode** step of the wizard]

At the **Data Center** step of the wizard, select an AWS Region where the restored EFS file system will reside.

If the selected location differs from the original location of the EFS file system, Veeam Backup & Replication will raise a warning notifying that the locations do not match. Click **Yes** to acknowledge the warning. Otherwise, you will not be able to proceed with the wizard.



The screenshot shows a wizard window titled "Restore to Amazon EFS". The current step is "Data Center", which is highlighted in the left-hand navigation pane. The main area of the wizard contains the following elements:

- Data Center**: Specify an Amazon data center to restore the file system to.
- Data center:** A dropdown menu with "Europe (Paris)" selected.
- Instruction:** Select an Amazon data center based on the geographical proximity or pricing.

At the bottom of the wizard, there are four buttons: "< Previous", "Next >" (which is highlighted with a blue border), "Finish", and "Cancel".


Step 5. Configure Restore Settings

[This step applies only if you have selected the **Restore to a new location, or with different settings** option at the **Restore Mode** step of the wizard]

At the **EFS Configuration** step of the wizard, you can change the configuration and encryption settings for the restored file system. To do that, select the file system and do the following:

1. Click **Redundancy**. Then, in the **Redundancy Settings** window:
 - a. Choose whether you want to redundantly store data of the restored file system across all Availability Zones within the selected AWS Region (*Regional*), or within a single Availability Zone (*One Zone*).
For more information on storage options, see [AWS Documentation](#)
 - b. [This step applies only if you have selected the **Regional** option] From the **Performance mode** drop-down list, choose whether the restored file system will use the General Purpose or Max I/O performance mode.
For more information on performance modes, see [AWS Documentation](#).
 - c. [This step applies only if you have selected the **One Zone** option] From the **Availability zone** drop-down list, select an Availability Zone where the restored file system will be located.
2. Click **Encryption**. Then, in the **File system encryption** window:
 - o Select the **Preserve the original encryption settings** option if you do not want to encrypt the file system or want to apply the existing encryption scheme.
 - o Select the **Use the following encryption password** option if you want to encrypt the file system with an AWS KMS key. Then, choose the necessary KMS key from the list.

For a KMS key to be displayed in the list of available encryption keys, it must be stored in the AWS Region selected at [step 4](#) of the wizard, and the IAM role specified for the restore operation must have permissions to access the key. For more information on KMS keys, see [AWS Documentation](#).




Step 6. Specify File System Name

[This step applies only if you have selected the **Restore to a new location, or with different settings** option at the **Restore Mode** step of the wizard]

At the **Name** step of the wizard, you can specify a new name for the restored EFS file system.

TIP

You can specify a single prefix or suffix and add it to the names of multiple restored EFS file systems. To do that, select the necessary file systems and click **Name**. In the **Change Name** window, select the **Add prefix** or **Add suffix** check box, and provide the text that you want to add. Then, click **OK**.



Step 7. Configure Network Settings

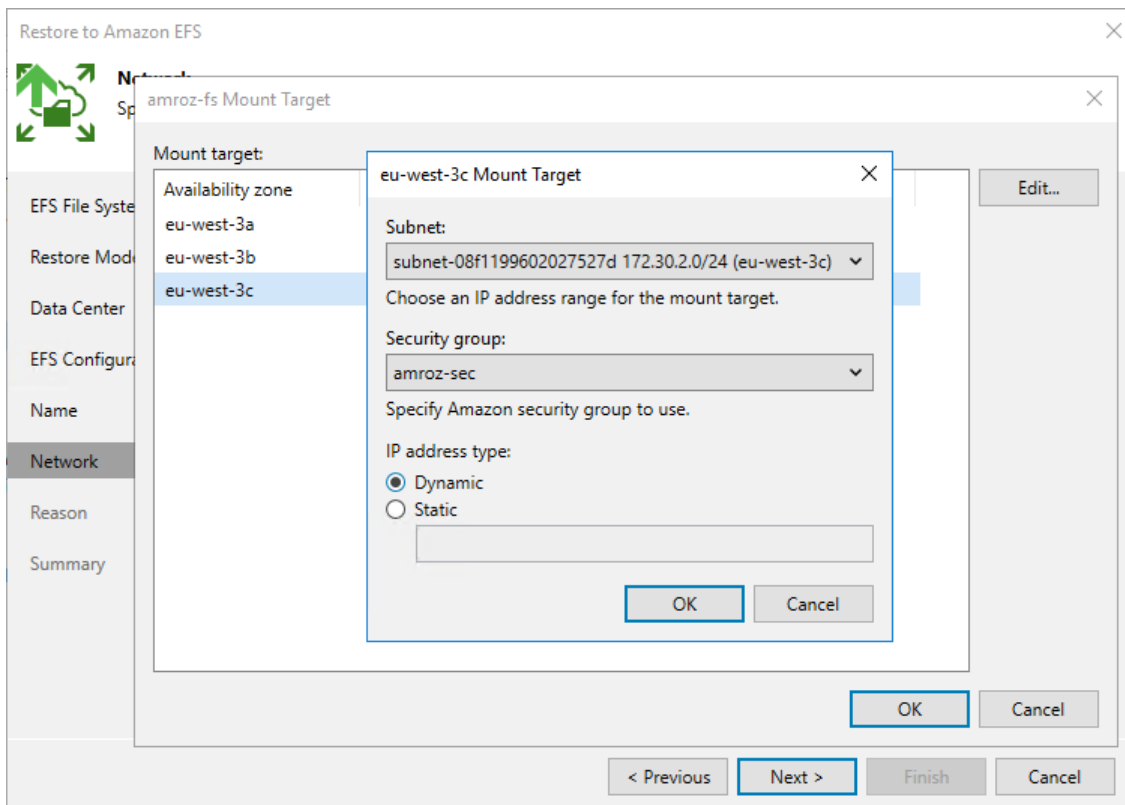
[This step applies only if you have selected the **Restore to a new location, or with different settings** option at the **Restore Mode** step of the wizard]

At the **Network** step of the wizard, you can configure specific network and mount target settings for the restored file system. To do that, select the file system and do the following:

1. Click **VPC** and select VPC to which the restored EFS file system will be connected.
For a VPC to be displayed in the list of available VPC networks, it must be created in AWS in the AWS Region specified at [step 4](#) of the wizard, as described in [AWS Documentation](#).
2. Click **Target**, select an Availability Zone where the mount target will be created and click **Edit**. Then, in the **Mount Target** window:
 - a. From the **Subnet** drop-down list, select a subnet to which the mount target will be connected.
For a subnet to be displayed in the list of available networks, it must be created in AWS as described in [AWS Documentation](#).
 - b. From the **Security group** drop-down list, select a security group that will be associated with the mount target.
For a security group to be displayed in the list of available groups, it must be created in AWS as described in [AWS Documentation](#).
 - c. In the **IP address type** section, choose whether you want Veeam Backup & Replication to assign a dynamic IP address to the mount target.

NOTE

If you have selected the *Regional* storage class at [step 5](#) of the wizard, it is required to configure at least one mount target for the restored EFS file system.



Step 8. Specify Restore Reason

At the **Reason** step of the wizard, specify a reason for restoring the Amazon EFS file system. The information you provide will be saved in the session history and you can reference it later.

Restore to Amazon EFS

Reason
Type in the reason for performing this restore operation. This information will be logged in the restore sessions history for later reference.

EFS File System
Restore Mode
Data Center
EFS Configuration
Name
Network
Reason
Summary


Restore reason:
Restore EFS

Do not show me this page again

< Previous Next > Finish Cancel

Step 9. Finish Working with Wizard

At the **Summary** step of the wizard, review summary information and click **Finish**.




Performing EFS File-Level Restore

You can perform EFS file-level restore only using the Veeam Backup for AWS Web UI. However, you can launch the EFS file-level recovery wizard directly from the Veeam Backup & Replication console. To do that, do the following:

1. In the Veeam Backup & Replication console, open the **Home** view.
2. Navigate to **Backups > Snapshots**.
3. Expand the EFS backup policy that protects a file system whose files and folders you want to restore, select the necessary file system and click **Amazon EFS file** on the ribbon.

Alternatively, you can right-click the file system and select **Restore to Amazon EFS files**.

Veeam Backup & Replication will open the **EFS File-level Recovery** wizard in a web browser. Complete the wizard as described in section [Performing File-Level Recovery](#).



EFS Restore Using Web UI

Veeam Backup for AWS offers the following restore options:

- [File system restore](#) – restores an entire Amazon EFS file system.
- [File-level recovery](#) – recovers individual files and folders stored in a file system.

You can restore EFS file system data to the most recent state or to any available restore point.

IMPORTANT

You can restore an EFS file system only to the same AWS account to which the source file system belongs.

Performing Entire File System Restore

In case of a disaster, you can restore an entire EFS file system from an EFS backup or backup copy. Veeam Backup for AWS allows you to restore one or more EFS file systems at a time, to the original location or to a new location.

How to Perform File System Restore

To restore a protected EFS file system, do the following:

1. [Launch the EFS Restore wizard.](#)
2. [Select a restore point.](#)
3. [Specify an IAM identity for restore.](#)
4. [Choose a restore mode.](#)
5. [Enable encryption for the restored file system.](#)
6. [Specify configuration settings.](#)
7. [Configure network settings.](#)
8. [Specify a restore reason.](#)
9. [Finish working with the wizard.](#)

Step 1. Launch EFS Restore Wizard

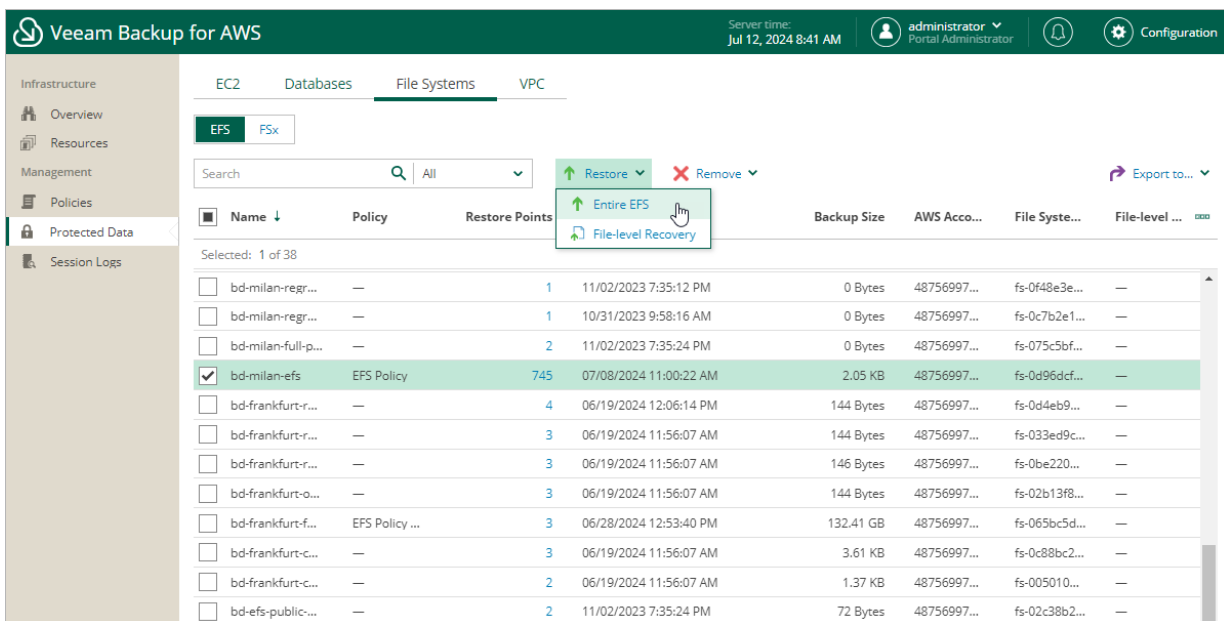
To launch the **EFS Restore** wizard, do the following:

1. Navigate to **Protected Data > File Systems > EFS**.
2. Select the EFS file system that you want to restore.
3. Click **Restore > Entire EFS**.

Alternatively, click the link in the **Restore Points** column. Then, in the **Available Restore Points** window, select the necessary restore point and click **Restore > Entire EFS**.

NOTE

You can restore multiple EFS file systems if they belong to same AWS account only.



The screenshot shows the Veeam Backup for AWS console interface. The top navigation bar includes the Veeam logo, 'Veeam Backup for AWS', server time 'Jul 12, 2024 8:41 AM', and user 'administrator'. The left sidebar shows navigation options like 'Infrastructure', 'Overview', 'Resources', 'Management', 'Policies', 'Protected Data', and 'Session Logs'. The main content area is titled 'File Systems' and shows a list of EFS file systems. The 'bd-milan-efs' file system is selected, and the 'Restore' dropdown menu is open, highlighting the 'Entire EFS' option.

Name	Policy	Restore Points	Backup Size	AWS Acco...	File System...	File-level ...
bd-milan-regr...	—	1	0 Bytes	48756997...	fs-0f48e3e...	—
bd-milan-regr...	—	1	0 Bytes	48756997...	fs-0c7b2e1...	—
bd-milan-full-p...	—	2	0 Bytes	48756997...	fs-075c5bf...	—
bd-milan-efs	EFS Policy	745	2.05 KB	48756997...	fs-0d96dcf...	—
bd-frankfurt-r...	—	4	144 Bytes	48756997...	fs-0d4eb9...	—
bd-frankfurt-r...	—	3	144 Bytes	48756997...	fs-033ed9c...	—
bd-frankfurt-r...	—	3	146 Bytes	48756997...	fs-0be220...	—
bd-frankfurt-o...	—	3	144 Bytes	48756997...	fs-02b13f8...	—
bd-frankfurt-f...	EFS Policy ...	3	132.41 GB	48756997...	fs-065bc5d...	—
bd-frankfurt-c...	—	3	3.61 KB	48756997...	fs-0c88bc2...	—
bd-frankfurt-c...	—	2	1.37 KB	48756997...	fs-005010...	—
bd-efs-public...	—	2	72 Bytes	48756997...	fs-02c38b2...	—

Step 2. Select Restore Point

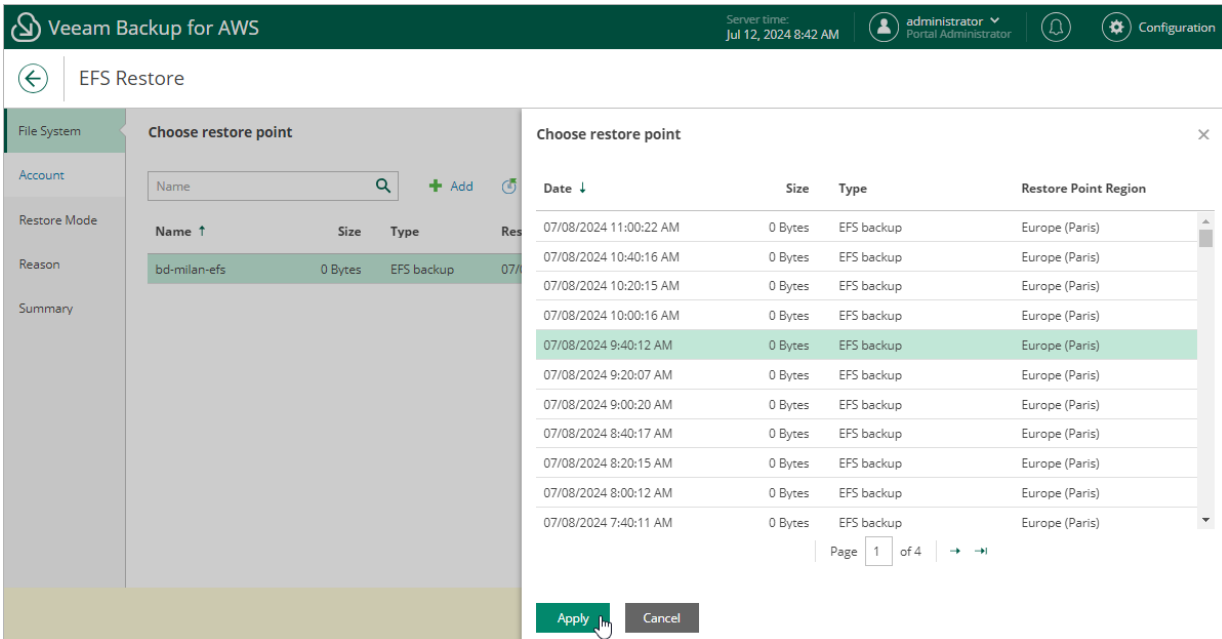
At the **File System** step of the wizard, you can add file systems to the restore session and select restore points to be used to perform the restore operation for each added EFS file system. By default, Veeam Backup for AWS uses the most recent valid restore point. However, you can restore a file system to an earlier state.

To select a restore point, do the following:

1. Select the EFS system and click **Restore Point**.
2. In the **Choose restore point** window, select the necessary restore point and click **Apply**.

To help you choose a restore point, Veeam Backup for AWS provides the following information on each available restore point:

- **Date** – the date when the restore point was created.
- **Size** – the size of the restore point.
- **Type** – the type of the restore point:
 - *EFS backup* – an EFS backup created by a backup policy.
 - *EFS backup copy* – a backup copy created by a backup policy.
 - *Manual backup* – an EFS backup created manually.
- **Restore Point Region** – the AWS Region where the restore point is stored.



The screenshot shows the Veeam Backup for AWS interface. The top bar displays the Veeam logo, the text 'Veeam Backup for AWS', the server time 'Jul 12, 2024 8:42 AM', and the user 'administrator Portal Administrator'. The main window is titled 'EFS Restore' and shows a 'Choose restore point' dialog box. The dialog box contains a table of restore points with the following columns: Date, Size, Type, and Restore Point Region. The selected restore point is 07/08/2024 9:40:12 AM, 0 Bytes, EFS backup, Europe (Paris). The 'Apply' button is highlighted.

Date ↓	Size	Type	Restore Point Region
07/08/2024 11:00:22 AM	0 Bytes	EFS backup	Europe (Paris)
07/08/2024 10:40:16 AM	0 Bytes	EFS backup	Europe (Paris)
07/08/2024 10:20:15 AM	0 Bytes	EFS backup	Europe (Paris)
07/08/2024 10:00:16 AM	0 Bytes	EFS backup	Europe (Paris)
07/08/2024 9:40:12 AM	0 Bytes	EFS backup	Europe (Paris)
07/08/2024 9:20:07 AM	0 Bytes	EFS backup	Europe (Paris)
07/08/2024 9:00:20 AM	0 Bytes	EFS backup	Europe (Paris)
07/08/2024 8:40:17 AM	0 Bytes	EFS backup	Europe (Paris)
07/08/2024 8:20:15 AM	0 Bytes	EFS backup	Europe (Paris)
07/08/2024 8:00:12 AM	0 Bytes	EFS backup	Europe (Paris)
07/08/2024 7:40:11 AM	0 Bytes	EFS backup	Europe (Paris)

Step 3. Specify IAM Identity

At the **Account** step of the wizard, choose whether you want to [use an IAM role](#) or [one-time access keys of an IAM user](#) to allow Veeam Backup for AWS to perform the restore operation. For information on the permissions that the IAM role or IAM user must have to perform the restore operation, see [EFS Restore IAM Permissions](#).

IMPORTANT

Make sure that the specified IAM role or one-time access keys belong to an AWS account where the source file system resides.

Specifying IAM Role

To specify an IAM role for restore, select the **IAM role** option and choose the necessary IAM role from the list.

For an IAM role to be displayed in the **IAM role** list, it must be added to Veeam Backup for AWS with the *Amazon EFS Restore* operation selected as described in section [Adding IAM Roles](#). If you have not added the necessary IAM role to Veeam Backup for AWS beforehand, you can do it without closing the **EFS Restore** wizard. To do that, click **Add** and complete the [Add IAM Role](#) wizard.

IMPORTANT

It is recommended that you check whether the selected IAM role has all the permissions required to perform the operation. If some permissions of the IAM role are missing, the restore operation will fail to complete successfully. To run the IAM role permission check, click **Check Permissions** and follow the instructions provided in section [Checking IAM Role Permissions](#).

The screenshot shows the 'EFS Restore' wizard in the 'Account' step. The interface includes a top navigation bar with the Veeam logo, server time (Jul 12, 2024 8:43 AM), and user information (administrator, Portal Administrator). The main content area is titled 'Choose IAM role' and contains the following elements:

- A left-hand navigation pane with tabs for 'File System', 'Account' (selected), 'Restore Mode', 'Reason', and 'Summary'.
- A heading 'Choose IAM role' followed by a descriptive text: 'Specify an IAM role that will be used to access resources for the restore operation or provide temporary access keys. The selected IAM role must belong to an AWS account where the source file system resides.'
- Two radio button options: 'IAM role' (selected) and 'Temporary access keys'.
- Under 'IAM role', a dropdown menu shows 'Default Backup Restore (Default Backup Restore)'. To its right are '+ Add' and 'Check Permissions' buttons.
- Under 'Temporary access keys', there are two text input fields labeled 'Access key:' and 'Secret key:'.
- An information icon (i) with a text box: 'The keys are used to perform this operation only. They are not saved or stored. To learn what permissions are required for performing the operation, see the [User Guide](#).'
- A bottom navigation bar with 'Previous', 'Next', and 'Cancel' buttons.

Specifying One-Time Access Keys

To specify one-time access keys for restore, select the **Temporary access keys** option and use the **Access key** and **Secret key** fields to provide the access key ID and the secret access key.

NOTE

Veeam Backup for AWS does not store one-time access keys in the configuration database.

Veeam Backup for AWS Server time: Jul 12, 2024 8:44 AM administrator Portal Administrator Configuration

EFS Restore

File System

Account

Restore Mode

Reason

Summary

Choose IAM role

Specify an IAM role that will be used to access resources for the restore operation or provide temporary access keys. The selected IAM role must belong to an AWS account where the source file system resides.

IAM role

Default Backup Restore (Default Backup Restore) + Add Check Permissions

Temporary access keys

Access key: ALDOIF18IDDUUWEU44656

Secret key:

i The keys are used to perform this operation only. They are not saved or stored. To learn what permissions are required for performing the operation, see the [User Guide](#).

Previous Next Cancel

Step 4. Choose Restore Mode

At the **Restore Mode** step of the wizard, choose whether you want to restore the selected EFS file system to the original or to a custom location. If you select the **Restore to a new location, or with different settings** option, specify the target AWS Region where the restored file system will reside.

The screenshot shows the 'EFS Restore' wizard in Veeam Backup for AWS. The interface includes a top navigation bar with the Veeam logo, 'Veeam Backup for AWS', server time 'Jul 12, 2024 8:45 AM', and user 'administrator Portal Administrator'. A left sidebar lists steps: File System, Account, Restore Mode (highlighted), Encryption, Settings, Network, Reason, and Summary. The main area is titled 'Choose restore mode' and contains the instruction: 'Specify whether you want to restore the file system to the original location or to a new one, or with different settings.' Two radio button options are present: 'Restore to original location' (unselected) and 'Restore to new location, or with different settings' (selected). Below the selected option is a dropdown menu showing 'Asia Pacific (Mumbai)'. At the bottom, there are three buttons: 'Previous', 'Next', and 'Cancel'.

Step 5. Enable Encryption

[This step applies only if you have selected the **Restore to new location, or with different settings** option at the **Restore Mode** step of the wizard]

At the **Encryption** step of the wizard, choose whether the restored file system will be encrypted with AWS KMS keys:

- If you do not want to encrypt the file system or want to apply the existing encryption scheme, select the **Use original encryption scheme** option.
- If you want to encrypt the file system, select the **Restore as encrypted file system** option and choose the necessary KMS key from the **Encryption key** drop-down list.

For a KMS key to be displayed in the list of available encryption keys, it must be stored in the AWS Region selected at [step 4](#) of the wizard, and the IAM role or user specified for the restore operation at [step 3](#) of the wizard must have permissions to access the key. For more information on KMS keys, see [AWS Documentation](#).

TIP

If the necessary KMS key is not displayed in the list, or if you want to use a KMS key from an AWS account other than the AWS account to which the specified IAM role belongs, you can select *Add custom key ARN* from the **Encryption key** drop-down list, and specify the amazon resource name (ARN) of the key in the **Add Custom Key ARN** window.

For Veeam Backup for AWS to be able to encrypt the restored file system using the provided KMS key, either the IAM role or user specified for the restore operation, or the IAM role used to create the restore point selected at [step 2](#) of the wizard must have permissions to access the key.

The screenshot shows the 'EFS Restore' wizard in Veeam Backup for AWS. The interface is in a dark theme. At the top, there's a header with the Veeam logo, 'Veeam Backup for AWS', server time 'Jul 12, 2024 8:45 AM', and user 'administrator Portal Administrator'. A navigation pane on the left lists steps: File System, Account, Restore Mode, Encryption (highlighted), Settings, Network, Reason, and Summary. The main area is titled 'Configure encryption settings' and contains the instruction: 'Choose whether you want to use the original encryption scheme or encrypt the restored file systems with a new key.' There are two radio buttons: 'Use original encryption scheme' (unselected) and 'Restore as encrypted file system' (selected). Below the radio buttons is an 'Encryption key:' dropdown menu with 'aws/elasticfilesystem' selected. An information icon with a link to a Veeam KB article is also present. At the bottom, there are three buttons: 'Previous', 'Next', and 'Cancel'.

Step 6. Configure General Settings

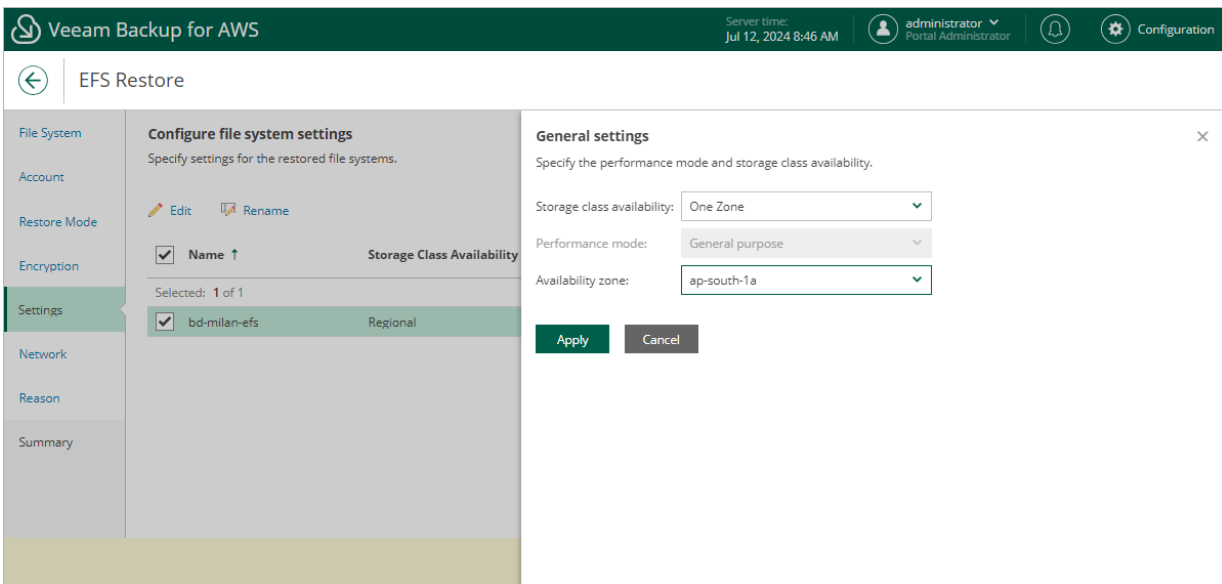
[This step applies only if you have selected the **Restore to new location, or with different settings** option at the **Restore Mode** step of the wizard]

At the **Settings** step of the wizard, you can specify new names and configuration settings for the restored file system.

To specify a new name, select the file system and click **Rename**. In the **File system name** window, specify the name and click **Apply**.

To specify configuration settings, do the following:

1. Select the file system and click **Edit**.
2. In the **General Settings** window, do the following:
 - a. From the **Storage class availability** drop-down list, select one of the following options:
 - *Regional* – if you want to redundantly store data of the restored file system across all Availability Zones within the selected AWS Region.
 - *One Zone* – if you want to redundantly store data of the restored file system within a single Availability Zone.
 - b. [Applies if you have selected the *Regional* option] From the **Performance mode** drop-down list, select a performance mode for the restored file system. For more information on performance modes, see [AWS Documentation](#).
 - c. [Applies if you have selected the *One Zone* option] From the **Availability zone** drop-down list, select an Availability Zone where the restored file system will be located.
3. To save changes made to the file system settings, click **Apply**.



Step 7. Configure Network Settings

[This step applies only if you have selected the **Restore to new location, or with different settings** option at the **Restore Mode** step of the wizard]

At the **Network** step of the wizard, configure network and mount target settings for the restored file system.

Choose Virtual Private Cloud

Specify an Amazon VPC to which the restored EFS file system must be connected:


1. In the **Network** section, click **Edit Network Settings**.
2. In the **Network specifications** window, select the necessary Amazon VPC.
For a VPC to be displayed in the **VPC** list, it must be created in the AWS Region specified at [step 4](#) of the wizard as described in [AWS Documentation](#).
3. Click **Apply**.

Configure Mount Targets

Configure settings for mount targets that will be created for the restored file system:

1. Click the link in the **Mount targets** section.
2. In the **Mount targets specification** window, click **Add**.
3. In the **Add Mount Target** window, do the following:
 - a. From the **Availability zone** drop-down list, select an Availability Zone where the mount target will be created.
 - b. From the **Subnet** drop-down list, select a subnet to which the mount target will be connected.
For a subnet to be displayed in the **Subnet** list, it must be created for the selected Availability Zone in the specified VPC as described in [AWS Documentation](#).
 - c. In the **IP address** section, choose one of the following options:
 - *Automatic* – if you want an IP address to be automatically assigned to the mount target.
 - *Static* – if you want to specify a static IP address for the mount target.
 - d. Add security groups to control inbound and outbound access to the restored file system. To do that, from the **Security groups** drop-down list, select a security group that will be associated with the mount target and click **Add**. Note that you cannot add more than 5 security groups.
For a security group to be displayed in the Security groups list, it must be created in the AWS Management Console as described in [AWS Documentation](#).
 - e. To save the mount target configuration, click **Add**.

4. To save the changes made to the mount target settings, click **Apply**.



Step 8. Specify Restore Reason

At the **Reason** step of the wizard, specify a reason for restoring the EFS file system. The information you provide will be saved in the session history and you can reference it later.

The screenshot shows the Veeam Backup for AWS interface. The top navigation bar includes the Veeam logo, the text "Veeam Backup for AWS", the server time "Jul 12, 2024 8:49 AM", the user "administrator Portal Administrator", and a "Configuration" link. The main content area is titled "EFS Restore" and features a left-hand navigation menu with options: File System, Account, Restore Mode, Encryption, Settings, Network, Reason (highlighted), and Summary. The "Restore reason" section contains the instruction "Specify a reason for performing the restore operation." and a text input field with the value "Restoring corrupted file system". At the bottom right, there are three buttons: "Previous", "Next", and "Cancel".

Step 9. Finish Working with Wizard

At the **Summary** step of the wizard, review summary information and click **Finish**.

The screenshot shows the 'EFS Restore' wizard in the 'Summary' step. The interface includes a top navigation bar with the Veeam logo, server time (Jul 12, 2024 8:50 AM), and user information (administrator, Portal Administrator). A left sidebar lists navigation options: File System, Account, Restore Mode, Encryption, Settings, Network, Reason, and Summary (highlighted). The main content area is titled 'Review configured settings' and contains the following information:

- Reason:** Restoring corrupted file system
- General settings:**
 - Restore mode: New location
 - Location name: Asia Pacific (Mumbai)
- IAM role:**
 - IAM role name: Default Backup Restore
- Encryption settings:**
 - Encryption: Encrypted file system
 - Encryption key: aws/elasticfilesystem

At the bottom right, there are three buttons: 'Previous' (disabled), 'Finish' (active), and 'Cancel' (disabled).

Performing File-Level Restore

In case a disaster strikes, you can recover corrupted or missing files of an EFS file system from an EFS backup or backup copy. Veeam Backup for AWS allows you to restore files and folders to the original file system or to another file system.

How to Perform EFS File-Level Recover

To recover files and folders of a protected file system, do the following:

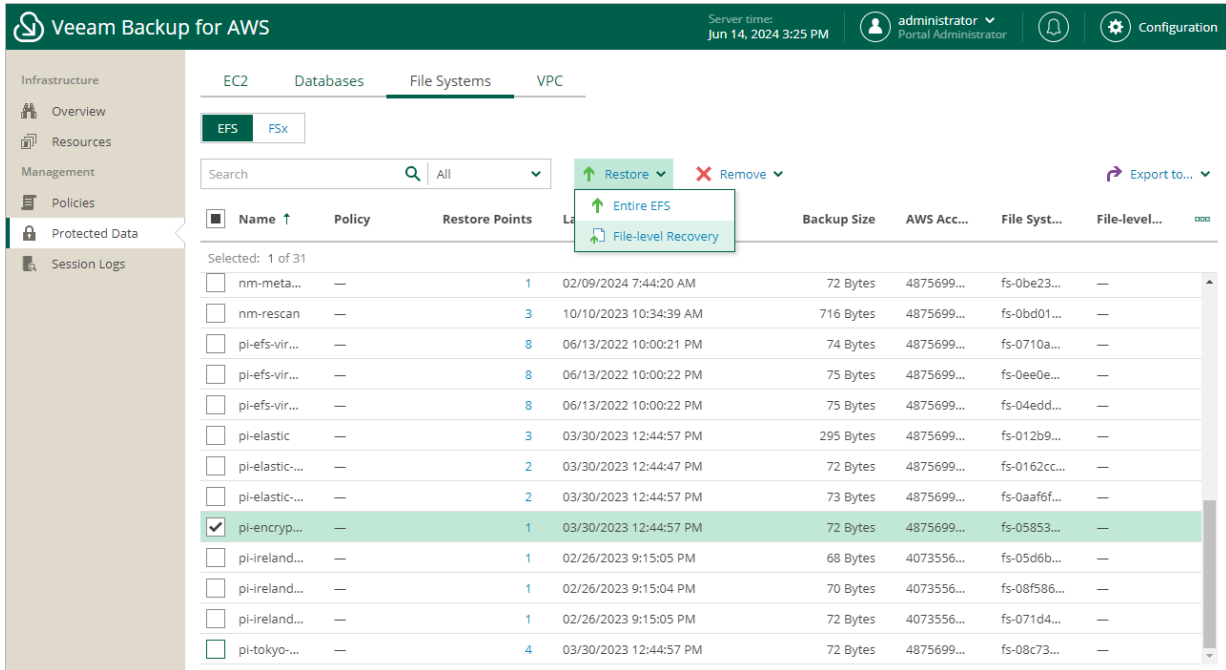
1. [Launch the EFS File-level Recovery wizard.](#)
2. [Choose a restore type.](#)
3. [Configure restore settings.](#)
4. [Specify an IAM identity for restore.](#)
5. [Choose a restore mode.](#)
6. [Specify a restore reason.](#)
7. [Finish working with the wizard.](#)
8. [Open the file-level recovery browser.](#)
9. [Select a restore point.](#)
10. [Choose files and folders to recover.](#)
11. [Stop the recovery session.](#)

Step 1. Launch EFS File-level Recovery Wizard

To launch the **EFS File-level Recovery** wizard, do the following:

1. Navigate to **Protected Data > File Systems > EFS**.
2. Select the file system whose files and folders you want to recover, and click **Restore > File-level Recovery**.

Alternatively, click the link in the **Restore Points** column. Then, in the **Available Restore Points** window, select the necessary restore point and click **Restore > File-level Recovery**.



Step 2. Choose Restore Type

At the **Restore Type** step of the wizard, choose whether you want to specify the exact paths to files and folders that you want to recover, or to select specific files and folders in the file-level recovery browser.

IMPORTANT

If you select the **Browse files** option, Veeam Backup for AWS will launch the EFS FLR session after you complete the **EFS File-level Recovery** wizard. Depending on the number of files stored in the file system, this session can consume up to 4 GB of RAM on the backup appliance.

The screenshot shows the Veeam Backup for AWS interface. The top navigation bar includes the Veeam logo, the text "Veeam Backup for AWS", the server time "Jun 14, 2024 3:26 PM", the user "administrator Portal Administrator", and a "Configuration" link. The main content area is titled "EFS File-level Recovery" and contains a sidebar with navigation options: "Restore Type" (selected), "Restore List", "Account", "Restore Mode", "Reason", and "Summary". The "Choose restore type" section is active, with the instruction "Specify whether you want to select files from an index or restore files using the exact file paths." Two radio button options are present: "Browse files" (unselected) and "Specify file paths" (selected). The "Specify file paths" option has a sub-instruction: "Specify direct paths to the files that will be restored to the file system." At the bottom of the form, there are "Next" and "Cancel" buttons.

Step 3. Configure Restore Settings

[This step applies only if you have selected the **Specify file paths** option at the **Restore Type** step of the wizard]

At the **Restore List** step of the wizard, do the following:

1. [Specify a restore point that will be used to restore the selected items.](#)
2. [Specify files and folders that you want to recover.](#)

Step 2.1 Select Restore Point

By default, Veeam Backup for AWS uses the most recent valid restore point. However, you can restore files and folders to an earlier state.

To select a restore point:

1. In the **Restore point** section of the **Restore List** step of the wizard, click the link to the right of **Restore point**.
2. In the **Choose restore point** window, select the necessary restore point and click **Apply**.

To help you choose a restore point, Veeam Backup for AWS provides the following information on each available restore point:

- **Date** – the date when the restore point was created.
- **Size** – the size of the restore point.
- **Type** – the type of the restore point:
 - *EFS backup* – an EFS backup created by a backup policy.
 - *EFS backup copy* – a backup copy created by a backup policy.
 - *Manual backup* – an EFS backup created manually.
- **Restore Point Region** – an AWS Region where the restore point is stored.

The screenshot shows the Veeam Backup for AWS interface. The top navigation bar includes the Veeam logo, 'Veeam Backup for AWS', server time 'Jun 14, 2024 3:26 PM', user 'administrator Portal Administrator', and a 'Configuration' link. The main content area is titled 'EFS File-level Recovery'. On the left, a sidebar contains 'Restore Type', 'Restore List', 'Account', 'Restore Mode', 'Reason', and 'Summary'. The 'Restore List' section is active, showing 'Choose restore point and items to restore'. The 'Restore point' section displays 'Restore point: 11/16/2023 11:01:23 AM'. The 'Restore list' section shows 'No items added yet' with an 'Edit' link. The 'Choose restore point' dialog box is open, displaying a table of restore points:

Date ↓	Size	Type	Restore Point Region
11/16/2023 11:01:23 AM	0 Bytes	EFS backup	Asia Pacific (Singapore)
11/16/2023 11:01:23 AM	0 Bytes	EFS backup copy	Asia Pacific (Sydney)
11/15/2023 4:35:15 PM	0 Bytes	EFS backup	Asia Pacific (Singapore)
11/15/2023 3:44:00 PM	31.75 KB	EFS backup	Asia Pacific (Singapore)
11/14/2023 5:54:26 PM	529.72 KB	EFS backup	Asia Pacific (Singapore)
11/14/2023 1:52:55 PM	0 Bytes	EFS backup	Asia Pacific (Singapore)
11/14/2023 1:52:55 PM	0 Bytes	EFS backup	Asia Pacific (Singapore)
11/14/2023 1:04:20 PM	519.25 KB	Manual backup	Asia Pacific (Mumbai)
11/14/2023 12:02:40 PM	0 Bytes	EFS backup	Asia Pacific (Singapore)
11/14/2023 12:02:40 PM	0 Bytes	EFS backup copy	Asia Pacific (Seoul)
11/14/2023 10:13:25 AM	0 Bytes	EFS backup	Asia Pacific (Singapore)

At the bottom of the dialog, there are 'Apply' and 'Cancel' buttons.

Step 2.2 Specify Items to Restore

To add files and folders to the restore list:

1. In the **Restore list** section, click **Edit**.
2. In the **Edit restore list** window, do the following:
 - a. For each file or folder you want to recover, specify a path in the **Item path** field and click **Add**. Note that you cannot add more than 5 paths.

Paths are case sensitive and cannot contain wild cards and regex strings. The following characters are not supported: ? * : " < > ` .

NOTE

The specified paths must be related to the mount point of the file system. For example, if the file system is mounted to the `/user/mydocs/efs` point and the file path is `/user/mydocs/efs/file1`, specify `/file1`.

- b. Review the restore list and click **Apply**.

The screenshot displays the Veeam Backup for AWS interface for EFS File-level Recovery. The top navigation bar includes the Veeam logo, 'Veeam Backup for AWS', server time 'Jun 14, 2024 3:26 PM', user 'administrator Portal Administrator', and a 'Configuration' button. The main window is titled 'EFS File-level Recovery' and features a sidebar with options: Restore Type, Restore List (selected), Account, Restore Mode, Reason, and Summary. The 'Restore List' section shows 'Choose restore point and items to restore' with a 'Restore point' of '11/15/2023 4:35:15 PM' and an 'Items' section with 'No items added yet' and an 'Edit' button. The 'Edit restore list' modal window is open, providing instructions: 'To restore a specific file, specify the file path related to the mount point. For example, if the file system is mounted to /user/mydocs/efs and the file path is /user/mydocs/efs/filename, specify /filename. Paths are case sensitive and cannot contain special characters, wild cards and regex strings.' It includes an 'Item path' input field with '/ortiz/accounts' and an '+ Add' button, a 'Remove from Restore List' button, a checkbox for 'Item', and a list of selected items showing 'Selected: 0 of 1' and one item: '/ortiz/balance.docx'. At the bottom of the modal are 'Apply' and 'Cancel' buttons.

Step 4. Specify IAM Identity

At the **Account** step of the wizard, choose whether you want to [use an IAM role](#) or [one-time access keys of an IAM user](#) to allow Veeam Backup for AWS to perform the restore operation. For information on the permissions that the IAM role or IAM user must have to perform the restore operation, see [EFS Restore IAM Permissions](#).

IMPORTANT

Make sure that the specified IAM role or one-time access keys belong to an AWS account where the source file system resides.

Specifying IAM Role

To specify an IAM role for restore, select the **IAM role** option and choose the necessary IAM role from the list.

For an IAM role to be displayed in the **IAM role** list, it must be added to Veeam Backup for AWS with the *Amazon EFS Restore* operation selected as described in section [Adding IAM Roles](#). If you have not added the necessary IAM role to Veeam Backup for AWS beforehand, you can do it without closing the **EFS Restore** wizard. To do that, click **Add** and complete the [Add IAM Role](#) wizard.

IMPORTANT

It is recommended that you check whether the selected IAM role has all the permissions required to perform the operation. If some permissions of the IAM role are missing, the restore operation will fail to complete successfully. To run the IAM role permission check, click **Check Permissions** and follow the instructions provided in section [Checking IAM Role Permissions](#).

The screenshot shows the 'EFS File-level Recovery' wizard in the 'Account' step. The 'Choose IAM role' section is active, with instructions: 'Specify an IAM role that will be used to access resources for the restore operation or provide temporary access keys. The selected IAM role must belong to an AWS account where the source file system resides.' The 'IAM role' radio button is selected. A dropdown menu is open, showing three options: 'Default Backup Restore (Default Backup Restore)', 'Default Backup Restore (Default Backup Restore)', and 'EFS restore role (Created by bd-regress-2 at 11/14/2023 6:45 PM)'. The 'EFS restore role' is highlighted. To the right of the dropdown are '+ Add' and 'Check Permissions' buttons. Below the dropdown is a 'Secret key:' field. An information icon and text state: 'The keys are used to perform this operation only. They are not saved or stored. To learn what permissions are required for performing the operation, see the User Guide.' At the bottom are 'Previous', 'Next', and 'Cancel' buttons.

Specifying One-Time Access Keys

To specify one-time access keys for restore, select the **Temporary access keys** option and use the **Access key** and **Secret key** fields to provide the access key ID and the secret access key.

NOTE

Veeam Backup for AWS does not store one-time access keys in the configuration database.

Veeam Backup for AWS Server time: Nov 16, 2023 2:27 PM administrator Portal Administrator Configuration

EFS File-level Recovery

Restore Type

Restore List

Account

Restore Mode

Reason

Summary

Choose IAM role

Specify an IAM role that will be used to access resources for the restore operation or provide temporary access keys. The selected IAM role must belong to an AWS account where the source file system resides.

IAM role

Default Backup Restore (Default Backup Restore) + Add Check Permissions

Temporary access keys

Access key: AKIAY4ZWOU4WMVRAGEVN

Secret key:

i The keys are used to perform this operation only. They are not saved or stored. To learn what permissions are required for performing the operation, see the [User Guide](#).

Previous Next Cancel

Step 5. Choose Restore Mode

At the **Restore Mode** step of the wizard, choose whether you want to restore files and folders to the original or to a custom location. If you select the **Restore to new location, or with different settings** option, specify the target AWS Region and the file system to which the files and folders will be restored.

The screenshot shows the 'EFS File-level Recovery' wizard in Veeam Backup for AWS. The interface includes a top navigation bar with the Veeam logo, server time (Jun 14, 2024 3:27 PM), and user information (administrator, Portal Administrator). A left sidebar contains navigation links: Restore Type, Restore List, Account, Restore Mode (highlighted), Reason, and Summary. The main content area is titled 'Choose restore mode' and includes the instruction: 'Specify whether you want to restore files to the original location or to a new one, or with different settings.' Two radio button options are present: 'Restore to original location' (unselected) and 'Restore to new location, or with different settings' (selected). Below the selected option, the 'Region' is set to 'Asia Pacific (Singapore)' and the 'File system' is 'bd-efs-singapore'. At the bottom, there are three buttons: 'Previous', 'Next', and 'Cancel'.

Step 6. Specify Restore Reason


At the **Reason** step of the wizard, you can specify a reason for restoring the files and folders. The information you provide will be saved in the session history and you can reference it later.

The screenshot shows the 'EFS File-level Recovery' wizard in the Veeam Backup for AWS interface. The top navigation bar includes the Veeam logo, 'Veeam Backup for AWS', server time 'Jun 14, 2024 3:28 PM', user 'administrator Portal Administrator', and a 'Configuration' link. The main content area is titled 'EFS File-level Recovery' and features a left-hand navigation pane with options: 'Restore Type', 'Restore List', 'Account', 'Restore Mode', 'Reason' (highlighted), and 'Summary'. The 'Reason' section is active, displaying the heading 'Restore reason' and the instruction 'Specify a reason for performing the restore operation.' Below this is a text input field containing the text 'Restoring corrupted files'. At the bottom of the wizard, there are three buttons: 'Previous', 'Next' (highlighted), and 'Cancel'.

Step 7. Finish Working with Wizard

At the **Summary** step of the wizard, review summary information and click **Finish**.

[Applies only if you have selected the **Browse files** option at the **Restore Type** step of the wizard] As soon as you click **Finish**, Veeam Backup for AWS will close the **EFS File-level Recovery** wizard, start a recovery session and display the **FLR Running Sessions** window. To select file and folders that you want to recover, follow the instructions provided in steps 8-10.



Step 8. Open FLR Browser

[This step applies only if you have selected the **Browse files** option at the **Restore Type** step of the wizard]

TIP

If you accidentally close the **FLR Running Sessions** window, navigate to **Protected Data > File Systems > EFS** and click the link in the **File-Level Recovery URL** column to open the window again.

In the **FLR Running Sessions** window you can track the progress of the recovery session. In the **URL** column of the window, Veeam Backup for AWS will display a link to the file-level recovery browser. You can use the link in either of the following ways:


- Click the link to open the file-level recovery browser on your local machine while the recovery session is running.
- Copy the link, close the **FLR Running Sessions** window and open the file-level recovery browser on another machine.

IMPORTANT

When you click **Copy URL**, Veeam Backup for AWS copies the following information to the clipboard:

- A link to the file-level recovery browser includes a public DNS name or an IP address of the backup appliance hosting the browser and authentication information used to access the browser.
- A thumbprint of a TLS certificate installed on the appliance hosting the file-level recovery browser.

To avoid a man-in-the-middle attack, before you start recovering files and folders, check that the certificate thumbprint displayed in the web browser from which you access the file-level recovery browser matches the provided certificate thumbprint.



Step 9. Select Restore Point

[This step applies only if you have selected the **Browse files** option at the **Restore Type** step of the wizard]

By default, Veeam Backup for AWS uses the most recent valid restore point. However, you can restore files and folders to an earlier state.

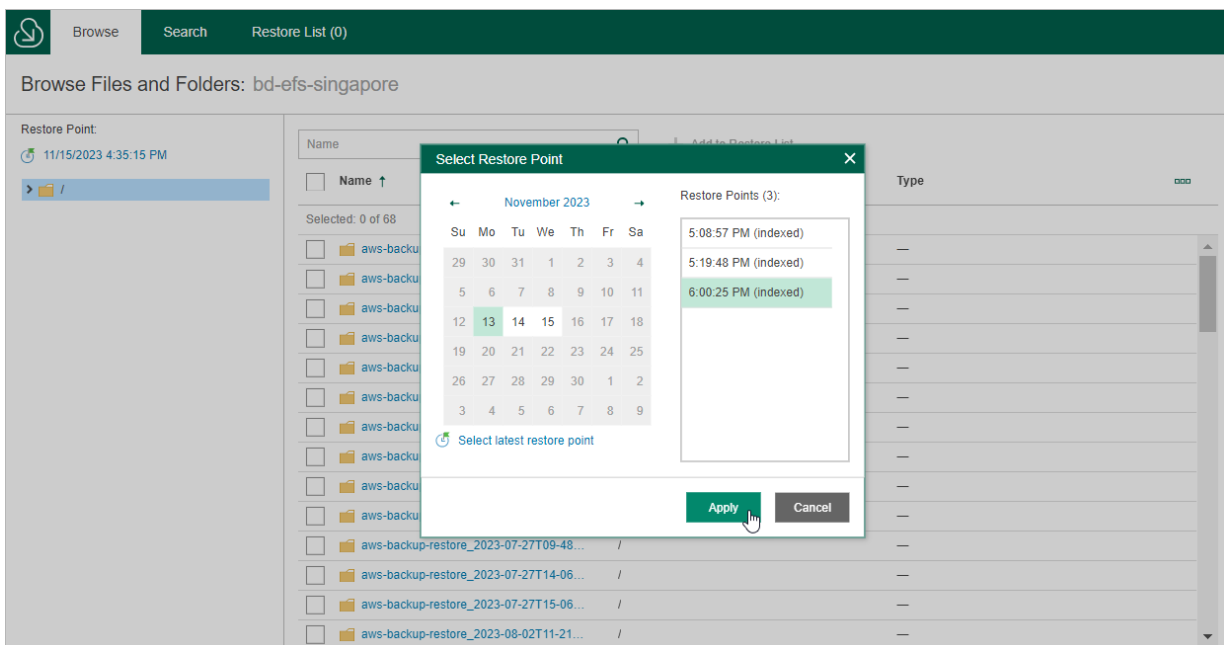
To select a restore point in the file-level recovery browser, do the following:

1. On the **Browse** tab, click the link in the **Restore Point** field.
2. In the **Select Restore Point** window, choose a date when the restore point was created, select the necessary restore point from the **Restore Points** list and click **Apply**.

The **Restore Points** list shows only restore points that are associated with created EFS indexes.

TIP

You can search for the necessary files in all indexed restore points simultaneously. To do that, switch to the **Search** tab, specify the file or folder name, its location and click **Search**.



Step 10. Choose Items to Recover

[This step applies only if you have selected the **Browse files** option at the **Restore Type** step of the wizard]

In the file-level recovery browser, you can find and recover items (files and folders) of the selected EFS file system. All recovered items are saved to the specified file system.

To select files and folders from the specific folder, do the following:

1. On the **Browse** tab, navigate to the folder that contains the necessary files.
2. In the working area, select check boxes next to the files and folders that you want to restore and click **Add to Restore List**.
3. Repeat steps 1-2 for all other files and folders that you want to restore.

If you want to restore different versions of a specific file or folder, select a new restore point as described in [Step 9. Select Restore Point](#), and then repeat steps 1-2.

TIP

You can search for the necessary files in all indexed restore points simultaneously. To do that, switch to the **Search** tab, specify the file or folder name, its location and click **Search**.

4. Switch to the **Restore List** tab.
5. On the **Restore List** tab, review the list files and folders, select check boxes next to the items that you want to recover and click **Restore**.

As soon as you click **Restore**, Veeam Backup for AWS will restore the selected files to the file system that you have specified at [step 4](#) of the **EFS File-level Recovery** wizard. You can track the progress and view the results of the restore operation in the **Session Log** section of the **Restore List** tab.

Restore List: bd-efs-singapore

Restore Status: All [Green] [Yellow] [Red]

Restore Remove

<input checked="" type="checkbox"/> Name ↑	Location	Type	Restore Point	Restore Date	Restore Status	...
<input checked="" type="checkbox"/> aws-backup-restore_2...	/	—	11/13/2023 6:00:25 PM	—	—	
<input checked="" type="checkbox"/> aws-backup-restore_2...	/	—	11/15/2023 4:35:15 PM	—	—	

Selected: All 2 items

Session Log

Status: All [Green] [Yellow] [Red]

Action	Status	Start Time	End Time	Duration	...
Select a single item to view sessions details					

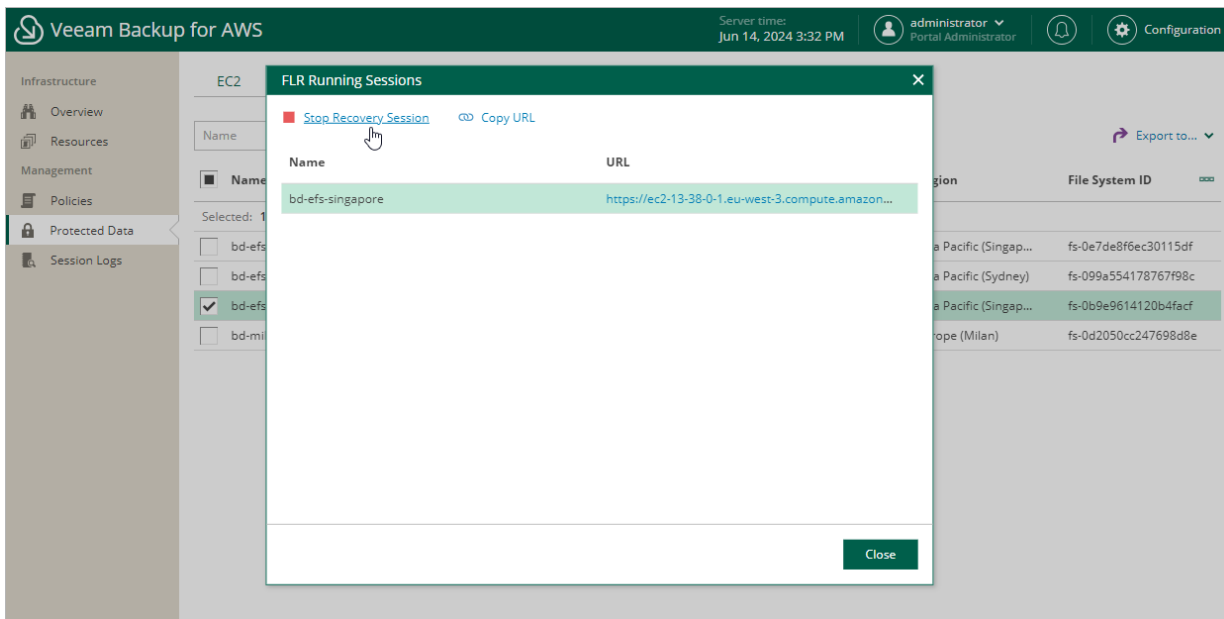
Step 11. Stop Recovery Session

[This step applies only if you have selected the **Browse files** option at the **Restore Type** step of the wizard]

After you finish working with the file-level recovery browser, it is recommended that you stop the recovery session. To stop the recovery session, click **Stop Recovery Session** in the **FLR Running Sessions** window. If you do not perform any actions in the file-level recovery browser for 30 minutes, Veeam Backup for AWS will stop the recovery session automatically.

TIP

If you accidentally close the **FLR Running Sessions** window, navigate to **Protected Data > File Systems > EFS** and click the link in the **File-Level Recovery URL** column to open the window again.



FSx Restore

The actions that you can perform with restore points of FSx file systems depend on whether you access the restore points using the Veeam Backup & Replication console or the Veeam Backup for AWS Web UI.

FSx Restore Using Console

You can recover corrupted FSx file systems in the Veeam Backup for AWS Web UI only. However, you can launch the **FSx Restore** wizard directly from the Veeam Backup & Replication console to start the restore operation:

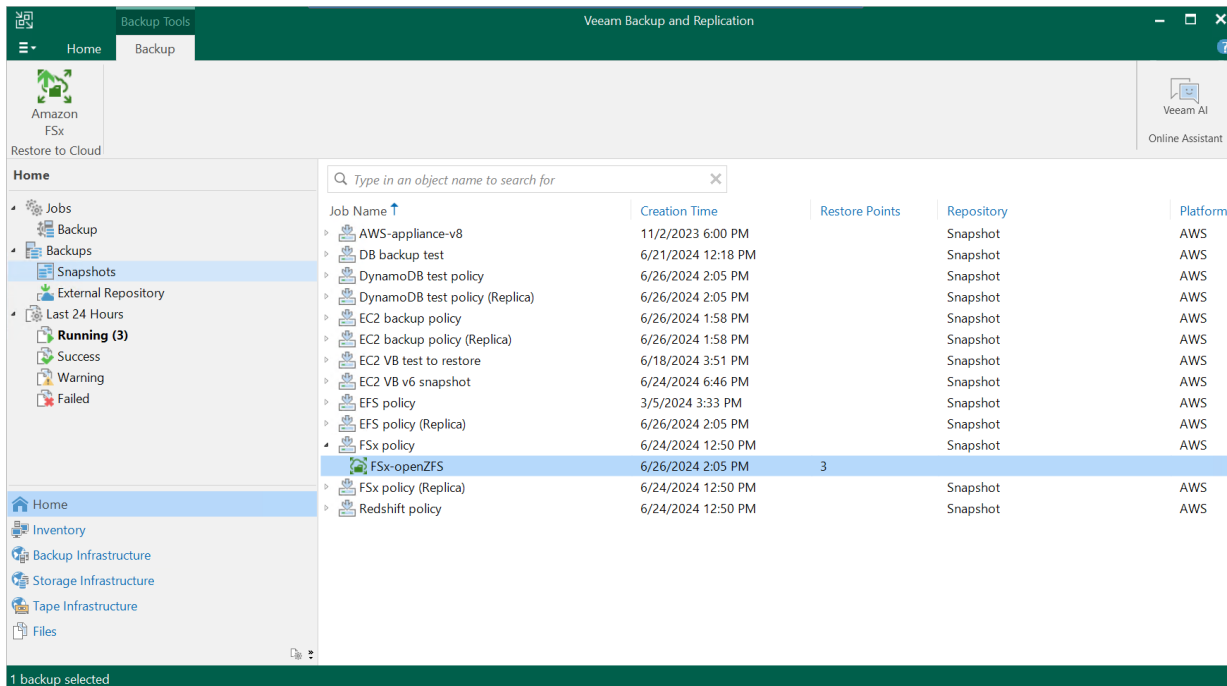
1. In the Veeam Backup & Replication console, open the **Home** view.
2. Navigate to **Backups > Snapshots**.
3. Expand the backup policy that protects the FSx file systems that you want to recover, select the necessary file system and click **Amazon FSx** on the ribbon.

Alternatively, you can right-click the selected file system and click **Restore to Amazon FSx**.

IMPORTANT

You cannot restore multiple FSx file systems from the Veeam Backup & Replication console.

Veeam Backup & Replication will open the **FSx Restore** wizard in a web browser. Complete the wizard as described in section [FSx Restore Using Web UI](#).



FSx Restore Using Web UI

In case of a disaster, you can restore a FSx file system from a FSx backup or backup copy. Veeam Backup for AWS allows you to restore one or more file systems at a time, to the original location or to a new location. To learn how FSx restore works, see [FSx Restore](#).

IMPORTANT

Consider the following:

- Veeam Backup for AWS supports restore of FSx file systems only to the same AWS accounts to which the source file systems belong.
- Veeam Backup for AWS supports restore of Amazon FSx for Windows File Server file systems. However, before you start a restore operation, it is recommended that you use the Amazon FSx Active Directory Validation tool to check the connection between the file systems that you plan to restore and the Microsoft Active Directories to which these file systems will be joined. To learn how to use the validation tool, see [AWS Documentation](#).

How to Perform File System Restore

To restore a protected FSx file system, do the following:

1. [Launch the FSx Restore wizard](#).
2. [Select a restore point](#).
3. [Specify an IAM identity for restore](#).
4. [Choose a restore mode](#).
5. [Enable encryption for the restored table](#).
6. [Configure General Settings](#).
7. [Configure network settings](#).
8. [Specify a restore reason](#).
9. [Finish working with the wizard](#).

Step 1. Launch FSx Restore Wizard

To launch the **FSx Restore** wizard, do the following:

1. Navigate to **Protected Data > File Systems > FSx**.
2. Select the FSx file system that you want to restore.
3. Click **Restore**.

Alternatively, click the link in the **Restore Points** column. Then, in the **Available Restore Points** window, select the necessary restore point and click **Restore**.

NOTE

You can restore multiple FSx file systems if they belong to same AWS account only.

The screenshot shows the Veeam Backup for AWS interface. The top navigation bar includes the Veeam logo, the text 'Veeam Backup for AWS', the server time 'Jun 17, 2024 2:24 PM', and the user 'administrator Portal Administrator'. The left sidebar shows the navigation menu with 'Protected Data' selected. The main content area is titled 'File Systems' and shows a list of file systems. The table has the following columns: Name, Policy, Restore Points, Latest Restore Point, AWS Account, and File System ID. Two rows are selected, highlighted in green. The 'Restore' button is visible above the table.

Name	Policy	Restore Points	Latest Restore Point	AWS Account	File System ID
Selected: 2 of 88					
<input type="checkbox"/> bd-virginia-openzfs-...	—	2	06/03/2024 12:10:15 PM	487569979969	fs-03615bbfbffb9e29f
<input type="checkbox"/> bd-virginia-openzfs-...	—	9	06/07/2024 4:00:19 PM	487569979969	fs-0ac252290813f4d...
<input type="checkbox"/> bd-virginia-openzfs-...	—	2	06/03/2024 12:10:09 PM	487569979969	fs-02061c3c2e9d6115f
<input type="checkbox"/> bd-virginia-openzfs-...	—	2	06/03/2024 12:10:10 PM	487569979969	fs-0aa20e90d2cd02a...
<input checked="" type="checkbox"/> bd-virginia-openzfs-...	—	2	06/03/2024 12:10:09 PM	487569979969	fs-06c5c2167dfb604ab
<input type="checkbox"/> bd-virginia-openzfs-...	—	2	06/03/2024 12:10:13 PM	487569979969	fs-0d4d81a6256ec66...
<input type="checkbox"/> bd-virginia-openzfs-...	—	9	06/07/2024 4:00:18 PM	487569979969	fs-0841492217ab2ccf7
<input type="checkbox"/> bd-virginia-openzfs-...	—	2	06/03/2024 12:10:14 PM	487569979969	fs-09b5cb137917618...
<input type="checkbox"/> bd-virginia-openzfs-...	—	2	06/03/2024 12:10:14 PM	487569979969	fs-029994710f99485...
<input type="checkbox"/> bd-virginia-openzfs-...	—	2	06/03/2024 12:10:13 PM	487569979969	fs-0b0b1aa7bf7f699f3
<input type="checkbox"/> bd-windows-aws-ma...	—	1	03/01/2024 5:33:11 PM	487569979969	fs-065229328e1980b...
<input type="checkbox"/> bd-windows-aws-ma...	—	2	04/04/2024 12:00:18 PM	487569979969	fs-06f547be35007e1...
<input type="checkbox"/> bd-windows-london-...	—	1	03/01/2024 5:33:11 PM	487569979969	fs-0e42c14d64ced94...
<input type="checkbox"/> bd-windows-restore...	—	1	04/03/2024 5:00:14 PM	487569979969	fs-0900673213caf579a
<input checked="" type="checkbox"/> bd-windows-self-ma...	—	2	04/04/2024 12:00:17 PM	487569979969	fs-0f85347810f94232c

Step 2. Select Restore Point

At the **File System** step of the wizard, you can add file systems to the restore session and select restore points to be used to perform the restore operation for each added file system. By default, Veeam Backup for AWS uses the most recent valid restore point. However, you can restore a file system to an earlier state.

To select a restore point, do the following:

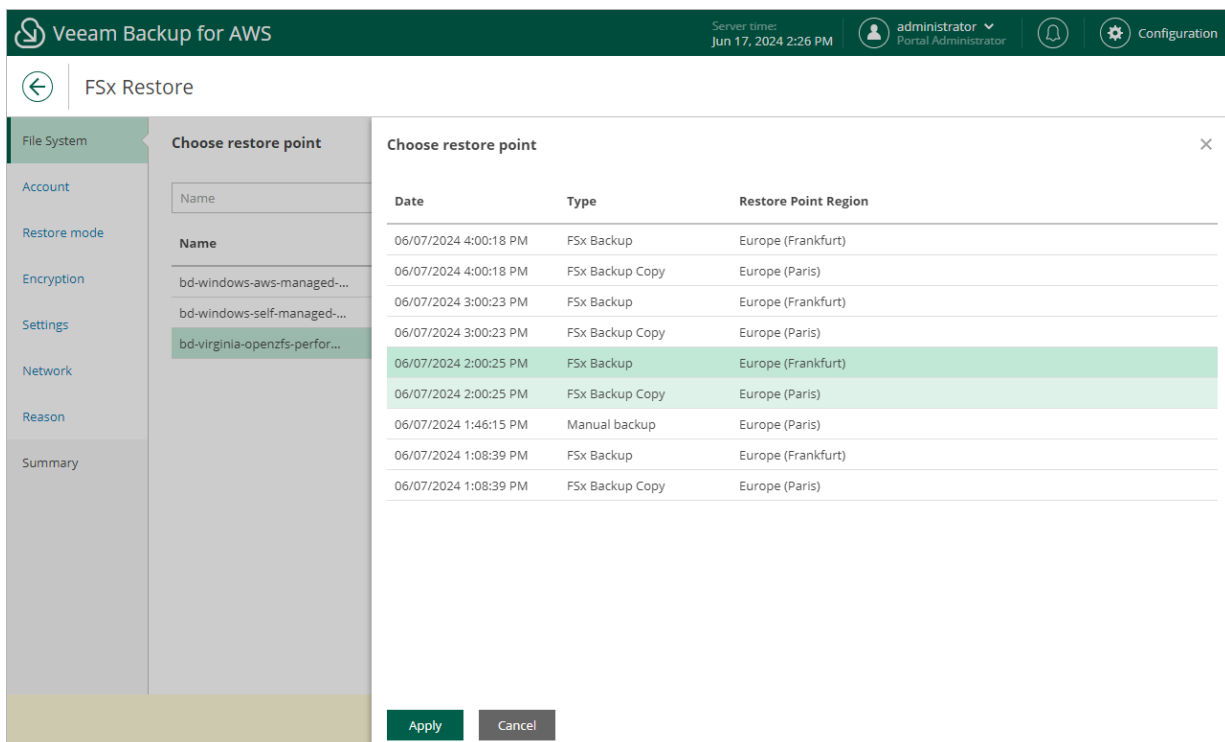
1. Select the file system and click **Restore Point**.
2. In the **Choose restore point** window, select the necessary restore point and click **Apply**.

To help you choose a restore point, Veeam Backup for AWS provides the following information on each available restore point:

- **Date** – the date when the restore point was created.
- **Type** – the type of the restore point:
 - *FSx backup* – a FSx backup created by a backup policy.
 - *FSx backup copy* – a FSx backup copy created by a backup policy.
 - *Manual backup* – a FSx backup created manually.
- **Restore Point Region** – the AWS Region where the restore point is stored.

IMPORTANT

Keep in mind that since Veeam Backup for AWS does not support cross-region copying of FSx backups for [opt-in Regions](#), some of the [restore options](#) may not be available. To work around the issue, it is recommended that you select restore points stored in the same opt-in Region or the same default AWS Region (that is, one of the AWS Regions activated for your AWS account by default) if you plan to perform restore either to a new location or to the original location but with different settings.



The screenshot shows the Veeam Backup for AWS interface. At the top, there is a navigation bar with the Veeam logo, the text 'Veeam Backup for AWS', and user information: 'administrator Portal Administrator'. Below the navigation bar, the main window is titled 'FSx Restore'. On the left, there is a sidebar with various settings categories: File System, Account, Restore mode, Encryption, Settings, Network, Reason, and Summary. The 'File System' category is selected, and the 'Choose restore point' dialog box is open. The dialog box has a title bar 'Choose restore point' and a close button 'X'. Inside the dialog, there is a table with the following columns: Date, Type, and Restore Point Region. The table contains the following data:

Date	Type	Restore Point Region
06/07/2024 4:00:18 PM	FSx Backup	Europe (Frankfurt)
06/07/2024 4:00:18 PM	FSx Backup Copy	Europe (Paris)
06/07/2024 3:00:23 PM	FSx Backup	Europe (Frankfurt)
06/07/2024 3:00:23 PM	FSx Backup Copy	Europe (Paris)
06/07/2024 2:00:25 PM	FSx Backup	Europe (Frankfurt)
06/07/2024 2:00:25 PM	FSx Backup Copy	Europe (Paris)
06/07/2024 1:46:15 PM	Manual backup	Europe (Paris)
06/07/2024 1:08:39 PM	FSx Backup	Europe (Frankfurt)
06/07/2024 1:08:39 PM	FSx Backup Copy	Europe (Paris)

At the bottom of the dialog, there are two buttons: 'Apply' and 'Cancel'.

Step 3. Specify IAM Identity

At the **Account** step of the wizard, choose whether you want to [use an IAM role](#) or [one-time access keys of an IAM user](#) to allow Veeam Backup for AWS to perform the restore operation. For information on the permissions that the IAM role or IAM user must have to perform the restore operation, see [FSx Restore IAM Permissions](#).

IMPORTANT

Make sure that the specified IAM role or one-time access keys belong to an AWS account where the source file system resides.

Specifying IAM Role

To specify an IAM role for restore, select the **IAM role** option and choose the necessary IAM role from the list.

For an IAM role to be displayed in the **IAM role** list, it must be added to Veeam Backup for AWS with the *Amazon FSx Restore* operation selected as described in section [Adding IAM Roles](#). If you have not added the necessary IAM role to Veeam Backup for AWS beforehand, you can do it without closing the **FSx Restore** wizard. To do that, click **Add** and complete the [Add IAM Role](#) wizard.

IMPORTANT

It is recommended that you check whether the selected IAM role has all the permissions required to perform the operation. If some permissions of the IAM role are missing, the restore operation will fail to complete successfully. To run the IAM role permission check, click **Check Permissions** and follow the instructions provided in section [Checking IAM Role Permissions](#).

The screenshot shows the Veeam Backup for AWS interface during the FSx Restore wizard. The top navigation bar includes the Veeam logo, the text "Veeam Backup for AWS", the server time "Jun 17, 2024 2:26 PM", the user "administrator Portal Administrator", a notification bell, and a "Configuration" gear icon. The main content area is titled "FSx Restore" and has a left-hand navigation menu with options: "File System", "Account" (selected), "Restore mode", "Settings", "Reason", and "Summary". The "Account" section is active, displaying the heading "Choose IAM role" and a descriptive text: "Specify an IAM role that will be used to access resources for the restore operation or provide temporary access keys. The selected IAM role must belong to an AWS account where the source file system resides." Below this text, the "IAM role" radio button is selected. A list of IAM roles is shown, with "Default Backup Restore (Default Backup Restore)" selected and highlighted in green. Other roles include "bd-fsx-autogen-new-permission-list", "bd-fsx-restore-role-regress", "FSx Backup Role final list", "FSx Restore role", "FSx Restore Role final list", and "FSx Restore role from wizard". To the right of the list are "+ Add" and "Check Permissions" buttons. At the bottom of the wizard, there are "Previous", "Next", and "Cancel" buttons.

Specifying One-Time Access Keys

To specify one-time access keys for restore, select the **Temporary access keys** option and use the **Access key** and **Secret key** fields to provide the access key ID and the secret access key.

NOTE

Veeam Backup for AWS does not store one-time access keys in the configuration database.

The screenshot shows the Veeam Backup for AWS interface for an FSx Restore operation. The top navigation bar includes the Veeam logo, the text 'Veeam Backup for AWS', the server time 'Jun 17, 2024 2:26 PM', and the user 'administrator Portal Administrator'. The main content area is titled 'FSx Restore' and has a sidebar with options: File System, Account, Restore mode, Reason, and Summary. The 'Account' section is active, showing the 'Choose IAM role' step. It instructs the user to specify an IAM role or provide temporary access keys. The 'Temporary access keys' option is selected. The 'Access key' field contains 'ALDOHVBGJK68KJSHGGF' and the 'Secret key' field is masked with dots. A warning icon and text state: 'The keys are used to perform this operation only. They are not saved or stored. To learn what permissions are required for performing the operation, see the User Guide.' At the bottom, there are 'Previous', 'Next', and 'Cancel' buttons.

Step 4. Choose Restore Mode

At the **Restore Mode** step of the wizard, choose whether you want to restore the selected FSx file system to the original or to a custom location.

IMPORTANT

Consider the following:

- If any of the restore options are not available, make sure that the selected restore points meet the requirements listed at [step 2](#).
- If the file system that you want to restore is an Amazon FSx for OpenZFS system with the Single-AZ 1 (HA) or Single-AZ 2 (HA) deployment type, Veeam Backup for AWS will not be able to restore these file systems to a new location or to the original location but with different settings.
- If the file system that you want to restore is an Amazon FSx for Lustre system with the Persistent 2 deployment type or an Amazon for OpenZFS system with the Single-AZ 2 deployment type, make sure the target AWS Region is supported by this deployment type. Otherwise, the restore operation will fail to complete successfully. For the list of supported AWS Regions for Amazon FSx for Lustre file systems, see [AWS Documentation](#). For the list of supported AWS Regions for Amazon OpenZFS file systems, see [AWS Documentation](#).
- If the file system that you want to restore is an FSx for Windows File Server file system associated with DNS aliases, keep in mind that Veeam Backup for AWS will not be able to restore these aliases to a new location or to the original location but with different settings. To learn how to manage DNS aliases, see [AWS Documentation](#).

If you select the **Restore to a new location, or with different settings** option, specify the target AWS Region where the restored file system will reside. Keep in mind that the list of available AWS Regions depends on the location of the restore point that you select at [step 2](#) of the wizard due to AWS limitations for opt-in regions. That is, if the restore point is stored in an AWS Region activated for your AWS account by default, you will be able to select any of the default AWS Regions; if the restore point is stored in an opt-in Region, you will be able to select the source opt-in Region only.

The screenshot shows the 'FSx Restore' wizard in the Veeam Backup for AWS console. The 'Restore mode' step is active, showing two options: 'Restore to original location' and 'Restore to new location, or with different settings'. The second option is selected, and a dropdown menu shows 'Europe (Stockholm)' as the target region. The interface includes a navigation sidebar on the left, a top status bar with server time and user information, and navigation buttons at the bottom.

Choose restore mode
Specify whether you want to restore the file system to the original location or to a new one, or with different settings.

Restore to original location
Quickly restore the selected file system to the original location, with the same name and settings as the source file system.

Restore to new location, or with different settings
Perform additional configuration steps to restore the selected file system to a new location or to use settings that differ from the source settings.

Europe (Stockholm)

Previous Next Cancel

Step 5. Enable Encryption

[This step applies only if you have selected the **Restore to new location, or with different settings** option at the **Restore Mode** step of the wizard]

At the **Encryption** step of the wizard, choose whether the restored file system will be encrypted with AWS KMS keys:

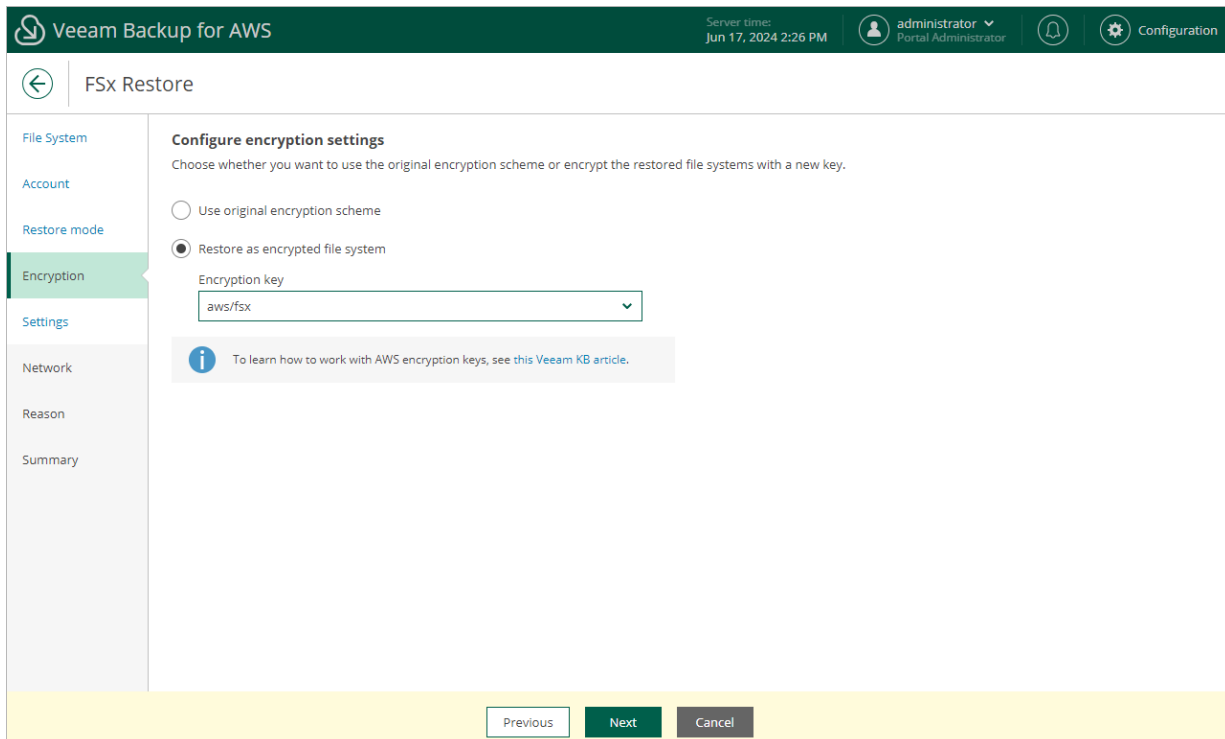
- If you want to apply the existing encryption scheme, select the **Use original encryption scheme** option.
- If you want to change the key that is used for file system encryption, select the **Restore as encrypted file system** option and choose the necessary KMS key from the **Encryption key** drop-down list.

For a KMS key to be displayed in the list of available encryption keys, it must be stored in the AWS Region selected at [step 4](#) of the wizard, and the IAM role or user specified for the restore operation at [step 3](#) of the wizard must have permissions to access the key. For more information on KMS keys, see [AWS Documentation](#).

TIP

If the necessary KMS key is not displayed in the list, or if you want to use a KMS key from an AWS account other than the AWS account to which the specified IAM role belongs, you can select *Add custom key ARN* from the **Encryption key** drop-down list, and specify the Amazon resource name (ARN) of the key in the **Add Custom Key ARN** window.

For Veeam Backup for AWS to be able to encrypt the restored file system using the provided KMS key, either the IAM role or user specified for the restore operation, or the IAM role used to create the restore point selected at [step 2](#) of the wizard must have permissions to access the key.



Step 6. Configure General Settings

The list of settings that you can configure for a restored file system depend both on the file system type and the option you choose at the **Choose Restore Mode** step of the wizard.

IMPORTANT

Veeam Backup for AWS does not allow you to change the type of Microsoft Active Directory to which the restored file system will be joined.

In This Section

- [Restoring to Original Location](#)
- [Restoring to New Location](#)

Restoring to Original Location


[This step applies only if you have selected the **Restore to original location** option at the **Restore Mode** step of the wizard, and if the selected file system is an Amazon FSx for Windows File Server file system joined to self-managed Microsoft Active Directories (AD)]

At the **Settings** step of the wizard, check Active Directory settings that will be used to authenticate against the Microsoft AD to which the restored file system will be connected, and provide the password of the service account. To do that, select the file system and click **Edit**.

NOTE

Veeam Backup for AWS does not store client secrets in the configuration database.

For more information on Microsoft Active Directory in FSx for Windows File Server, see [AWS Documentation](#).



Restoring to New Location

[This step applies only if you have selected the **Restore to new location, or with different settings** option at the **Restore Mode** step of the wizard]

At the **Settings** step of the wizard, you can provide a new name and specify Active Directory details for the restored file system, depending on the file system type:

- To specify a new name, select the file system and click **Rename**. In the **File system name** window, specify the name and click **Apply**.
- To specify Active Directory settings for an Amazon FSx for Windows File Server file system with AWS Managed Microsoft AD, select the file system and click **Edit**. Then, in the **General Settings** window, do the following:
 - a. From the **Active Directory** drop-down list, select the AWS Managed Microsoft AD to which the restored file system will be joined.

For an AWS Managed Microsoft AD to be displayed in the list of available directories, it must be created in the AWS Region specified at [step 4](#) of the wizard as described in [AWS Documentation](#).
 - b. To save changes made to the file system settings, click **Apply**.
- To specify Active Directory settings for an Amazon FSx for Windows File Server file system with self-managed Microsoft AD, select the file system and click **Edit**. Then, in the **General Settings** window, do the following:
 - a. In the **Active Directory domain name** field, enter the domain name of an AD to which file system will be joined.
 - b. In the **DNS server IP addresses** field, enter the IPv4 address of DNS servers configured for the domain.
 - c. In the **Service account username** field, enter the name for a service account (without a domain prefix or suffix) that has access to the restored file system.
 - d. In the **Service account password** field, enter the password of the service account.
 - e. [Optional] In the **Organizational unit** field, enter the path name of an organizational unit in which you want to connect your file system.
 - f. [Optional] In the **System administrators group** field, enter the name of an AD group that has privileges to manage the restored file system.
 - g. Click **Apply**.

NOTE

Veeam Backup for AWS does not store client secrets in the configuration database.

For more information on Microsoft Active Directory in FSx for Windows File Server, see [AWS Documentation](#).

The screenshot shows the Veeam Backup for AWS interface. At the top, the header includes the Veeam logo, the text "Veeam Backup for AWS", the server time "Jun 17, 2024 2:26 PM", and the user "administrator Portal Administrator". A navigation bar on the left contains "FSx Restore" and a list of menu items: "File System", "Account", "Restore mode", "Encryption", "Settings" (highlighted), "Network", "Reason", and "Summary".

The main content area is titled "Configure restore settings" and includes a sub-section "General settings". The "General settings" section is titled "Specify the Active Directory settings and provide additional details if required." and contains the following fields:

- Active Directory type: Self-managed Microsoft Active Directory
- Active Directory domain name: ⓘ
- DNS server IP addresses: ⓘ
- Service account username: ⓘ
- Service account password: ⓘ
- Organizational unit (optional): ⓘ
- System administrators group (optional): ⓘ

At the bottom of the "General settings" section are two buttons: "Apply" and "Cancel".

Step 7. Configure Network Settings

[This step applies only if you have selected the **Restore to new location, or with different settings** option at the **Restore Mode** step of the wizard]

At the **Network** step of the wizard, configure network settings for the restored file system. To do that:

1. Select the file system and click **Edit**.
2. In the **Network settings** window, do the following:
 - a. From the **VPC** drop-down list, select an Amazon VPC to which the restored FSx file system will be connected.

For a VPC to be displayed in the list of available VPCs, it must be created in the AWS Region specified at [step 4](#) of the wizard as described in [AWS Documentation](#).
 - b. [Applies only to the file systems with the Single-AZ deployment type] From the **Subnet** drop-down list, select a subnet in which the elastic network interface of the file system will reside.

For a subnet to be displayed in the list of available subnets, it must be created in the AWS Region specified at [step 4](#) of the wizard as described in [AWS Documentation](#).
 - c. [Applies only to file systems with the Multi-AZ deployment type] From the **Preferred subnet** and **Standby subnet** drop-down lists, select subnets in which the network interfaces of the primary and standby file servers will reside.

For a subnet to be displayed in the list of available subnets, it must be created for the preferred Availability Zone in the specified VPC as described in [AWS Documentation](#).
 - d. Click the link in the **Security group** filed. In the opened window, select security groups that will be associated with the restored file system. Note that you cannot associate more than 5 security groups.

For a security group to be displayed in the list of available groups, it must be created in the AWS Region specified at [step 4](#) of the wizard as described in [AWS Documentation](#).
 - e. [Applies only to file systems of the Amazon FSx for OpenZFS file system type] Click the link in the **Route tables** filed. In the opened window, select a route table that will be associated with the subnet of the specified VPC.

For a route table to be displayed in the list of available tables, it must be created in the AWS Region specified at [step 4](#) of the wizard as described in [AWS Documentation](#).
 - f. Click **Apply**.

IMPORTANT

Consider the following:

- Preferred and standby subnets must reside in different Availability Zones within the same AWS Region selected at [step 4](#) of the wizard.
- The selected security group must meet requirements described in section [Protecting FSx File Systems](#).

The screenshot shows the Veeam Backup for AWS interface during the FSx Restore process. The top navigation bar includes the Veeam logo, the text 'Veeam Backup for AWS', the server time 'Jun 17, 2024 2:27 PM', and the user 'administrator Portal Administrator'. The main content area is titled 'FSx Restore' and features a left-hand navigation menu with options: File System, Account, Restore mode, Encryption, Settings, Network (highlighted), Reason, and Summary. The 'Configure network settings' section is active, with a sub-section 'Network settings' that prompts the user to 'Choose a VPC, subnet and security group.' The configuration fields are as follows: VPC is set to 'vpc-1a555373'; Preferred subnet is 'subnet-0d1cb2d2dfc7253c5 172.31.1'; Standby subnet is 'subnet-634a0e18 172.31.16.0/20 (eu-'; and Security group is '1 security group selected'. A 'Browse...' button is visible next to each field. An information icon and a message box state: 'The selected security groups must allow inbound and/or outbound traffic, depending on the FSx file system type. For more information, see [User Guide](#).' At the bottom of the configuration area are 'Apply' and 'Cancel' buttons.

Step 8. Specify Restore Reason

At the **Reason** step of the wizard, you can specify a reason for restoring the FSx file system. The information you provide will be saved in the session history, and you can reference it later.

The screenshot shows the Veeam Backup for AWS interface during the FSx Restore process. The top navigation bar includes the Veeam logo, the text "Veeam Backup for AWS", the server time "Jun 17, 2024 2:28 PM", the user "administrator Portal Administrator", a notification bell, and a "Configuration" gear icon. Below the navigation bar, the page title is "FSx Restore" with a back arrow. A left-hand sidebar lists the wizard steps: "File System", "Account", "Restore mode", "Encryption", "Settings", "Network", "Reason" (which is highlighted in green), and "Summary". The main content area is titled "Reason" and contains the instruction "Specify a reason for performing the restore operation." Below this is a text input field labeled "Restore reason:" containing the text "corrupted fsx file system". At the bottom of the wizard, there are three buttons: "Previous", "Next" (highlighted in green), and "Cancel".

Step 9. Finish Working with Wizard

At the **Summary** step of the wizard, review summary information and click **Finish**.

The screenshot shows the 'FSx Restore' wizard in the 'Summary' step. The interface includes a top navigation bar with the Veeam logo, server time (Jun 17, 2024 2:28 PM), user information (administrator, Portal Administrator), and a Configuration icon. A left sidebar lists steps: File System, Account, Restore mode, Encryption, Settings, Network, Reason, and Summary (highlighted). The main content area is titled 'Review configured settings' and contains the following information:

- Reason:** corrupted fsx file system
- General settings:**
 - Restore mode: New location
 - Location name: Europe (Stockholm)
- IAM role:**
 - IAM role name: bd-fsx-autogen-new-permission-list
- Encryption:**
 - Encryption: Encrypted file system
 - Encryption key: aws/fsx

At the bottom of the wizard, there are three buttons: 'Previous', 'Finish' (highlighted in green), and 'Cancel'.

VPC Configuration Restore

The actions that you can perform with restore points of VPC configurations depend on whether you access the restore points using the Veeam Backup & Replication console or the Veeam Backup for AWS Web UI.

Performing VPC Configuration Restore Using Console

You can recover corrupted Amazon VPC configurations in the Veeam Backup for AWS UI only. However, you can launch the **VPC Restore** wizard directly from the Veeam Backup & Replication console to start the restore operation:

1. In the Veeam Backup & Replication console, open the **Home** view.
2. Navigate to **Backups > External Repository**.
3. Expand the AWS account in which VPC configuration has been backed up, select the AWS Region whose VPC configuration you want to restore and click **Amazon VPC** on the ribbon.

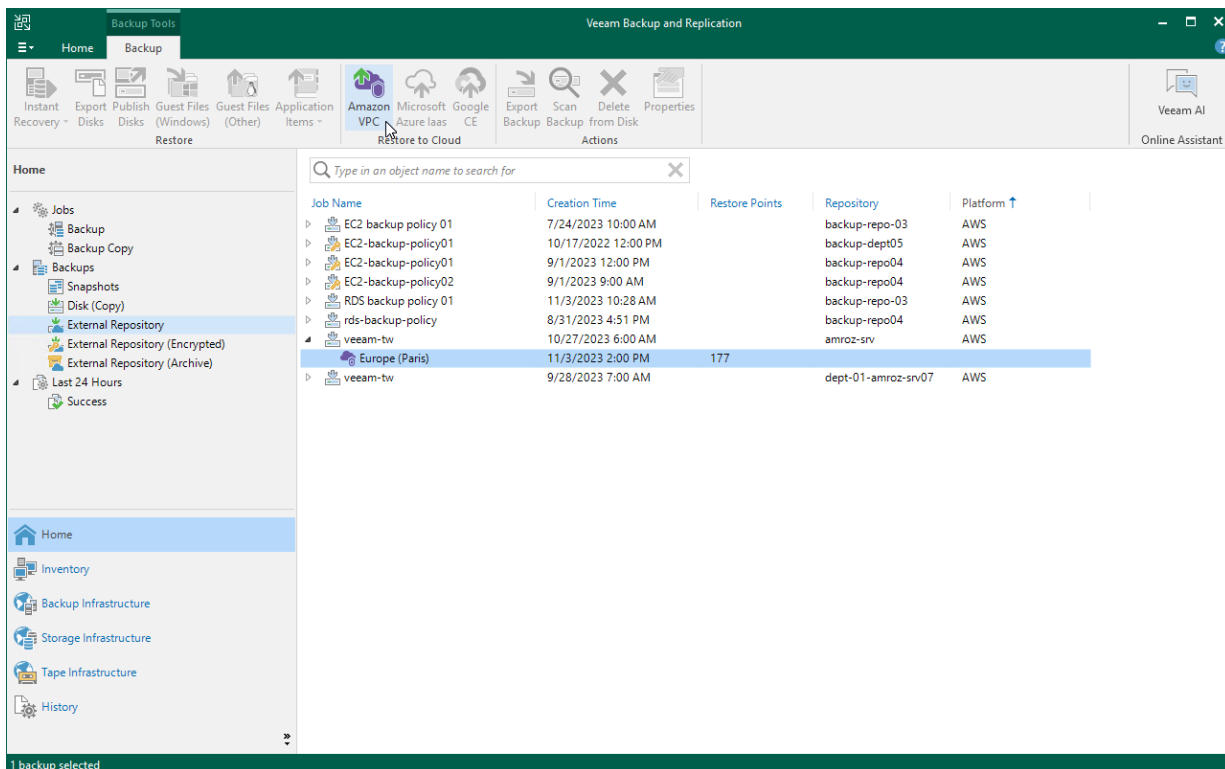
Alternatively, you can right-click the selected region and click **Restore to Amazon VPC**.

Veeam Backup & Replication will open the **VPC Restore** wizard in a web browser. Complete the wizard as described in section [VPC Configuration Restore](#).

IMPORTANT

Consider the following:

- VPC configuration restore is available only if you have logged in to the Veeam Backup & Replication console under a user account with the Veeam Backup Administrator role. For more information on user roles, see the Veeam Backup & Replication User Guide, section [Roles and Users](#).
- Selected items restore of the virtual network configuration is not available from the Veeam Backup & Replication console – you can perform it using the Veeam Backup for AWS Web UI only.



VPC Configuration Restore Using Web UI

Veeam Backup for AWS offers the following disaster recovery operations:

- [VPC configuration restore](#) – restores an entire VPC configuration.
- [Selected items restore](#) – restores the selected VPC configuration items.

You can restore the VPC configuration data to the most recent state or to any available restore point.

IMPORTANT

When restoring VPC route tables, consider that routes that had the `blackhole` state when a restore point was created will not be restored and a restore session will complete with warning. In this case, it is recommended to check the restored target route table configurations in the AWS Management Console to ensure that all traffic flows correctly. To learn how to configure routes in route tables, see [AWS Documentation](#).

Performing Entire Configuration Restore

In case of unexpected configuration changes, you can restore entire Amazon VPC configuration from a VPC configuration backup. Veeam Backup for AWS allows you to restore the VPC configuration to the original location or to a new location.

IMPORTANT

Restore to a new location is not supported for the following VPC configuration items:

- Client VPN endpoints.
- Customer gateways and load balancer listeners that use authentication certificates.
- In route tables, for core networks and routes to AWS Outpost local gateways, network interfaces, instances and carrier gateways.

How to Perform Entire VPC Configuration Restore

To restore the entire VPC configuration, do the following:

1. [Launch the VPC Restore wizard](#).
2. [Select a restore point and VPCs to restore](#).
3. [Specify an IAM identity for restore](#).
4. [Choose a restore mode](#).
5. [Configure mapping for Availability Zones](#).
6. [Review settings of VPC peering connections](#).
7. [Specify a restore reason](#).
8. [Finish working with the wizard](#).

Step 1. Launch VPC Restore Wizard

To launch the **VPC Restore** wizard, do the following:

1. Navigate to **Protected Data > VPC**.
2. Select the configuration record for an AWS Region whose VPC configuration you want to restore.
3. Click **Restore > Entire VPC**.

The screenshot shows the Veeam Backup for AWS interface. The top navigation bar includes the Veeam logo, server time (Jun 14, 2024 2:49 PM), user (administrator), and Configuration settings. The left sidebar shows navigation options: Infrastructure, Overview, Resources, Management, Policies, Protected Data, and Session Logs. The main content area is divided into tabs: EC2, Databases, File Systems, and VPC. The VPC tab is active, displaying a table of VPC configurations. A search bar and action buttons (Restore, Export, Compare, Remove) are visible. The 'Restore' dropdown menu is open, showing 'Entire VPC' and 'Selected Items' options. Below the table, the 'Configuration Details' section shows a list of VPC components with columns for Name, ID, Type, Modification Date, and State.

Account	Region	Latest Backup	Latest Changes	Restore Points
487569979969 (veeam-qa-b...)	Europe (London)	06/05/2024 1:08:47 PM	1 Vpc, 5 Subnet, 1 Endpoint, ...	1
407355668422 (veeam-qa-d...)	Europe (London)	06/05/2024 1:08:47 PM	3 Vpc, 5 Subnet, 1 Endpoint, ...	1

Name	ID ↑	Type	Modification Date	State
—	acl-086f91540d0ab2e21	Acl	06/05/2024 1:08:47 PM	Created
—	acl-0cf462f3a18242324	Acl	06/05/2024 1:08:47 PM	Created
—	acl-e315308bb	Acl	06/05/2024 1:08:47 PM	Created
—	dopt-5ada8832	DHCPOptionsSet	06/05/2024 1:08:47 PM	Created
nm-vb-v3 InternetGateway	igw-01345a5cbcac3a109	InternetGateway	06/05/2024 1:08:47 PM	Created
—	igw-a3b25fcb	InternetGateway	06/05/2024 1:08:47 PM	Created
com.amazonaws.eu-west-2...	pl-05cadb937455b8b91	ManagedPrefixList	06/05/2024 1:08:47 PM	Created
com.amazonaws.global.gr...	pl-082b2916fa6d7b4b7	ManagedPrefixList	06/05/2024 1:08:47 PM	Created

Step 2. Select Restore Point

At the **Restore List** step of the wizard, select a restore point that will be used to restore the selected VPC configuration. By default, Veeam Backup for AWS uses the most recent valid restore point. However, you can restore the VPC configuration data to an earlier state.

To select a restore point, do the following:

1. In the **Choose restore point** section, click the link to the right of **Restore point**.
2. In the **Available restore points** window, select the necessary restore point and click **Apply**.
3. In the **Choose VPCs to restore** section, select VPCs whose configuration you want to restore.

The screenshot shows the Veeam Backup for AWS interface during the VPC Restore process. The top navigation bar includes the Veeam logo, 'Veeam Backup for AWS', the server time 'Jun 14, 2024 2:50 PM', and user information 'administrator Portal Administrator'. The main content area is titled 'VPC Restore' and is divided into three sections: 'Restore List', 'Choose restore point', and 'Choose VPCs to restore'. The 'Available restore points' window is open, displaying a table of restore points. The selected restore point is '07/12/2023 11:02:59 AM', which shows '1 Vpc, 2 Subnet, 2 RouteTable, 5 SecurityGroup, 1 Acl, 1 NatGateway, 1 Interne...'. The 'Choose VPCs to restore' section shows a table with columns for Name and ID, and the 'jf_vpc_nat' VPC is selected.

Date	Changed Objects
07/28/2023 3:18:27 PM	No changes detected
07/28/2023 2:27:49 PM	No changes detected
07/27/2023 5:40:08 PM	1 SecurityGroup, 3 ManagedPrefixList
07/12/2023 11:02:59 AM	1 Vpc, 2 Subnet, 2 RouteTable, 5 SecurityGroup, 1 Acl, 1 NatGateway, 1 Interne...
06/05/2023 12:14:03 PM	No changes detected
06/02/2023 2:18:25 PM	No changes detected
06/02/2023 1:55:23 PM	No changes detected
05/10/2023 3:09:53 PM	No changes detected
05/10/2023 3:08:41 PM	No changes detected

Name ↑	ID
—	vpc-0daee8ed77c1d7de
<input checked="" type="checkbox"/> jf_vpc_nat	vpc-00296d95c10925b6
<input type="checkbox"/> pi-vpc	vpc-050ab318fb18ad39

Step 3. Specify IAM Identity

At the **Account** step of the wizard, choose whether you want to [use an IAM role](#) or [one-time access keys of an IAM user](#) to allow Veeam Backup for AWS to perform the restore operation. For information on the permissions that the IAM role or IAM user must have to perform the restore operation, see [VPC Configuration Restore IAM Permissions](#).

IMPORTANT

Make sure that the specified IAM role or one-time access keys belong to an AWS account in which you plan to restore the VPC configuration.

Specifying IAM Role

To specify an IAM role for restore, select the **IAM role** option and choose the necessary IAM role from the list.

For an IAM role to be displayed in the **IAM role** list, it must be added to Veeam Backup for AWS with the *Amazon VPC Restore* operation selected as described in section [Adding IAM Roles](#). If you have not added the necessary IAM role to Veeam Backup for AWS beforehand, you can do it without closing the **VPC Restore** wizard. To do that, click **Add** and complete the [Add IAM Role](#) wizard.

IMPORTANT

It is recommended that you check whether the selected IAM role has all the permissions required to perform the operation. If some permissions of the IAM role are missing, the restore operation will fail to complete successfully. To run the IAM role permission check, click **Check Permissions** and follow the instructions provided in section [Checking IAM Role Permissions](#).

The screenshot shows the 'VPC Restore' wizard in the 'Account' step. The interface includes a top navigation bar with the Veeam logo, 'VPC Restore', and server time information. A left sidebar contains navigation links: 'Restore List', 'Account' (selected), 'Restore Mode', 'Reason', and 'Summary'. The main content area is titled 'Select IAM role' and contains the following elements:

- A heading: 'Select IAM role'.
- A sub-heading: 'Specify an IAM role that will be used to access resources for the restore operation or provide temporary access keys.'
- Two radio button options: 'IAM role' (selected) and 'Temporary access keys'.
- Under 'IAM role', a dropdown menu showing 'acc_5393', a '+ Add' button, and a 'Check Permissions' button.
- Under 'Temporary access keys', two text input fields labeled 'Access key:' and 'Secret key:'.
- An information icon (i) with a text box: 'The keys are used to perform this operation only. They are not saved or stored. To learn what permissions are required for performing the operation, see the [User Guide](#).'

At the bottom of the wizard, there are three buttons: 'Previous', 'Next', and 'Cancel'.

Specifying One-Time Access Keys

To specify one-time access keys for restore, select the **Temporary access keys** option and use the **Access key** and **Secret key** fields to provide the access key ID and the secret access key.

NOTE

Veeam Backup for AWS does not store one-time access keys in the configuration database.

The screenshot shows the 'VPC Restore' configuration page in the Veeam Backup for AWS interface. The top navigation bar includes the Veeam logo, 'Veeam Backup for AWS', the server time 'Nov 16, 2023 3:16 PM', and the user 'administrator Portal Administrator'. A left sidebar contains navigation options: 'Restore List', 'Account', 'Restore Mode', 'Reason', and 'Summary'. The main content area is titled 'Select IAM role' and includes the instruction: 'Specify an IAM role that will be used to access resources for the restore operation or provide temporary access keys.' There are two radio button options: 'IAM role' (unselected) and 'Temporary access keys' (selected). Under 'IAM role', there is a dropdown menu showing 'acc_5393' and buttons for '+ Add' and 'Check Permissions'. Under 'Temporary access keys', there are two text input fields: 'Access key:' containing 'AKIAY4ZWOU4WMVRAGEVN' and 'Secret key:' containing a masked password. A blue information icon is present next to a note: 'The keys are used to perform this operation only. They are not saved or stored. To learn what permissions are required for performing the operation, see the User Guide.' At the bottom of the page, there are three buttons: 'Previous', 'Next', and 'Cancel'.

Step 4. Choose Restore Mode

At the **Restore Mode** step of the wizard, choose whether you want to restore the selected VPC configuration to the original or to a custom location. If you select the **Restore to new location, or with different settings** option, specify the target AWS Region where to restore the VPC configuration.

IMPORTANT

If you select the **Restore to a new location, or with different settings** option, consider that AWS Regions have different lists of the supported AWS services. VPC endpoints created using an AWS service that is not available in the target AWS Region will not be restored.

The screenshot shows the Veeam Backup for AWS interface. The top bar includes the Veeam logo, the text "Veeam Backup for AWS", the server time "Jun 14, 2024 2:50 PM", and user information "administrator Portal Administrator". The main content area is titled "VPC Restore" and contains a sidebar with navigation options: "Restore List", "Account", "Restore Mode" (highlighted), "Availability Zones", "Peering Connection", "Reason", and "Summary". The "Restore Mode" section has the following content:

Restore Mode
Choose whether you want to restore to the original location or to a new location, or with different settings.

- Restore to original location
Quickly restore the selected VPCs to the original location, with the same settings as the source VPCs.
- Restore to new location, or with different settings
Perform additional configuration steps to restore VPCs to a new location or to use settings that differ from the source settings.

A dropdown menu below the second option is set to "Europe (Paris)". At the bottom of the wizard, there are three buttons: "Previous", "Next", and "Cancel".

Step 5. Configure Availability Zone Mapping

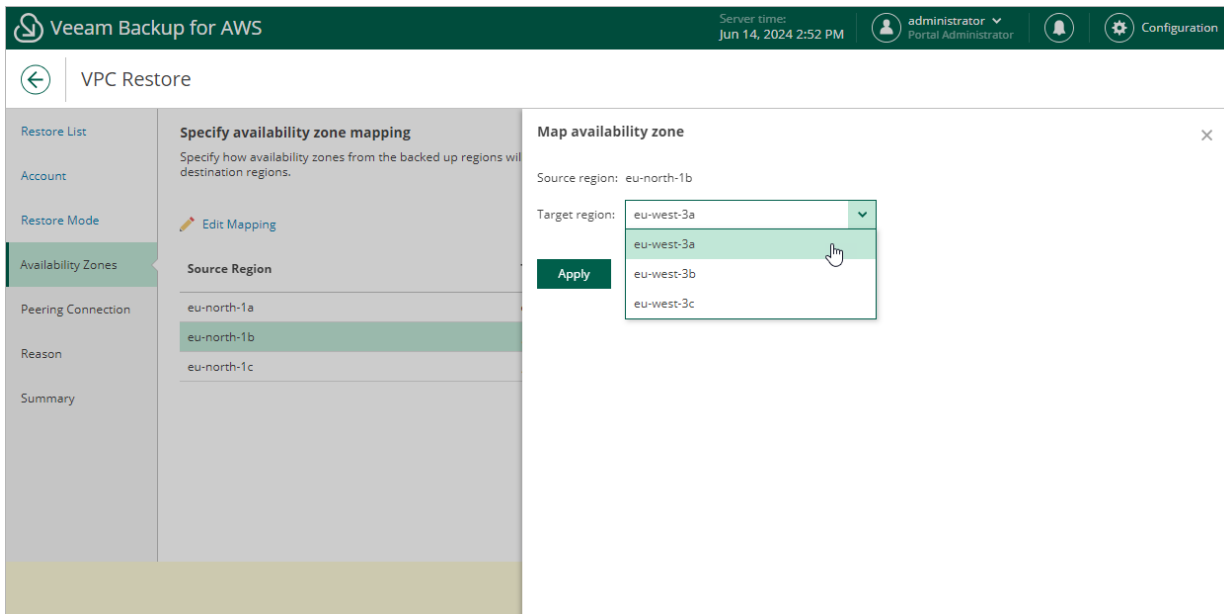
[This step applies only if you have selected the **Restore to new location, or with different settings** option at the **Restore Mode** step of the wizard]

At the **Availability Zones** step of the wizard, for each source Availability Zone, choose an Availability Zone in the target AWS Region to which VPC configuration items of the source Availability Zone will be restored:

1. Choose an Availability Zone from the list and click **Edit Mapping**.
2. In the **Map availability zone** window, select the target Availability Zone from the **Target region** drop-down list.
3. Click **Apply**.

IMPORTANT

The source and target AWS Regions may have different number of Availability Zones. In this case, Veeam Backup for AWS will automatically change subnet configuration for transit gateway VPC attachments, VPC endpoints and load balancers. After restoring, you can modify the subnet configuration manually in the AWS Management Console. To learn how to modify subnet configuration for VPC networking components, see [AWS Documentation](#).



Step 6. Review Peering Connection Settings

[This step applies only if you have selected the **Restore to a new location, or with different settings** option at the **Restore Mode** step of the wizard]

At the **Peering Connection** step of the wizard, review preconfigured VPC peering connection settings. You cannot modify the settings for the restored VPC configuration – by default, Veeam Backup for AWS will restore VPC peering connections as follows:

- If you restore both VPCs between which you have created a peering connection, Veeam Backup for AWS will create a peering connection between the restored VPCs in the target AWS Region.
- If you restore a VPC that has a peering connection to a VPC in the same AWS Region, Veeam Backup for AWS will create an inter-region peering connection between the restored VPC in the target AWS Region and the VPC with which the source VPC is peered in the source AWS Region.
- If you restore a VPC that has a peering connection to a VPC in another AWS Region, Veeam Backup for AWS will create an inter-region peering connection between the restored VPC in the target AWS Region and the VPC with which the source VPC is peered in the other AWS Region.

NOTE

VPC peering connections will have the *Pending Acceptance* status after restoring. To accept the restored VPC peering connections, use the AWS Management Console. For more information, see [AWS Documentation](#).

The screenshot shows the 'VPC Restore' wizard in Veeam Backup for AWS. The current step is 'Review peering connection settings'. The interface includes a navigation menu on the left with options: Restore List, Account, Restore Mode, Availability Zones, Peering Connection (selected), Reason, and Summary. The main content area displays a table of peering connections with the following data:

Name	ID	Requested VPC	Accepted VPC
NAT_Cali_S3_9969	pcx-03176746ba86189f4	vpc-00296d95c10925b0f	vpc-0a51f37ed81f
NAT_Oregon_s3	pcx-0fbc94abe5de1dd62	vpc-00296d95c10925b0f	vpc-08dc9acb4ac7

At the bottom of the interface, there are three buttons: 'Previous', 'Next', and 'Cancel'.

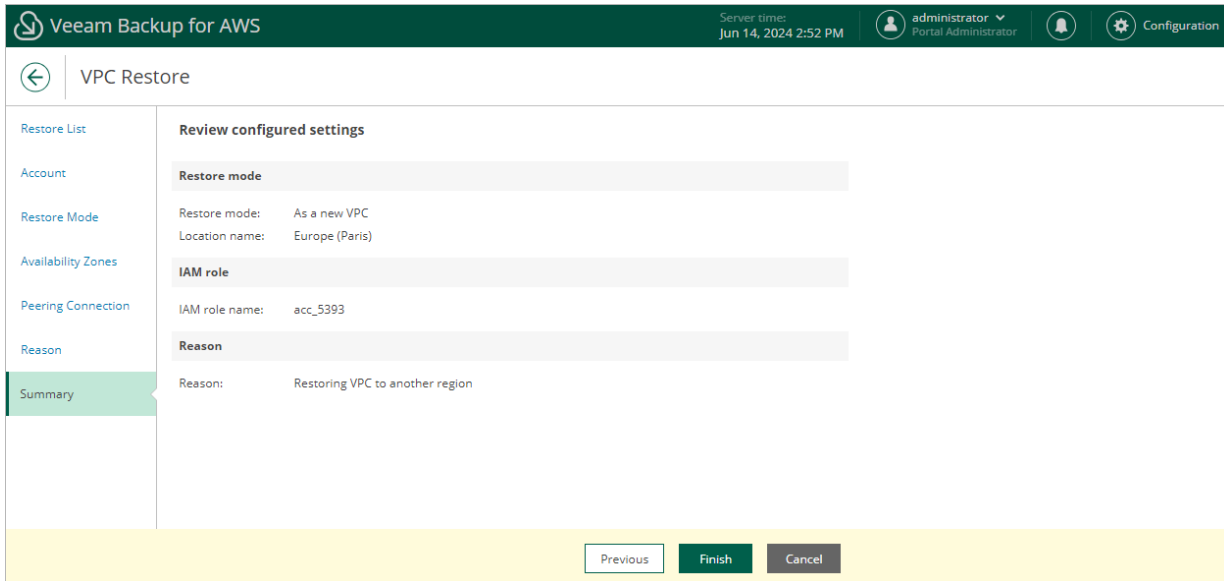
Step 7. Specify Restore Reason

At the **Reason** step of the wizard, you can specify a reason for restoring VPC configuration. The information you provide will be saved in the session history and you can reference it later.

The screenshot shows the Veeam Backup for AWS interface. The top navigation bar includes the Veeam logo, the text "Veeam Backup for AWS", the server time "Jun 14, 2024 2:52 PM", the user "administrator Portal Administrator", and a "Configuration" icon. The main content area is titled "VPC Restore" and features a left-hand navigation menu with the following items: "Restore List", "Account", "Restore Mode", "Availability Zones", "Peering Connection", "Reason" (which is highlighted in green), and "Summary". The "Reason" step is active, displaying the heading "Restore reason" and the instruction "Specify a reason for performing the restore operation." Below this is a text input field containing the text "Restoring VPC to another region". At the bottom of the wizard, there are three buttons: "Previous", "Next" (highlighted in green), and "Cancel".

Step 8. Finish Working with Wizard

At the **Summary** step of the wizard, review summary information and click **Finish**.



The screenshot shows the 'VPC Restore' wizard in the 'Summary' step. The interface includes a top navigation bar with the Veeam logo, server time (Jun 14, 2024 2:52 PM), user information (administrator, Portal Administrator), and a Configuration icon. A left sidebar lists the wizard steps: Restore List, Account, Restore Mode, Availability Zones, Peering Connection, Reason, and Summary (which is highlighted). The main content area is titled 'Review configured settings' and displays the following information:

- Restore mode:** As a new VPC
- Location name:** Europe (Paris)
- IAM role:** IAM role name: acc_5393
- Reason:** Restoring VPC to another region

At the bottom of the wizard, there are three buttons: 'Previous' (disabled), 'Finish' (active), and 'Cancel' (disabled).

Performing Selected Items Restore

In case of unexpected configuration changes, you can restore only specific items of the Amazon VPC configuration from a VPC configuration backup. Veeam Backup for AWS allows you to restore these items to the original location only.

How to Perform Selected Items Restore

To restore specific items of the VPC configuration, do the following:

1. [Launch the VPC Restore wizard.](#)
2. [Select a restore point and items to restore.](#)
3. [Specify an IAM identity for restore.](#)
4. [Specify a restore reason.](#)
5. [Finish working with the wizard.](#)

Step 1. Launch VPC Restore Wizard

To launch the **VPC Restore** wizard, do the following:

1. Navigate to **Protected Data > VPC**.
2. Select the configuration record for an AWS Region whose VPC configuration you want to restore.
3. Click **Restore > Selected Items**.

The screenshot shows the Veeam Backup for AWS interface. The top navigation bar includes the Veeam logo, server time (Jun 14, 2024 3:12 PM), and user information (administrator, Portal Administrator). The left sidebar shows the navigation menu with 'Protected Data' selected. The main content area is divided into tabs: EC2, Databases, File Systems, and VPC. The VPC tab is active, displaying a table of configuration records. The second record is selected, and a context menu is open over it, showing 'Restore' with sub-options 'Entire VPC' and 'Selected Items'. Below the table, the 'Configuration Details' section is visible, showing a list of resources with columns for Name, ID, Type, Modification Date, and State.

Name	ID ↑	Type	Modification Date	State
—	acl-086f91540d0ab2e21	Acl	06/05/2024 1:08:47 PM	Created
—	acl-0cf462f3a18242324	Acl	06/05/2024 1:08:47 PM	Created
—	acl-e315308bb	Acl	06/05/2024 1:08:47 PM	Created
—	dopt-5ada8832	DHCPOptionsSet	06/05/2024 1:08:47 PM	Created
nm-vb-v3 InternetGateway	igw-01345a5cbcac3a109	InternetGateway	06/05/2024 1:08:47 PM	Created
—	igw-a3b25fcb	InternetGateway	06/05/2024 1:08:47 PM	Created
com.amazonaws.eu-west-2...	pl-05cadb937455b8b91	ManagedPrefixList	06/05/2024 1:08:47 PM	Created
com.amazonaws.global.gr...	pl-082b2916fa6d7b4b7	ManagedPrefixList	06/05/2024 1:08:47 PM	Created

Step 2. Select Restore Point and Items to Restore

At the **Restore List** step of the wizard, select VPC configuration items you want to restore and a restore point that will be used to restore the selected items. By default, Veeam Backup for AWS uses the most recent valid restore point. However, you can restore the VPC configuration data to an earlier state.

1. To select the restore point:
 - a. In the **Choose restore point** section, click the link to the right of **Restore point**.
 - b. In the **Available restore points** window, select the necessary restore point and click **Apply**.
2. To select the VPC configuration items:
 - a. In the **Create restore list** section, click **Edit** and select an Amazon VPC resource that you want to restore.
 - b. In the **Edit restore list** window, click **Add to Restore List**.
 - c. In the **Item List** window, select check boxes next to the items that you want to restore, and click **Add**.
 - d. In the **Edit restore list** window, review the restore list and click **Apply**.

The screenshot displays the Veeam Backup for AWS interface during the VPC Restore process. The main window is titled 'VPC Restore' and is divided into several sections. On the left, there is a navigation pane with 'Restore List' selected. The main content area is split into three panels: 'Choose restore point', 'Create restore list', and 'Edit restore list (VPC)'. The 'Edit restore list (VPC)' panel is the active window, showing a search bar for 'Name or ID' and a 'State' dropdown. Below this is a table of VPC resources. The table has columns for 'Name', 'ID', 'Type', and 'State'. Two VPCs are listed: 'nm-5a-temp-lab VPC' with ID 'vpc-0386662d76b02e288' and 'vpc-6c2d5704'. Both are of type 'VPC' and have a state of 'Created'. The 'Apply' button is highlighted in green.

<input type="checkbox"/>	Name	ID ↑	Type	State
<input type="checkbox"/>	nm-5a-temp-lab VPC	vpc-0386662d76b02e288	VPC	Created
<input type="checkbox"/>	--	vpc-6c2d5704	VPC	Created

Step 3. Specify IAM Identity

At the **Account** step of the wizard, choose whether you want to [use an IAM role](#) or [one-time access keys of an IAM user](#) to allow Veeam Backup for AWS to perform the restore operation. For information on the permissions that the IAM role or IAM user must have to perform the restore operation, see [VPC Configuration Restore IAM Permissions](#).

IMPORTANT

After you click **Next**, Veeam Backup for AWS will use the permissions of the specified IAM role or IAM user to validate the restore list created at [step 2](#) of the wizard. If any of the VPC configuration items on which the selected items depend are missing from the current VPC configuration, Veeam Backup for AWS will open the **Missing Configuration Items** window with the list of the missing items. To proceed to the next step, click **Add**. The missing items will be automatically added to the restore list.

Specifying IAM Role

To specify an IAM role for restore, select the **IAM role** option and choose the necessary IAM role from the list.

For an IAM role to be displayed in the **IAM role** list, it must be added to Veeam Backup for AWS with the *Amazon VPC Restore* operation selected as described in section [Adding IAM Roles](#). If you have not added the necessary IAM role to Veeam Backup for AWS beforehand, you can do it without closing the **VPC Restore** wizard. To do that, click **Add** and complete the [Add IAM Role](#) wizard.

IMPORTANT

It is recommended that you check whether the selected IAM role has all the permissions required to perform the operation. If some permissions of the IAM role are missing, the restore operation will fail to complete successfully. To run the IAM role permission check, click **Check Permissions** and follow the instructions provided in section [Checking IAM Role Permissions](#).

The screenshot shows the Veeam Backup for AWS VPC Restore wizard at the Account step. The interface includes a top navigation bar with the Veeam logo, server time (Jun 14, 2024 3:12 PM), and user information (administrator, Portal Administrator). The main content area is titled "VPC Restore" and features a left-hand navigation menu with options: Restore List, Account (selected), Reason, and Summary. The "Select IAM role" section is active, with instructions: "Specify an IAM role that will be used to access resources for the restore operation or provide temporary access keys." There are two radio button options: "IAM role" (selected) and "Temporary access keys". Under "IAM role", a dropdown menu shows "admin-407355668422 (Created by nm at 9/26/2023...)" with a "+ Add" button and a "Check Permissions" button. Under "Temporary access keys", there are input fields for "Access key:" and "Secret key:". An information icon (i) is present with a note: "The keys are used to perform this operation only. They are not saved or stored. To learn what permissions are required for performing the operation, see the User Guide." At the bottom, there are three buttons: "Previous", "Next", and "Cancel".

Specifying One-Time Access Keys

To specify one-time access keys for restore, select the **Temporary access keys** option and use the **Access key** and **Secret key** fields to provide the access key ID and the secret access key.

NOTE

Veeam Backup for AWS does not store one-time access keys in the configuration database.

The screenshot shows the 'VPC Restore' configuration interface in Veeam Backup for AWS. The top navigation bar includes the Veeam logo, 'VPC Restore', server time 'Nov 16, 2023 3:33 PM', and user 'administrator Portal Administrator'. The left sidebar has 'Restore List', 'Account', 'Reason', and 'Summary'. The main area is titled 'Select IAM role' with the instruction 'Specify an IAM role that will be used to access resources for the restore operation or provide temporary access keys.' There are two radio button options: 'IAM role' (unselected) and 'Temporary access keys' (selected). Under 'IAM role', there is a dropdown menu showing 'ec2-restore-test-1' and buttons for '+ Add' and 'Check Permissions'. Under 'Temporary access keys', there are two text input fields: 'Access key:' containing 'AKIAY4ZWOU4WMVRAGEVN' and 'Secret key:' containing a masked string of dots. A blue information icon and a text box below the fields state: 'The keys are used to perform this operation only. They are not saved or stored. To learn what permissions are required for performing the operation, see the User Guide.' At the bottom, there are three buttons: 'Previous', 'Next', and 'Cancel'.

Step 4. Specify Restore Reason

At the **Reason** step of the wizard, you can specify a reason for the restore of VPC configuration items. The information you provide will be saved in the session history and you can reference it later.

The screenshot shows the Veeam Backup for AWS interface. The top navigation bar includes the Veeam logo, the text "Veeam Backup for AWS", the server time "Jun 14, 2024 3:13 PM", the user "administrator Portal Administrator", and a "Configuration" link. The main content area is titled "VPC Restore" and features a left-hand navigation pane with tabs for "Restore List", "Account", "Reason", and "Summary". The "Reason" tab is currently selected. The main area displays the heading "Restore reason" followed by the instruction "Specify a reason for performing the restore operation." Below this is a text input field with the label "Restore reason:" and the text "restoring vpcs" entered. At the bottom of the interface, there are three buttons: "Previous", "Next", and "Cancel".

Step 5. Finish Working with Wizard

At the **Summary** step of the wizard, review summary information and click **Finish**.

The screenshot shows the Veeam Backup for AWS interface during the VPC Restore process. The top navigation bar includes the Veeam logo, the text "Veeam Backup for AWS", the server time "Jun 14, 2024 3:13 PM", the user "administrator Portal Administrator", and a "Configuration" link. The main content area is titled "VPC Restore" and features a sidebar with navigation options: "Restore List", "Account", "Reason", and "Summary" (which is highlighted). The main panel displays "Review configured settings" with three sections: "Restore mode" (set to "Original location"), "IAM role" (set to "admin-407355668422 (Created by nm at 9/26/2023 7:19 AM)"), and "Reason" (set to "restoring vpcs"). At the bottom, there are three buttons: "Previous", "Finish" (highlighted in green), and "Cancel".

Instant Recovery

Veeam Backup & Replication allows you to use the Instant Recovery feature to restore EC2 instances from image-level backups to VMware vSphere and Microsoft Hyper-V environments, and to Nutanix AHV clusters. For more information, see the [Veeam Backup & Replication User Guide for VMware vSphere](#), [Veeam Backup & Replication User Guide for Microsoft Hyper-V](#) and [Veeam Backup for Nutanix AHV User Guide](#), section *Instant Recovery*.

IMPORTANT


Instant Recovery can be performed only using backup files stored in standard backup repositories for which you have specified access keys of an IAM user whose permissions are used to access the repositories. To learn how to specify credentials for the repositories, see sections [Creating New Repositories](#) and [Connecting to Existing Appliances](#).

Before you start the restore operation, make sure to add to the backup infrastructure a vCenter Server, a Microsoft Hyper-V server or a Nutanix AHV cluster that will manage restored EC2 instances, as described in the Veeam Backup & Replication User Guide, section [Adding VMware vSphere Servers](#), [Adding Microsoft Hyper-V Servers](#) or [Adding Nutanix AHV Cluster](#).

To perform Instant Recovery, do the following:

1. In the Veeam Backup & Replication console, open the **Home** view.
2. Navigate to **Backups > External Repository**.
3. Expand the backup policy that protects an EC2 instance that you want to recover, select the necessary EC2 instance and click **Instant Recovery** on the ribbon.
4. Select **VMware vSphere**, **Microsoft Hyper-V** or **Nutanix AHV**.

- Depending on the selected **Instant Recovery** option, complete the **Instant Recovery** wizard as described in the Veeam Backup & Replication User Guide, section [Performing Instant Recovery of Workloads to VMware vSphere VMs](#), [Performing Instant Recovery of Workloads to Hyper-V VMs](#) or [Performing Instant Recovery of Workloads to Nutanix AHV](#).



Exporting Disks

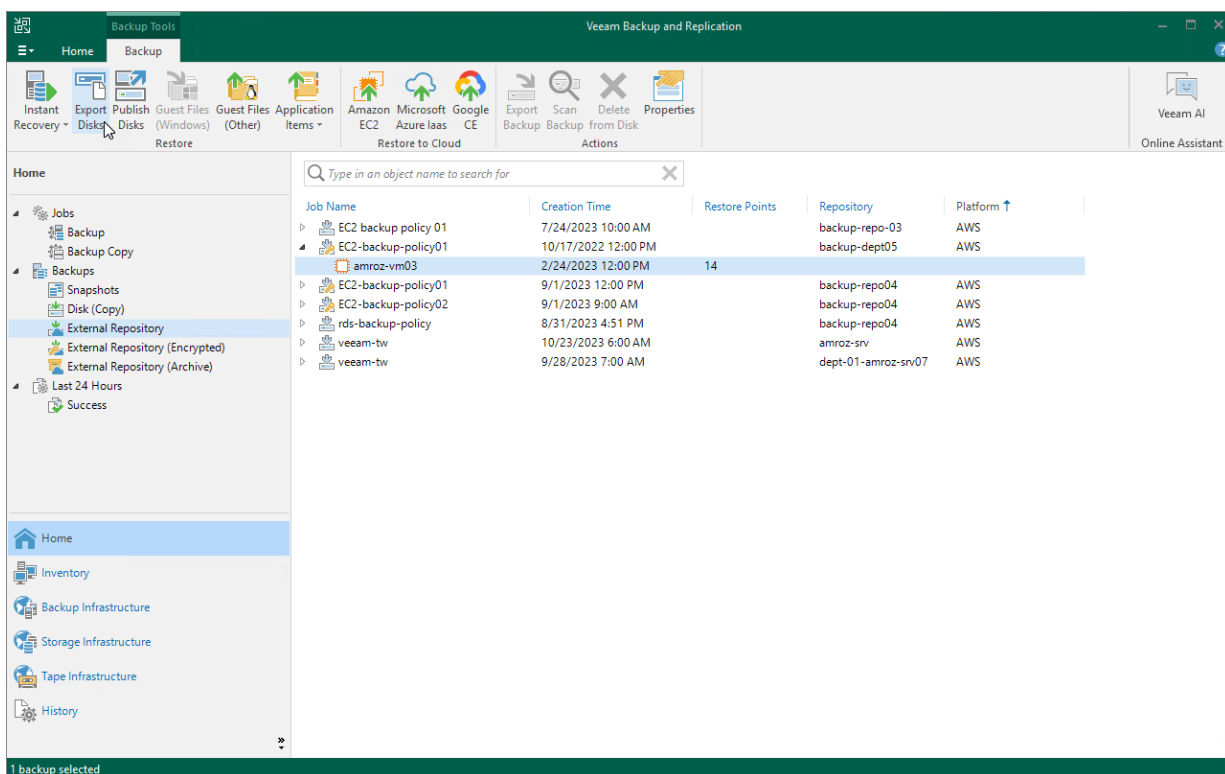
Veeam Backup & Replication allows you to export disks, that is, to restore EBS volumes of EC2 instances from image-level backups created by Veeam Backup for AWS and to convert them to the VMDK, VHD and VHDX formats. You can save the converted disks to any server added to the backup infrastructure or place the disks on a datastore connected to an ESXi host (for the VMDK disk format only). For more information, see the Veeam Backup & Replication, section [Disk Export](#).

IMPORTANT

Exporting Disks can be performed only using backup files stored in standard backup repositories for which you have specified access keys of an IAM user whose permissions are used to access the repositories. To learn how to specify credentials for the repositories, see sections [Creating New Repositories](#) and [Connecting to Existing Appliances](#).

To export EBS volumes of EC2 instance to the VMDK, VHD or VHDX format, do the following:

1. In the Veeam Backup & Replication console, open the **Home** view.
2. Navigate to **Backups > External Repository**.
3. Expand the backup policy that protects an EC2 instance whose volume you want to restore, select the necessary instance and click **Export Disk** on the ribbon.
4. Complete the **Export Disk** wizard as described in the Veeam Backup & Replication User Guide, section [Exporting Disks](#).



Publishing Disks


Veeam Backup & Replication allows you to publish point-in-time disks, that is, to mount specific EBS volumes of backed-up EC2 instances to any server to instantly access data in the read-only mode. You can copy the necessary files and folders to the target server, and perform an antivirus scan of the backed-up data. For more information, see the Veeam Backup & Replication User Guide, section [Disk Publishing \(Data Integration API\)](#).

IMPORTANT

Publishing Disks can be performed only using backup files stored in standard backup repositories for which you have specified access keys of an IAM user whose permissions are used to access the repository. To learn how to specify credentials for the repositories, see sections [Creating New Repositories](#) and [Connecting to Existing Appliances](#).

To publish volumes of an EC2 instance, do the following:

1. In the Veeam Backup & Replication console, open the **Home** view.
2. Navigate to **Backups > External Repository**.
3. Expand the necessary backup policy, select the EC2 instance whose volumes you want to publish and click **Publish Disks** on the ribbon.
4. Complete the **Publish Disks** wizard as described in the Veeam Backup & Replication User Guide, section [Publishing Disks](#).



Restoring to Microsoft Azure

Veeam Backup & Replication allows you to restore Amazon EC2 instances from image-level backups created with Veeam Backup for AWS to Microsoft Azure as Azure VMs. You can restore EC2 instances to any available restore point. For more information, see the Veeam Backup & Replication User Guide, section [Restore to Microsoft Azure](#).

IMPORTANT

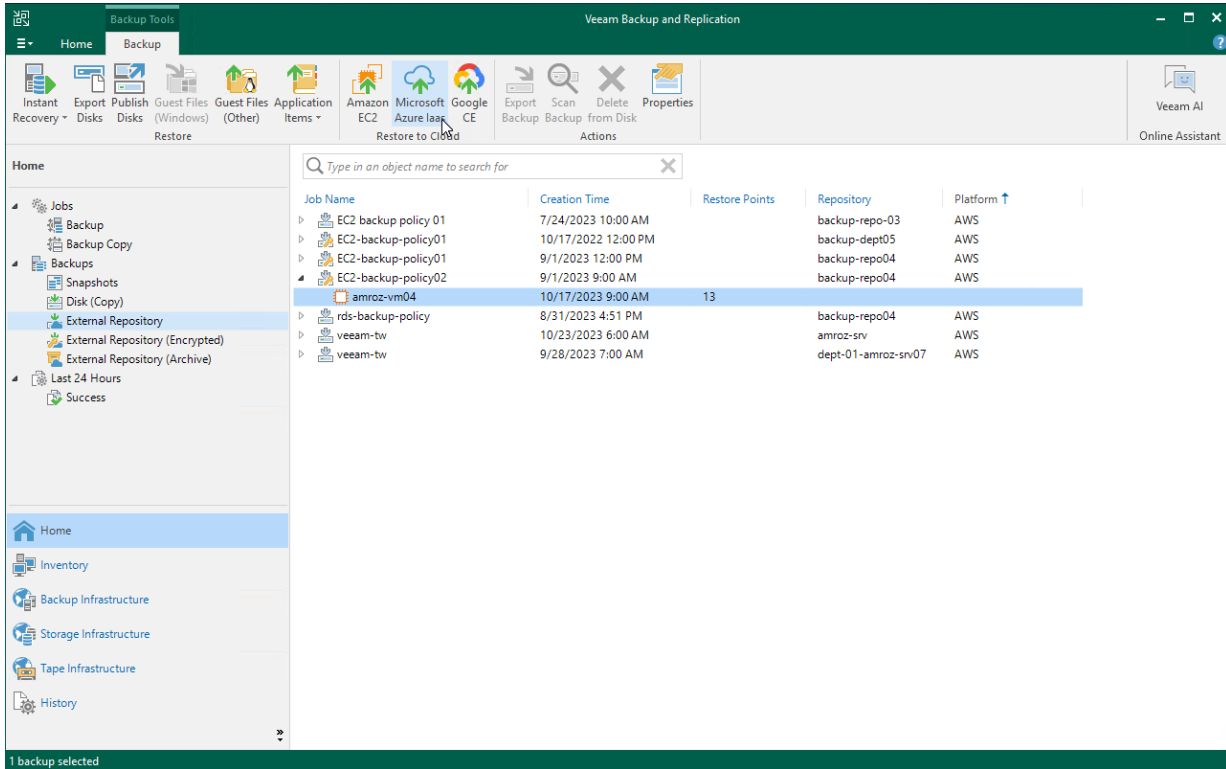
Restore to Microsoft Azure can be performed only using backup files stored in standard backup repositories for which you have specified access keys of an IAM user whose permissions are used to access the repositories. To learn how to specify credentials for the repositories, see sections [Creating New Repositories](#) and [Connecting to Existing Appliances](#).

Before you start the restore operation:

- Configure the initial settings of an Azure account or Azure Stack account as described in the Veeam Backup & Replication User Guide, section [Configuring Initial Settings](#).
- Check the limitations and prerequisites described in the Veeam Backup & Replication User Guide, section [Before You Begin](#).

To restore an EC2 instance to Microsoft Azure, do the following:

1. In the Veeam Backup & Replication console, open **Home** view.
2. Navigate to **Backups > External Repository**.
3. Expand the backup policy that protects an EC2 instance that you want to restore, select the necessary instance and click **Microsoft Azure IaaS** on the ribbon.
4. Complete the **Restore to Microsoft Azure** wizard as described in the Veeam Backup & Replication User Guide, section [Restoring to Microsoft Azure](#).



Restoring to Google Cloud

Veeam Backup & Replication allows you to restore Amazon EC2 instances from image-level backups created with Veeam Backup for AWS to Google Cloud as VM instances. You can restore EC2 instances to any available restore point. For more information, see the Veeam Backup & Replication User Guide, section [Restore to Google Compute Engine](#).


IMPORTANT

Consider the following:

- Restore to Google Cloud can be performed only using backup files stored in standard backup repositories for which you have specified access keys of an IAM user whose permissions are used to access the repositories. To learn how to specify credentials for the repositories, see sections [Creating New Repositories](#) and [Connecting to Existing Appliances](#).
- Before you start the restore operation, check the limitations and prerequisites described in the Veeam Backup & Replication User Guide, section [Before You Begin](#).

To restore an EC2 instance to Google Cloud, do the following:

1. In the Veeam Backup & Replication console, open the **Home** view.
2. Navigate to **Backups > External Repository**.
3. Expand the backup policy that protects an EC2 instance that you want to restore, select the necessary instance and click **Google CE** on the ribbon.
4. Complete the **Restore to Google Compute Engine** wizard as described in the Veeam Backup & Replication User Guide, section [Restoring to Google Compute Engine](#).



Restoring to Nutanix AHV

Veeam Backup & Replication allows you to restore EC2 instances from image-level backups created with Veeam Backup for AWS to Nutanix AHV as Nutanix AHV VMs. You can restore EC2 instances to any available restore point. For more information, see the Veeam Backup for Nutanix AHV User Guide, section [Performing Restore](#).

IMPORTANT

Restore to Nutanix AHV can be performed only using backup files stored in standard backup repositories for which you have specified access keys of an IAM user whose permissions are used to access the repositories. To learn how to specify credentials for the repositories, see sections [Creating New Repositories](#) and [Connecting to Existing Appliances](#).


Before you start the restore operation:

- Configure the backup infrastructure as described in the Veeam Backup for Nutanix AHV User Guide, section [Deployment](#).
- If you restore EC2 instances from a standard backup, make sure that this backup have been copied to an on-premises backup repository as described in the Veeam Backup & Replication User Guide, section [Creating Backup Copy Jobs for VMs and Physical Machines](#).
- If you restore EC2 instances from an archived backup stored in a scale-out backup repository, make sure that this backup have been retrieved from an archive as described in the Veeam Backup & Replication User Guide, section [Retrieving Backup Files](#).

To restore an EC2 instance to a Nutanix AHV cluster, do the following:

1. In the Veeam Backup & Replication console, open the **Home** view.
2. Navigate to **Backups > Disk (Copy)**.
3. Expand the necessary backup policy, select the EC2 instance that you want to restore and click **Entire VM** on the ribbon.

4. Complete the **Restore to Nutanix AHV** wizard as described in the Veeam Backup for Nutanix AHV User Guide, section [Restoring VMs Using Veeam Backup & Replication Console](#).



Reviewing Dashboard

Veeam Backup for AWS comes with an **Overview** dashboard that provides at-a-glance real-time overview of the protected AWS resources and allows you to estimate the overall backup performance. The dashboard includes the following widgets:

- **Sessions in Last 24 Hours** – displays the number of sessions started for data protection or disaster recovery operations during the past 24 hours that completed successfully, the number of sessions that completed with warnings, the number of sessions that completed with errors, and the number of sessions that are currently running.

To get more information on the sessions, click either **View Session Logs** or any of the widget rows. In the latter case, the **Session Logs** page will show only those sessions that have the same status as that clicked in the widget.

For more information on the **Session Logs** page, see [Viewing Session Statistics](#).

- **Successful Policy Tasks** – displays the number of snapshots, snapshot replicas, backups and archived backups successfully created by backup policies during a specific time period (the past 24 hours by default).

To specify the time period, click the link next to the **Schedule** icon. To get more information on the created snapshots, backups or archived backups, click any of the widget rows. In the latter case, the **Session Logs** page will show only those sessions during which Veeam Backup for AWS created the same items as that clicked in the widget.

For more information on the **Session Logs** page, see [Viewing Session Statistics](#).

- **Protected Workloads** – displays the number of AWS resources that got protected by Veeam Backup for AWS during a specific time period (the past 24 hours by default).

To specify the time period, click the link next to the **Schedule** icon. To get more information on the protected resources, click any of the widget rows.

For more information on the available resources, their properties and the actions you can perform for the resources, see [Viewing Available Resources](#).

- **Storage Usage** – displays the amount of storage space that is currently consumed by restore points created by Veeam Backup for AWS in Amazon S3 buckets. The widget also displays the total amount of storage space used in the S3 Standard, S3 Glacier Flexible Retrieval and S3 Glacier Deep Archive storage classes explicitly.
- **Top Policies** – shows top backup policies for execution time (including retries). For each policy, the widget also calculates the growth rate to detect whether it took less or more time for the policy to complete in comparison with the previous policy run.
- **Bottlenecks Overview** – is designed to help you avoid possible backup bottlenecks.

The **Policy sizing** widget verifies whether the appliance CPU and memory resources are enough to process all enabled backup policies and whether the backup policies are sized correctly. Note that one backup policy should not protect more than 250 resources for Veeam Backup for AWS to work properly.

The **CPU quota** widget analyzes the amount of CPU quota across all regions to detect whether the quota has already been reached in any of the regions, and if Veeam Backup for AWS could not deploy a worker instance in that region during a backup or restore process. For more information on worker profiles, see [Managing Worker Profiles](#).

The **Appliance disk usage** widget analyzes memory usage on the backup appliance, and displays a warning if the memory usage keeps breaching the preconfigured threshold (80%) for 60 minutes in a row. If the problem persists, increase the EBS volume size of the backup appliance or open a [support case](#) to remove the unnecessary data from the configuration database.

TIP

To prevent occasional runtime issues caused by multiple concurrent operations running on the backup appliance, you can allow the system to allocate additional resources in case of memory shortage. For more information, see [Appendix D. Enabling Swap Partition](#).

The screenshot displays the Veeam Backup for AWS dashboard with the following data:

Sessions in Last 24 Hours

Failed	1
Warning	2
Success	93
Running now	0

Successful Policy Tasks (Last 24 hours)

Snapshots	4 of 4	100%
Replicas	4 of 4	100%
Backups	4 of 4	100%
Archives	4 of 4	100%

Protected Workloads (Last 24 hours)

EC2 instances	2 of 31	6%
Databases	0 of 14	0%
File systems	0 of 4	0%

Storage Usage

- Snapshots: 118
- Replicas: 9
- Backups: 103 GB
- Archives: 11 GB
- Total: 114 GB
- S3 Standard: 103 GB
- S3 Glacier: 11 GB
- S3 Glacier Deep Archive: 0 Bytes

Top Policies (By duration increase)

Policy	Duration	Start Time	Percentage
policy-fsx-487569979969	4 min 23 sec	06/09 08:29 AM	+74.06%
policy-redshift-487569979969	2 min 18 sec	06/04 09:21 PM	+25%
policy-dynamo-487569979969	2 min 47 sec	06/05 12:32 PM	-19%
policy-ec2-487569979969	4 min 54 sec	06/12 01:58 PM	-47%

Bottlenecks Overview

- Policy sizing: 06/13/2024 9:24 AM OK
- CPU quota: 06/12/2024 2:05 PM Available
- Appliance disk usage: 06/13/2024 9:30 AM OK

Viewing Session Statistics

For each performed data protection or disaster recovery operation, Veeam Backup for AWS starts a new session and stores its records in the configuration database.

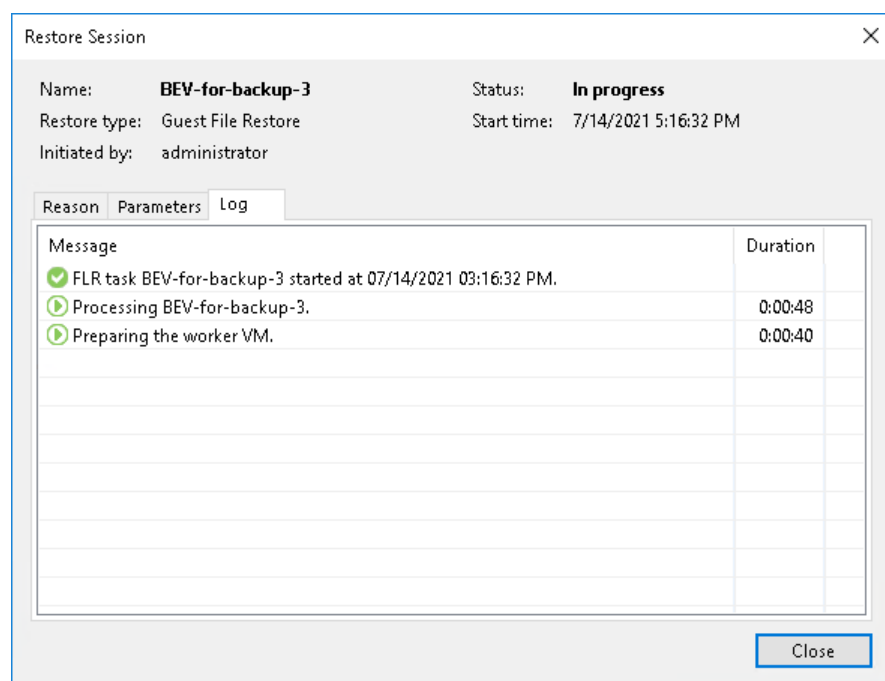
Viewing Session Statistics Using Console

You can track real-time statistics of all running and completed operations on the **Jobs**, **Last 24 hours** and **Running** nodes. For more information, see Veeam Backup & Replication User Guide, sections [Viewing Real-Time Statistics](#) and [Viewing Job Session Results](#).

Veeam Backup & Replication also allows you track statistics of data recovery operations initiated from Veeam Backup for AWS. To do that, do either of the following:

- In the Veeam Backup & Replication console, open the **Home** view and navigate to **Last 24 hours**. In the working area, double-click the necessary restore session.
Alternatively, select the session and click **Statistics** on the ribbon.
- In the Veeam Backup & Replication console, open the **History view** and navigate to **Restore**. In the working area, double-click the necessary restore session.
Alternatively, select the session and click **Statistics** on the ribbon.

The **Restore Session** window will display restore session details such as the name of the VM instance whose data is being restored, the account under which the session has started, the session status and duration, information on the restore point selected for the restore operation, and the list of tasks performed during the session.



Viewing Session Statistics Using Web UI

You can track real-time statistics of all running and completed operations on the **Session Logs** page. To view the full list of tasks executed during an operation, click the link in the **Status** column. To view the full list of instances processed during an operation, click the link in the **Items** column.

TIPS

- You can filter sessions by the name of the operation. To do that, enter a name in the search field. In addition to the operation name, you can also filter sessions by their status and type, as well as by the resource for which the session was run. To do that, click **Filter** and select the necessary options.
- You can save the full list of sessions as a .CSV or .XML file. To do that, click **Export to** and select the necessary format.
- If you want to specify the time period during which Veeam Backup for AWS must keep session records in the configuration database, follow the instructions provided in section [Configuring Global Retention Settings](#).

The screenshot displays the Veeam Backup for AWS interface. The main window is titled "RDS Policy Snapshot: RDS Snapshot and Replicas". The left sidebar shows the navigation menu with "Session Logs" selected. The main content area is divided into two sections: "Session Status" and "Session Log".

Session Status

Result	Start Time	End Time	Duration
Warning	09/10/2021 11:00:01 AM	09/10/2021 11:00:03 AM	2 sec
Warning	09/10/2021 5:00:09 AM	09/10/2021 5:00:12 AM	3 sec

Session Log

Start Time	Status	Description	Execution Duration
09/10/2021 11:00:01 AM	Warning	The resource is already protected by another policy: le-mariadb	—
09/10/2021 11:00:02 AM	Success	Snapshot policy started at 09/10/2021 11:00:01 AM.	—
09/10/2021 11:00:02 AM	Warning	There are no resources to process	—
09/10/2021 11:00:03 AM	Success	All instances have been queued for processing	1 sec
09/10/2021 11:00:03 AM	Warning	Session finished with warning at 09/10/2021 11:00:03 AM.	—

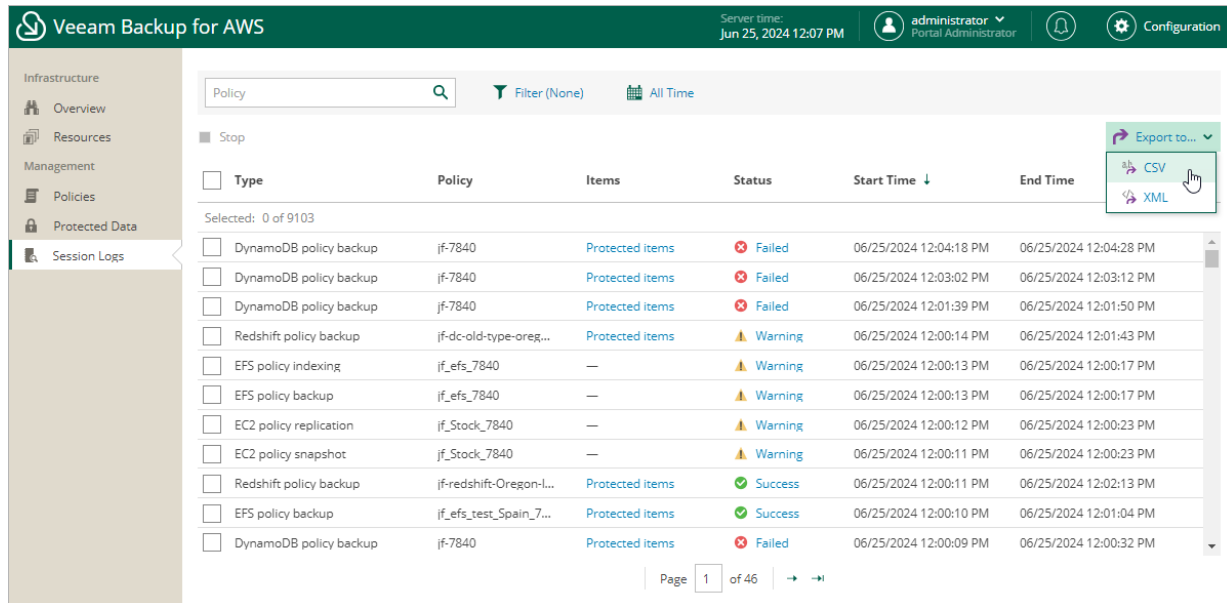
A "Close" button is located at the bottom right of the dialog box.

Collecting Object Properties

You can export properties of objects managed by Veeam Backup for AWS as a single file in the CSV or XML format. To do that, navigate to the necessary tab and click **Export to**. Veeam Backup for AWS will save the file with the exported data to the default download directory on the local machine.

NOTE

Even if you try to export properties of a specific object, Veeam Backup for AWS will still export all properties of all objects present on the currently opened tab.



The screenshot displays the Veeam Backup for AWS web interface. The top navigation bar includes the Veeam logo, the text 'Veeam Backup for AWS', the server time 'Jun 25, 2024 12:07 PM', the user 'administrator Portal Administrator', and a 'Configuration' button. The left sidebar shows navigation options: Infrastructure, Overview, Resources, Management, Policies, Protected Data, and Session Logs. The main content area features a search bar with 'Policy' entered, a 'Filter (None)' button, and an 'All Time' view selector. Below this is a table with columns: Type, Policy, Items, Status, Start Time, and End Time. The table lists various backup policies with their respective statuses (Failed, Warning, Success). An 'Export to...' dropdown menu is open in the top right corner, showing options for 'CSV' and 'XML'. The bottom of the table shows 'Page 1 of 46'.

Type	Policy	Items	Status	Start Time	End Time	
<input type="checkbox"/>	DynamoDB policy backup	if-7840	Protected items	Failed	06/25/2024 12:04:18 PM	06/25/2024 12:04:28 PM
<input type="checkbox"/>	DynamoDB policy backup	if-7840	Protected items	Failed	06/25/2024 12:03:02 PM	06/25/2024 12:03:12 PM
<input type="checkbox"/>	DynamoDB policy backup	if-7840	Protected items	Failed	06/25/2024 12:01:39 PM	06/25/2024 12:01:50 PM
<input type="checkbox"/>	Redshift policy backup	if-dc-old-type-oreg...	Protected items	Warning	06/25/2024 12:00:14 PM	06/25/2024 12:01:43 PM
<input type="checkbox"/>	EFS policy indexing	if_efs_7840	—	Warning	06/25/2024 12:00:13 PM	06/25/2024 12:00:17 PM
<input type="checkbox"/>	EFS policy backup	if_efs_7840	—	Warning	06/25/2024 12:00:13 PM	06/25/2024 12:00:17 PM
<input type="checkbox"/>	EC2 policy replication	if_Stock_7840	—	Warning	06/25/2024 12:00:12 PM	06/25/2024 12:00:23 PM
<input type="checkbox"/>	EC2 policy snapshot	if_Stock_7840	—	Warning	06/25/2024 12:00:11 PM	06/25/2024 12:00:23 PM
<input type="checkbox"/>	Redshift policy backup	if-redshift-Oregon-l...	Protected items	Success	06/25/2024 12:00:11 PM	06/25/2024 12:02:13 PM
<input type="checkbox"/>	EFS policy backup	if_efs_test_Spain_7...	Protected items	Success	06/25/2024 12:00:10 PM	06/25/2024 12:01:04 PM
<input type="checkbox"/>	DynamoDB policy backup	if-7840	Protected items	Failed	06/25/2024 12:00:09 PM	06/25/2024 12:00:32 PM

Updating Veeam Backup for AWS

Veeam Backup for AWS allows you to check for new product versions and available package updates. It is recommended that you timely install available package updates to avoid performance issues while working with the product. For example, timely installed security updates may help you prevent potential security issues and reduce the risk of compromising sensitive data.

IMPORTANT

Updating the backup appliance in the unattended mode or using third-party tools is not supported.

Updating Appliances Using Console

Starting from version 6a, you can upgrade backup appliances from the Veeam Backup & Replication console only. Upgrade to Veeam Backup for AWS version 8.0 is supported from Veeam Backup for AWS version 4.0 or later. To upgrade from an earlier version, you must first perform upgrade to version 4.0 as described in section [Installing Updates](#).

IMPORTANT

Before you upgrade a backup appliance, check whether the Veeam Backup for AWS version is compatible with the current version of AWS Plug-in for Veeam Backup & Replication. For more information, see [System Requirements](#).

AWS Plug-in for Veeam Backup & Replication allows you to download and install new available Veeam Backup for AWS versions and product updates:


1. In the Veeam Backup & Replication console, open the **Backup Infrastructure** view.
2. Navigate to **Managed Servers**.
3. Select the necessary backup appliance and click **Upgrade Appliance** on the ribbon.

Alternatively, right-click the appliance and select **Upgrade**.

During upgrade, Veeam Backup & Replication updates permissions only of the *Default Backup Restore* IAM role created on the backup appliance. However, the Amazon FSx and Amazon Redshift resources and operations for them will not be automatically selected for this role in the **Add IAM Role** wizard. This means that, Veeam Backup for AWS will not be able to filter the role in the backup and restore settings of FSx and Redshift resources. To work around the issue, modify the *Default Backup Restore* IAM role settings manually and select the necessary resources and operations, as described in section [Editing IAM Role Settings](#). To update permissions of custom IAM roles added to the appliance, follow the instructions provided in section [Updating IAM Roles](#).

NOTE

When you upgrade to Veeam Backup for AWS version 8.0 from Veeam Backup for AWS version 6.0 or earlier, the backup appliance operating system is upgraded to Ubuntu 22.04 LTS and the configuration database is upgraded to PostgreSQL 15. For more information on the upgrade process, see [Upgrading to 8.0 from Version 6.0 or Earlier](#).



Upgrading to 8.0 from Version 6.0 or Earlier

To upgrade Veeam Backup for AWS to version 8.0, a backup appliance must be running version 4.0 or later. To upgrade the appliance, check the [prerequisites](#) and follow the instructions provided in section [Updating Appliances Using Console](#).

When you perform upgrade to version 8.0 from Veeam Backup for AWS version 6.0 or earlier, the backup appliance operating system is upgraded from Ubuntu 18.04 LTS to Ubuntu 22.04 LTS, and the configuration database is upgraded to PostgreSQL 15. Consider that during upgrade the original root volume of the backup appliance will be replaced with the new one.

How Upgrade to Version 8.0 Works

When upgrading backup appliances to version 8.0 from Veeam Backup for AWS version 6.0 or earlier, Veeam Backup & Replication performs the following steps:

1. Instructs Veeam Backup for AWS to create a cloud-native snapshot of the original appliance. If the upgrade process fails, the appliance will be reverted to the created snapshot.

Consider that this snapshot will be automatically removed by Veeam Backup & Replication from AWS after the upgrade operation completes successfully.
2. Upgrades version of the appliance configuration database to PostgreSQL 15: creates a new PostgreSQL database on the data volume, copies all configuration data to this database and removes the old database.
3. Saves the following configuration files and settings to the data volume: the appliance configuration file (`/etc/awsbackup/config.ini`), nginx configuration files (`/etc/nginx/nginx.conf`, `/etc/nginx/proxy_params`), users, MFA and time zone settings, and Linux environment (`/etc/ssh/`, `/root/`, `/home/`).
4. Deploys a temporary EC2 instance from the Ubuntu 22.04 LTS image.
5. Detaches the root volume from the newly created EC2 instance and removes the EC2 instance.
6. Detaches the outdated root volume and attaches the new root volume to the original appliance.
7. Removes the outdated root volume from AWS infrastructure.
8. Restores the configuration files and settings saved at step 3 to the new root volume.

Limitations and Prerequisites

Before you start the upgrade process, consider the following requirements and limitations:

- The IAM user whose access keys specified when [deploying a backup appliance](#) or [connecting to the appliance](#) must be assigned permissions required to perform upgrade. For the list of required permissions, see [Plug-in Permissions](#).
- Outbound internet access must be allowed from the backup appliance to the PostgreSQL Apt Repository (`apt.postgresql.org`, `apt-archive.postgresql.org`) through port **80** over the HTTP protocol.
- Outbound internet access must be allowed from the backup appliance to the PostgreSQL through port **443** over the HTTPS protocol to download the repository key <https://www.postgresql.org/media/keys/ACCC4CF8.asc>.
- Outbound internet access must be allowed from the backup appliance to the [Veeam Update Repository](#) through port **443** over the HTTPS protocol.

- Outbound internet access must be allowed from the backup appliance to the Ubuntu Security Repository (security.ubuntu.com) through port **80** over the HTTP protocol.
- During upgrade, the data volume of the backup appliance will temporarily contain files of 2 databases. That is why the size of the data volume must be twice the total amount of storage space used by the configuration database.
- During upgrade, Veeam Backup & Replication will create a new root volume with the default settings. That is why if you have modified the root disk settings, for example, have increased the volume size or enabled volume encryption, these settings will not be transferred, and custom 3rd-party software installed on the backup appliance will not be migrated.
- During upgrade, Veeam Backup & Replication will overwrite custom settings of the `/etc/fstab` configuration file on the backup appliance with the default settings. That is why if you have attached an additional EBS volume to the backup appliance, you must re-mount the volume by adding its label or UUID to the `/etc/fstab` file.
- After the upgrade process completes, the original root volume will be automatically deleted from AWS.

Eliminating Warnings Received During Upgrade

During upgrade to version 8.0 from Veeam Backup for AWS version 6.0 or earlier, Veeam Backup & Replication will verify whether the IAM user whose access keys are used to connect to the appliance has sufficient permissions to upgrade the appliance. If some permissions are missing, you will receive a warning.

You can manually grant missing permissions to the IAM user in AWS or instruct Veeam Backup & Replication to do it:

- If you want to grant the missing permissions manually, do the following:
 - a. Click **Copy permissions to Clipboard**.
Note that the list of copied permissions will contain all the permissions required to perform the upgrade operation, not the list of missing permissions.
 - b. In AWS, create an IAM policy with the missing permissions and attach the policy to the IAM user whose access key are used to connect to the appliance.
To learn how to create IAM policies, see [Appendix B. Creating IAM Policies in AWS](#).
 - c. Back to the Veeam Backup & Replication console, click **Proceed**.
- If you want to instruct Veeam Backup & Replication to grant the missing permissions automatically, click **Grant** and provide one-time access keys of an IAM user that is [authorized to grant IAM permissions](#) in the opened window. Note that the specified user must belong to the same AWS account in which the backup appliance is deployed.

Veeam Backup & Replication will create an IAM policy with missing permissions and attach the policy to the IAM user whose permissions are used to connect to the appliance.

NOTE

Veeam Backup & Replication does not store the provided one-time access keys in the configuration database.

Updating Appliances Using Web UI

Veeam Backup for AWS automatically notifies you about newly released product versions and package updates available for the operating system running on the backup appliance. However, starting from Veeam Backup for AWS version 6a, you can use the Veeam Backup for AWS Web UI to install package updates only. To upgrade Veeam Backup for AWS to new versions, follow the instructions provided in section [Upgrading Appliances](#).

IMPORTANT

You can update the standalone backup appliance using the Veeam Updater service only. Updating the backup appliance in the unattended mode or using third-party tools is not supported.

Upgrading Appliances

Upgrade to Veeam Backup for AWS version 8.0 is supported from Veeam Backup for AWS version 4.0 or later. To upgrade from an earlier version, you must first perform upgrade to version 4.0 as described in section [Installing Updates](#).

IMPORTANT

Before you upgrade the backup appliance, make sure that all backup policies are both disabled and stopped, and no restore tasks are currently executing. Otherwise, the upgrade process will interrupt the running activities, which may result in data loss.

To upgrade the backup appliance, do the following:

1. Install AWS Plug-in for Veeam Backup & Replication as described in section [Deployment](#).
If you do not have a valid Veeam Backup & Replication license, you can download a [30-day trial version](#) of the product.
2. Add the backup appliance to the Veeam Backup & Replication infrastructure as described in section [Connecting to Existing Appliances](#).
When connecting to the backup appliance, Veeam Backup & Replication will display a warning notifying you that the appliance must be upgraded. Acknowledge the warning to allow Veeam Backup & Replication to automatically upgrade the appliance to the necessary version.

NOTE

When you add the backup appliance to the Veeam Backup & Replication infrastructure, the license installed on the appliance is replaced with the license installed on the backup server. Protected instances start consuming license units from the license installed on the Veeam Backup & Replication server. However, as soon as you remove the backup appliance from the Veeam Backup & Replication infrastructure, Veeam Backup for AWS will continue using the license that had been used before you added the Veeam Backup for AWS appliance to the Veeam Backup & Replication infrastructure.

For more information on licensing scenarios, see [Licensing](#).

3. [This step applies only if the backup appliance has not been upgraded at step 2] Upgrade the backup appliance as described in the section [Updating Appliances Using Web UI](#).
4. After the upgrade process completes, you can remove the backup appliance from the Veeam Backup & Replication infrastructure, as described in section [Removing Appliances](#), if you do not plan to further manage this appliance from the Veeam Backup & Replication console.
If you remove the backup appliance from the backup infrastructure, you will no longer be able to create image-level backups of PostgreSQL DB instances and protect Redshift clusters, DynamoDB tables and FSx file systems. For more information, see [Integration with Veeam Backup & Replication](#).


NOTE

When you upgrade to Veeam Backup for AWS version 8.0 from Veeam Backup for AWS version 6.0 or earlier, the backup appliance operating system is upgraded to Ubuntu 22.04 LTS and the configuration database is upgraded to PostgreSQL 15. For more information on the upgrade process, see [Upgrading to Veeam Backup for AWS 8.0 from Version 6.0 or Earlier](#).

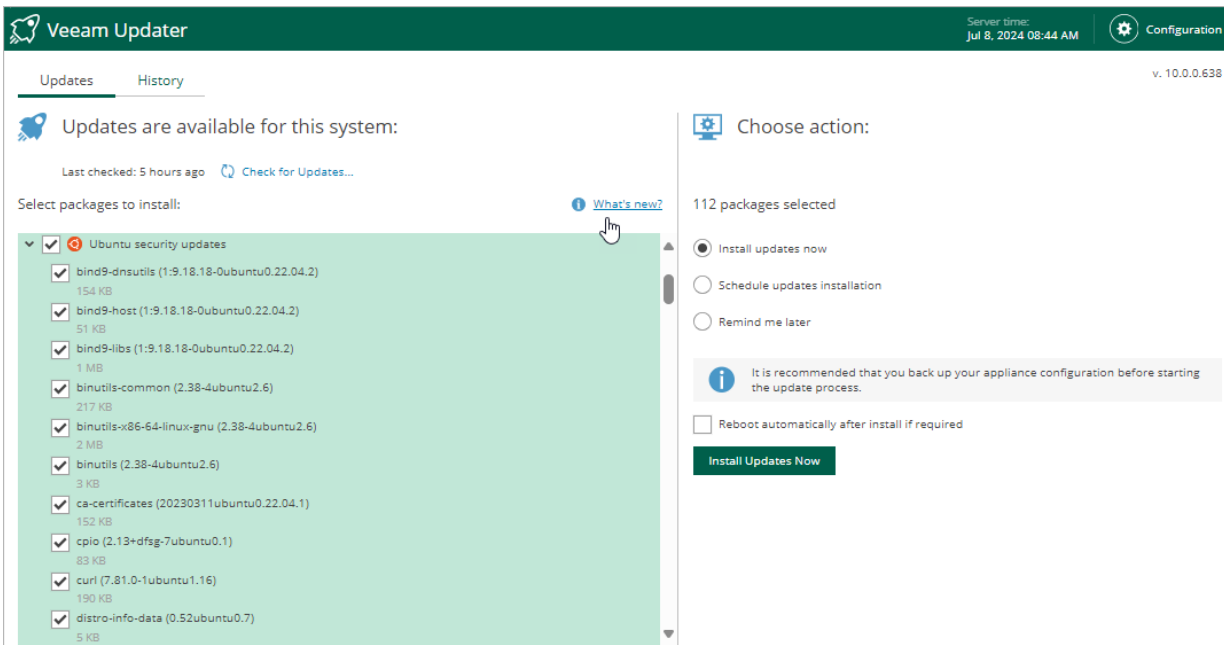
Checking for Updates

Veeam Backup for AWS automatically notifies you about newly released product versions and package updates available for the operating system running on the backup appliance. However, you can check for available updates manually if required:

1. Switch to the **Configuration** page.
2. Navigate to **Support Information > Updates**.
3. Click **Check and View Updates**.



If new updates are available, Veeam Backup for AWS will display them on the **Updates** tab of the **Veeam Updater** page. To view detailed information on an update, select the check box next to the update and click **What's new?**



Installing Updates

To download and install new available product and package updates using the Veeam Updater service, you can use either of the following options:

- [Install updates immediately](#)
- [Schedule update installation](#)

You can also [set a reminder to send update notifications](#).

IMPORTANT

Consider the following:

- You can update the standalone backup appliance using the Veeam Updater service only. Updating the backup appliance in the unattended mode or using third-party tools is not supported.
- You can update the backup appliance managed by a Veeam Backup & Replication server from the Veeam Backup & Replication console as described in section [Updating Appliances Using Console](#).

Installing Updates

IMPORTANT

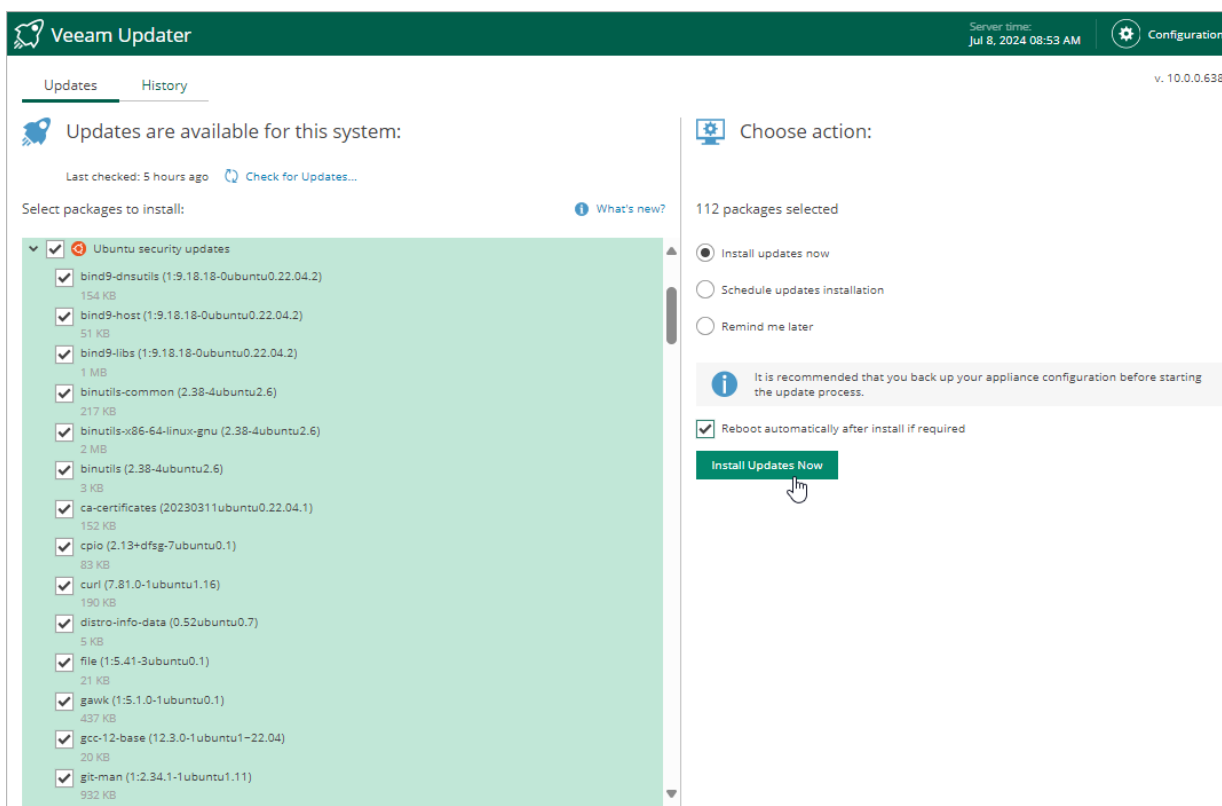
Before you install a product update, make sure that all backup policies are both disabled and stopped, and no restore tasks are currently executing. Otherwise, the update process will interrupt the running activities, which may result in data loss.

To download and install available product and package updates:

1. Open the **Veeam Updater** page. To do that:
 - a. Switch to the **Configuration** page.
 - b. Navigate to **Support Information**.
 - c. On the **Updates** tab, click **Check and View Updates**.
2. On the **Veeam Updater** page, do the following:
 - a. In the **Updates are available for this system** section, select check boxes next to the necessary updates.
 - b. In the **Choose action** section, select the **Install updates now** option, select the **Reboot automatically after install if required** check box to allow Veeam Backup for AWS to reboot the backup appliance if needed, and then click **Install Updates Now**.

NOTE

The updater may require you to read and accept the Veeam license agreement and the 3rd party components license agreement. If you reject the agreements, you will not be able to continue installation.



Veeam Backup for AWS will download and install the updates; the results of the installation process will be displayed on the **History** tab. Keep in mind that it may take several minutes for the installation process to complete.

NOTE

When installing product and package updates, Veeam Backup for AWS restarts all services running on the backup appliance, including the Web UI service. That is why Veeam Backup for AWS will log you out when the update process completes.

Scheduling Update Installation

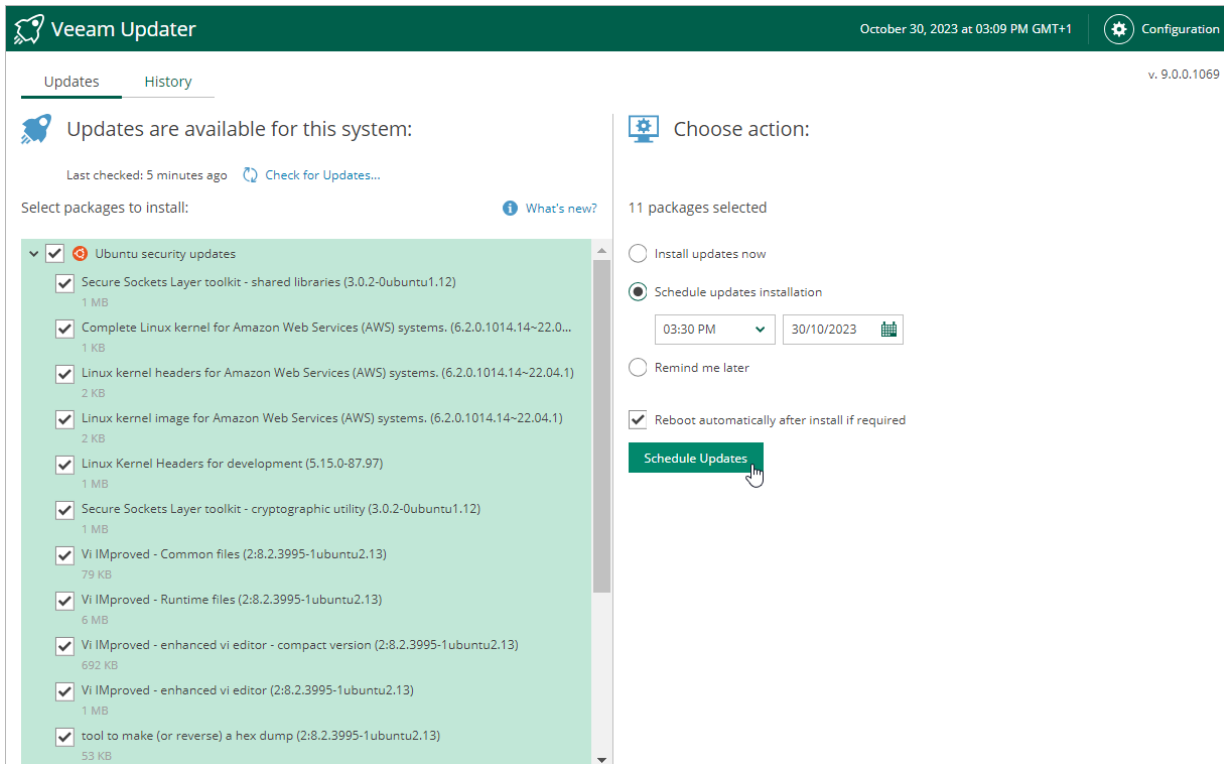
You can instruct Veeam Backup for AWS to automatically download and install available product and package updates on a specific date at a specific time:

1. On the **Veeam Updater** page, in the **Updates are available for this system** section, select check boxes next to the necessary updates.
2. In the **Choose action** section, do the following:
 - a. Select the **Schedule updates installation** option and configure the necessary schedule.

IMPORTANT

When selecting a date and time when updates must be installed, make sure no backup policies are scheduled to run on the selected time. Otherwise, the update process will interrupt the running activities, which may result in data loss.

- b. Select the **Reboot automatically after install if required** check box to allow Veeam Backup for AWS to reboot the backup appliance if needed.
- c. Click **Schedule Updates**.



Veeam Backup for AWS will automatically download and install the updates on the selected date at the selected time; the results of the installation process will be displayed on the **History** tab.

Setting Update Reminder

If you have not decided when to install updates, you can set an update reminder – instruct Veeam Backup for AWS to send an update notification later.

To do that, on the **Veeam Updater** page, in the **Choose action** section, do the following:

1. Select the **Remind me later** option and choose when you want to receive the reminder.

If you select the **Next Week** option, Veeam Backup for AWS will send the reminder in a week.

2. Click Remind me later.

The screenshot shows the Veeam Updater application interface. At the top, there is a dark green header with the Veeam logo, the text 'Veeam Updater', the date and time 'October 30, 2023 at 03:09 PM GMT+1', and a 'Configuration' button. Below the header, the main area is divided into two sections: 'Updates' and 'History'. The 'Updates' section is active and displays the following information:

- Updates are available for this system:** Last checked: 5 minutes ago. [Check for Updates...](#)
- Select packages to install:**
- Ubuntu security updates:** A list of 14 packages with checkboxes and file sizes:
 - Secure Sockets Layer toolkit - shared libraries (3.0.2-0ubuntu1.12) 1 MB
 - Complete Linux kernel for Amazon Web Services (AWS) systems. (6.2.0.1014.14~22.0... 1 KB
 - Linux kernel headers for Amazon Web Services (AWS) systems. (6.2.0.1014.14~22.04.1) 2 KB
 - Linux kernel image for Amazon Web Services (AWS) systems. (6.2.0.1014.14~22.04.1) 2 KB
 - Linux Kernel Headers for development (5.15.0-87.97) 1 MB
 - Secure Sockets Layer toolkit - cryptographic utility (3.0.2-0ubuntu1.12) 1 MB
 - Vi IMproved - Common files (2:8.2.3995-1ubuntu2.13) 79 KB
 - Vi IMproved - Runtime files (2:8.2.3995-1ubuntu2.13) 6 MB
 - Vi IMproved - enhanced vi editor - compact version (2:8.2.3995-1ubuntu2.13) 692 KB
 - Vi IMproved - enhanced vi editor (2:8.2.3995-1ubuntu2.13) 1 MB
 - tool to make (or reverse) a hex dump (2:8.2.3995-1ubuntu2.13) 53 KB

The 'Choose action:' section on the right shows:

- 0 packages selected
- Install updates now
- Schedule updates installation
- Remind me later
- A dropdown menu showing 'Next Week' with a downward arrow.
- A green button labeled 'Remind me later' with a mouse cursor hovering over it.

Updating IAM Roles

When you update the backup appliance to a newer version, the improvements and new features instantly become available in Veeam Backup for AWS. However, to meet new requirements, IAM roles must be assigned missing permissions manually either using the Veeam Backup for AWS UI or the AWS Management Console.

Updating Custom IAM Role

To update the custom IAM role, run a permission check for this role at the **IAM Roles** tab as described in section [Checking IAM Role Permissions](#). Veeam Backup for AWS will verify whether the IAM role is specified in any backup policy, repository or worker settings and check if all the permissions required to perform these operations are assigned to the role. If some of the permissions are missing, you will receive a warning in the **AWS Permission Check** window. You can grant the missing permissions to the IAM role using the AWS Management Console or [instruct Veeam Backup for AWS to do it](#). To learn how to grant permissions to IAM roles using the AWS Management Console, see [AWS Documentation](#).

NOTE

The permission check at the **IAM Roles** tab verifies only permissions of roles that are currently used by Veeam Backup for AWS. Permissions of IAM roles that are not specified in any settings on the backup appliance and are not used to perform any operations are not checked. That is why it is recommended that you additionally verify IAM role permissions using the built-in wizard permission check when specifying a role for the operation.

Updating Default Backup Restore IAM Role

You can update the *Default Backup Restore* IAM role [using the AWS Management Console](#) or instruct Veeam Backup for AWS to do it:

1. Click the warning.
2. In the **IAM Roles Update** window, provide one-time access keys of an IAM user that is authorized to update permissions of IAM roles, and then click **Apply**.

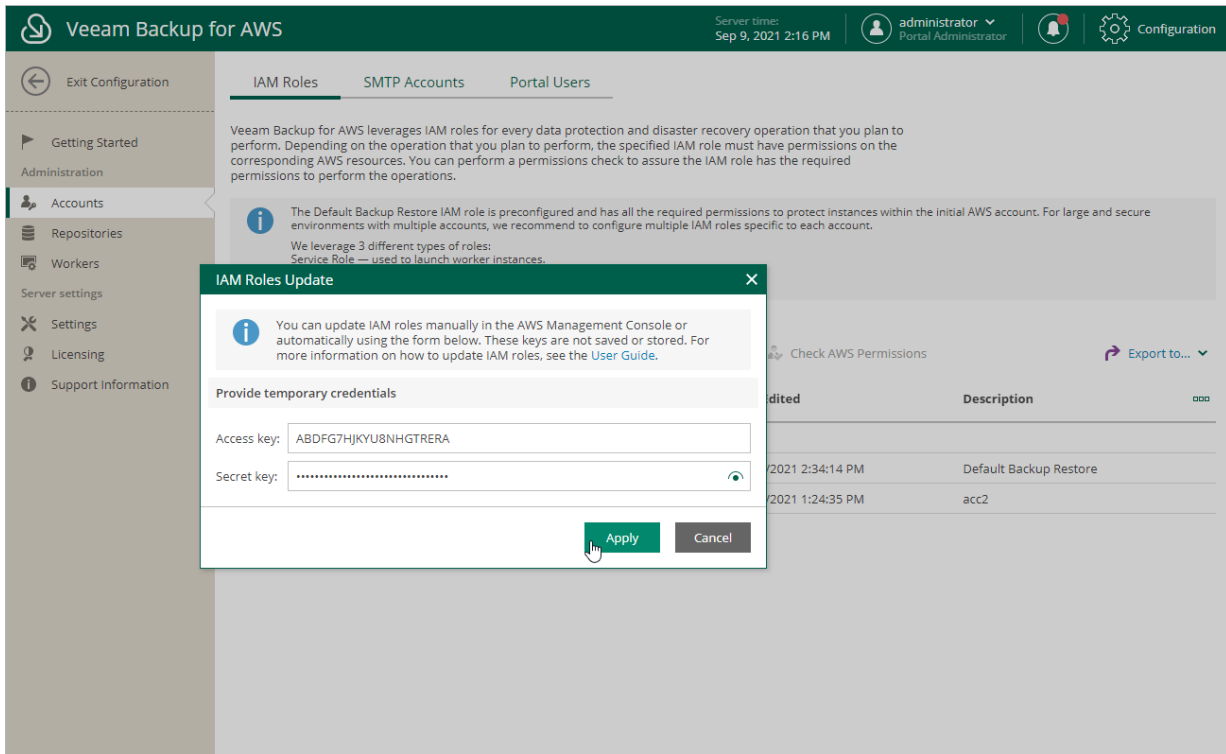
The IAM user must have the following permissions:

```
"iam:AttachRolePolicy",  
"iam:CreatePolicy",  
"iam:CreatePolicyVersion",  
"iam:CreateRole",  
"iam:GetAccountSummary",  
"iam:GetPolicy",  
"iam:GetPolicyVersion",  
"iam:GetRole",  
"iam:ListAttachedRolePolicies",  
"iam:ListPolicyVersions",  
"iam:SimulatePrincipalPolicy",  
"iam:UpdateAssumeRolePolicy"
```

NOTE

Veeam Backup for AWS does not store one-time access keys in the configuration database.

3. To make sure that the missing permissions have been successfully granted, navigate to **Accounts > IAM Roles**, select the *Default Backup Restore* IAM role and click **Check AWS Permissions**.



Viewing Updates History

To see the results of the update installation performed on the backup appliance, do the following:

1. Switch to the **Configuration** page.
2. Navigate to **Support Information > Updates**.
3. Click **Check and view updates**.
4. On the **Veeam Updater** page, switch to the **History** tab.

For each date when an update was installed, the **Veeam Updater** page will display the name of the update and its status (whether the installation process completed successfully or failed to complete).

To download logs for the installed updates, select the necessary date in the **Date** section, and click **View Full Log**. Veeam Backup for AWS will save the logs as a single file to the default download directory on the local machine.

The screenshot shows the Veeam Updater interface. At the top, there is a green header with the Veeam logo, the text 'Veeam Updater', the server time 'Jul 3, 2024 01:59 PM', and a 'Configuration' button. Below the header, there are two tabs: 'Updates' and 'History', with 'History' selected. The main content area is titled 'Update sessions history' and contains a table with two columns: 'Date' and 'Status'. The table lists 17 update sessions, all with a 'Success' status. The most recent session is highlighted in green. To the right of the table, there is a 'View Full Log' button. Below this button, there is a detailed log for the selected session, showing three messages, all with a 'Success' status.

Date ↓	Status
June 28, 2024 at 12:47 PM	Success
June 27, 2024 at 01:20 PM	Success
June 25, 2024 at 09:06 AM	Success
June 24, 2024 at 11:46 AM	Success
June 21, 2024 at 08:46 AM	Success
June 20, 2024 at 09:15 AM	Success
June 19, 2024 at 11:04 PM	Success
June 19, 2024 at 08:45 AM	Success
June 18, 2024 at 03:04 PM	Success
June 17, 2024 at 10:54 PM	Success
June 14, 2024 at 03:44 PM	Success
June 13, 2024 at 10:01 AM	Success
June 10, 2024 at 09:48 PM	Success
June 4, 2024 at 11:24 PM	Success
June 3, 2024 at 09:45 PM	Success
May 30, 2024 at 09:37 AM	Success

Message	Status
Successfully prepared the updates for installation	Success
Successfully installed package veeamawsbackup 8.0.0.845	Success
Successfully finalized the update process	Success

Configuring Web Proxy

To check for available package updates for Veeam Backup for AWS, the Veeam Updater service running on the backup appliance connects to the Veeam Update Repository over the internet. If the backup appliance is not connected to the internet, you can instruct the Veeam Backup for AWS to use a web proxy that will provide access to the required resources.

IMPORTANT


Veeam Backup for AWS does not support access to resources through HTTPS proxy.

To configure connection to the internet through a web proxy, do the following:

1. Open the **Veeam Updater** page:
 - a. Switch to the **Configuration** page.
 - b. Navigate to **Support Information**.
 - c. On the **Updates** tab, click **Check and View Updates**.
2. On the **Veeam Updater** page:
 - a. Switch to the **Configuration** page.
 - b. Navigate to **Proxy Server**.
 - c. Set the **Use HTTP Internet proxy** toggle to *On*.
 - d. In the **Host** field, enter the IP address or FQDN of the web proxy.
 - e. In the **Port** field, enter the port used on the web proxy for HTTP or HTTPS connections.
 - f. [Applies only if the web proxy requires authentication] In the **Username** and **Password** fields, enter credentials of the user account configured on the web proxy to access the internet.
 - g. Click **Save**.

IMPORTANT

You cannot modify the web proxy settings during checking for updates.



Getting Technical Support

If you have any questions or issues with Veeam Backup for AWS, you can search for a resolution on [Veeam R&D Forums](#) or submit a support case in the [Veeam Customer Support Portal](#).

When you submit a support case, it is recommended that you provide the Veeam Customer Support Team with the following information:

- [Version information for the product and its infrastructure components](#)
- The error message or an accurate description of the problem you are facing
- [Log files](#)

For information on Veeam Technical Support Tiers, SLAs and coverage, see the [Veeam Customer Support Policy](#).


Viewing Product Details Using Web UI

To view the product details:

1. Switch to the **Configuration** page.
2. Navigate to **Support Information**.

The **About** section of the **Updates** tab displays the following information:

- **Product version** – the currently installed version of Veeam Backup for AWS.
- **FLR service version** – the currently installed version of the File-level recovery service.
- **AWS ID** – the unique identification number of the AWS account where Veeam Backup for AWS is installed.
- **Support ID** – the unique identification number of the Veeam support contract.



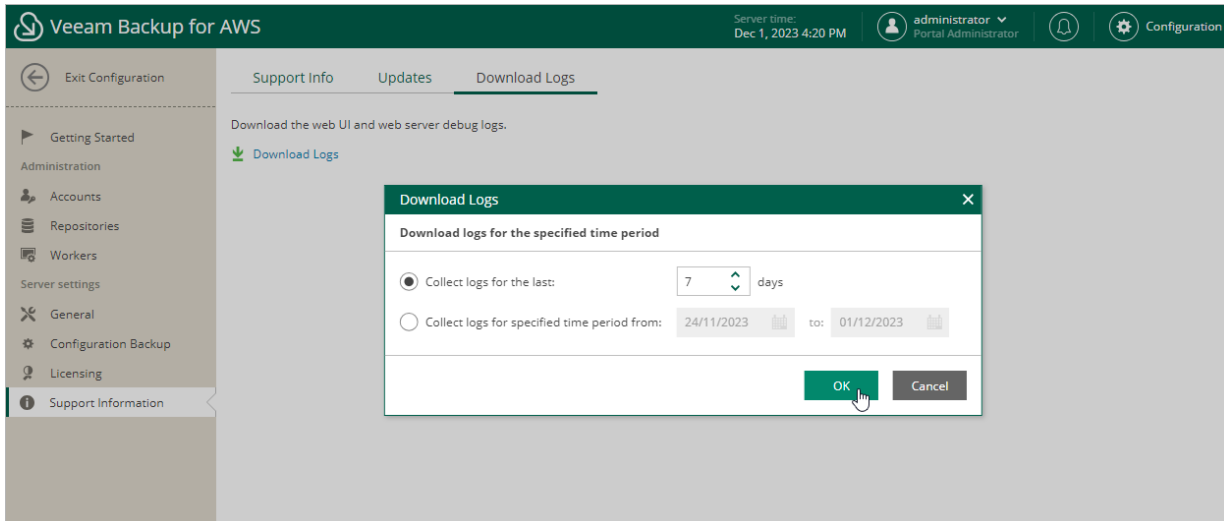
Downloading Product Logs Using Web UI

To download the product logs, do the following:

1. Switch to the **Download Logs** tab.
2. Click **Download Logs**.

3. In the **Download Logs** window, specify a time interval for which logs must be collected:
 - Select the **Collect logs for the last** option if you want to collect data for a specific number of days in the past.
 - Select the **Collect logs for specified time period** option if you want to collect data for a specific period of time in the past.
4. Click **OK**.

Veeam Backup for AWS will collect logs for the specified time interval and save them to the default download folder on the local machine in a single log.zip archive.



Downloading Product Logs Using Veeam Backup & Replication Console

To export the product logs, do the following:

1. In the Veeam Backup & Replication console, open the main menu and navigate to **Help > Support Information**.
2. In the **Export Logs** wizard, do the following:
 - a. At the **Scope** step, select the **Export all logs for selected components** option. Then, in the **Managed servers** list, select the backup server, backup appliances and other components for which you want to export logs.

b. Complete the wizard as described in the Veeam Backup & Replication User Guide, section [Export Logs](#).

Export Logs X

Scope
Specify the scope for logs export.

Scope | Date Range | Location | Export

Export logs for this job:
[Text Field] [Choose...]

Export logs for these objects:
[Text Field] [Choose...]

Export all logs for selected components (may result in a very large log package)
Managed servers:

Server ↑	Components
<input type="checkbox"/> 172.24.29.99	Installer, Tape Proxy, Transport
<input checked="" type="checkbox"/> srv12win16.tech.lo...	Installer, Mount Server, Transport, Veeam A...
<input type="checkbox"/> 172.24.182.46	Hyper-V Integration, Installer, Transport
<input checked="" type="checkbox"/> dept-01-amroz-sr...	AWS backup appliance

[Select All] [Clear All]

< Previous | **Next >** | Finish | Cancel

Appendices

This section provides additional information on how to configure AWS endpoints, AWS Identity and Access Management resources required for Veeam Backup for AWS to perform backup and restore operations.

Appendix A. Creating IAM Roles in AWS

NOTE

This section provides instructions on steps performed in a third-party application. Keep in mind that the instructions may become outdated. For up-to-date instructions, see [AWS Documentation](#).

You must specify an IAM role for each data protection and disaster recovery operation performed by Veeam Backup for AWS – the solution uses permissions of the specified IAM roles to access AWS services and resources. You can either [create an IAM role using Veeam Backup for AWS](#), or, first create the role in AWS using the AWS Management Console, [AWS CLI](#) or [AWS API](#), and then [add this role to Veeam Backup for AWS](#).

This section describes how to create an IAM role for Veeam Backup for AWS using the AWS Management Console. To do that:

1. Log in to the **AWS Management Console** using credentials of an AWS account in which you want to create the IAM role.
2. Navigate to **All Services > Security, Identity, & Compliance** and click **IAM**.
3. In the **IAM** console, navigate to **Access Management > Roles** and click **Create role**.
4. Complete the **Create role** wizard:
 - a. At the **Select trusted entity** step of the wizard, do either of the following:
 - If you want to create the IAM role in the initial AWS account to which the backup appliance belongs, select the **Custom trust policy** option. Then, in the **Custom trust policy** section, add the following statement to allow the [Impersonal IAM role](#) to assume the role you want to create:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "sts:AssumeRole",
      "Principal": {
        "AWS": "<Role ARN>"
      }
    }
  ]
}
```

Where `<Role ARN>` is the ARN of the *Impersonation* IAM role. To obtain the ARN of the *Impersonation* IAM role, you can look it up on the **Roles** page in the AWS Management Console.

- If you want to create the IAM role in another AWS account, select the **AWS account** option. Then, in the **An AWS account** section, select the **Another AWS account** option and enter the ID of the trusted account – the AWS account to which the backup appliance belongs.

If you want to increase the security of the role, select the **Require external ID** check box and enter a password. To learn how to use an external ID to increase security of an IAM role, see [AWS Documentation](#).

- b. At the **Add permissions** step of the wizard, select an IAM policy that must be attached to the IAM role.
For an IAM policy to be displayed in the list, it must be created beforehand as described in section [Appendix B. Creating IAM Policies in AWS](#).
 - c. At the **Role details** step of the wizard, specify a name and description for the IAM role.
 - d. At the **Tags** step of the wizard, specify AWS tags that will be assigned to the IAM role.
 - e. Click **Create role**.
5. Add the created IAM role to the Veeam Backup for AWS configuration database as described in section [Adding IAM Roles](#).

Appendix B. Creating IAM Policies in AWS

NOTE

This section provides instructions on steps performed in a third-party application. Keep in mind that the instructions may become outdated. For up-to-date instructions, see [AWS Documentation](#).

When you [create an IAM role](#), you must define permissions that the role will have in AWS. To define the role permissions, you must create an IAM policy and attach it to the IAM role. For more information on managing IAM identity permissions, see [AWS Documentation](#).

To create an IAM policy using the AWS Management Console, do the following:

1. Log in to the **AWS Management Console** using credentials of an AWS account in which you want to create the IAM policy.
2. Navigate to **All Services > Security, Identity, & Compliance** and click **IAM**.
3. In the **IAM** console, navigate to **Access Management > Policies** and click **Create policy**.
4. Complete the **Create policy** wizard:
 - a. At the **Editor** step of the wizard, switch to the **JSON** tab.
 - b. Type or paste a JSON policy document.

The JSON policy document must include permissions required for an IAM role to which you want to attach the policy. For more information on required permissions, see [IAM Permissions](#). To learn how to write JSON policy documents, see [AWS Documentation](#).

IMPORTANT

Consider the following AWS limitations on IAM policy sizing:

- The size of a managed IAM policy cannot exceed 6,144 characters. For more information on managed IAM policies, see [AWS Documentation](#).
- The total size of all inline IAM policies added to an IAM role cannot exceed 10,240 characters. For more information on inline IAM policies, see [AWS Documentation](#).

For more information on IAM character limits, see [AWS Documentation](#).

- c. At the **Tags** step of the wizard, specify AWS tags that will be assigned to the IAM policy.
- d. At the **Review** step of the wizard, specify a name and description for the IAM policy. Review the configured settings and click **Create policy**.

After you create a policy, you can attach it to IAM roles as described in section [Appendix A. Creating IAM Roles in AWS](#).

Appendix C. Configuring Endpoints in AWS

IMPORTANT

The provided instructions on configuring endpoints are not compatible with the [private network deployment](#) functionality. If you plan to use this functionality, follow the instructions provided in section [Configuring Private Networks](#).

If you want worker instances to operate in private environments, that is to use subnets with disabled auto-assignment of Public IPv4 addresses to deploy worker instances in AWS Regions, configure specific endpoints for services used by the backup appliance to perform backup and restore operations.

The following endpoints are required to perform Veeam Backup for AWS operations.

Operation	Interface Endpoints	S3 Gateway Endpoints
Creating EC2 image-level backups	<ul style="list-style-type: none"> com.amazonaws.<region>.ec2messages com.amazonaws.<region>.ssm com.amazonaws.<region>.sqs com.amazonaws.<region>.ebs 	<ul style="list-style-type: none"> com.amazonaws.<region>.s3
Restoring EC2 instances from image-level backups	<ul style="list-style-type: none"> com.amazonaws.<region>.ec2messages com.amazonaws.<region>.ssm com.amazonaws.<region>.sqs 	<ul style="list-style-type: none"> com.amazonaws.<region>.s3
Restoring EC2 volumes from image-level backups	<ul style="list-style-type: none"> com.amazonaws.<region>.ec2messages com.amazonaws.<region>.ssm com.amazonaws.<region>.sqs 	<ul style="list-style-type: none"> com.amazonaws.<region>.s3
Performing health check for EC2 backups	<ul style="list-style-type: none"> com.amazonaws.<region>.ec2messages com.amazonaws.<region>.ssm com.amazonaws.<region>.sqs 	<ul style="list-style-type: none"> com.amazonaws.<region>.s3
Creating EC2 archived backups	<ul style="list-style-type: none"> com.amazonaws.<region>.ec2messages com.amazonaws.<region>.ssm com.amazonaws.<region>.sqs 	<ul style="list-style-type: none"> com.amazonaws.<region>.s3
Creating RDS image-level backups	<ul style="list-style-type: none"> com.amazonaws.<region>.ssmmessages com.amazonaws.<region>.ssm com.amazonaws.<region>.sqs 	<ul style="list-style-type: none"> com.amazonaws.<region>.s3
Restoring PostgreSQL DB instances from image-level backups	<ul style="list-style-type: none"> com.amazonaws.<region>.ssmmessages com.amazonaws.<region>.ssm com.amazonaws.<region>.sqs 	<ul style="list-style-type: none"> com.amazonaws.<region>.s3
Performing health check for RDS backups	<ul style="list-style-type: none"> com.amazonaws.<region>.ssmmessages com.amazonaws.<region>.ssm com.amazonaws.<region>.sqs 	<ul style="list-style-type: none"> com.amazonaws.<region>.s3
Creating RDS archived backups	<ul style="list-style-type: none"> com.amazonaws.<region>.ssmmessages com.amazonaws.<region>.ssm com.amazonaws.<region>.sqs 	<ul style="list-style-type: none"> com.amazonaws.<region>.s3
Applying retention policy settings to created restore points	<ul style="list-style-type: none"> com.amazonaws.<region>.ec2messages com.amazonaws.<region>.ssm com.amazonaws.<region>.sqs 	<ul style="list-style-type: none"> com.amazonaws.<region>.s3

Operation	Interface Endpoints	S3 Gateway Endpoints
Performing file-level recovery from image-level backups	<ul style="list-style-type: none"> com.amazonaws.<region>.ec2messages com.amazonaws.<region>.ssm com.amazonaws.<region>.sqs [Applies only if you restore to the original location] com.amazonaws.<region>.kinesis-streams 	<ul style="list-style-type: none"> com.amazonaws.<region>.s3
Performing file-level recovery from cloud-native snapshots and replicated snapshots	<ul style="list-style-type: none"> com.amazonaws.<region>.ec2messages com.amazonaws.<region>.ssm com.amazonaws.<region>.sqs [Applies only if you restore to the original location] com.amazonaws.<region>.kinesis-streams 	<ul style="list-style-type: none"> com.amazonaws.<region>.s3
Performing EFS indexing	<ul style="list-style-type: none"> com.amazonaws.<region>.ssmmessages com.amazonaws.<region>.ssm com.amazonaws.<region>.sqs 	<ul style="list-style-type: none"> com.amazonaws.<region>.s3

To create these endpoints, use the specified endpoint names, where <region> is the name of an AWS Region in which worker instances will be deployed.

Creating Interface Endpoints

NOTE

This section provides instructions on steps performed in a third-party application. Keep in mind that the instructions may become outdated. For up-to-date instructions, see [AWS Documentation](#).

To allow Veeam Backup for AWS to create EC2 and RDS image-level backups and to perform restore operations and EFS indexing, configure interface VPC endpoints in AWS regions where worker instances are deployed for subnets to which worker instances must be connected. By default, Veeam Backup for AWS uses the default or the most appropriate network settings of AWS Regions to deploy worker instances. However, you can add specific worker configurations as described in section [Configuring Private Networks](#).

For more information on AWS regions in which worker instances are deployed to perform specific operations, see [Worker Instances in Private Environment](#).

To create an interface VPC endpoint, do the following:

1. Log in to the **AWS Management Console** using credentials of an AWS account in which you want to create the endpoint.
2. In the **AWS services** section, navigate to **All Services > Networking & Content Delivery** and click **VPC**. The **VPC** console will open.
3. Navigate to **Virtual Private Cloud > Endpoints** and click **Create Endpoint**. The **Create endpoint** wizard will open.
4. Complete the **Create endpoint** wizard:
 - a. At the **Endpoint settings** step of the wizard, do the following:

- i. [Optional] In the **Name tag** field, specify a name for the endpoint.
 - ii. In the **Service category** section, select **AWS services**.
- b. At the **Services** step of the wizard, use the following filter *Type: Interface* and select a service for which you want to create a VPC endpoint.
- c. At the **VPC** step of the wizard, do the following:
- i. From the **VPC drop-down** list, select a VPC to which the deployed worker instances will be connected.
 - ii. In the **Additional settings** section, select the **Enable DNS name** check box.
- d. At the **Subnets** step of the wizard, select one subnet for each Availability Zone where worker instances will be deployed.
- e. At the **Security groups** step of the wizard, select security groups that will be associated with the endpoint network interfaces.
- Ensure that the security group that is associated with the endpoint network interface allows communication between the endpoint network interface and the resources in your VPC that communicate with the service. If the security group restricts inbound HTTPS traffic (port 443) from resources in the VPC, you will not be able to send traffic through the endpoint network interface.
- f. At the **Policy** step of the wizard, select **Full access** to allow full access to the service. Alternatively, select **Custom** and attach a VPC endpoint policy that will control permissions on resources available over the VPC endpoint.
- g. Click **Create Endpoint**.

For more information on interface VPC endpoints, see [AWS Documentation](#).

Creating S3 Gateway Endpoints

NOTE

This section provides instructions on steps performed in a third-party application. Keep in mind that the instructions may become outdated. For up-to-date instructions, see [AWS Documentation](#).

To allow Veeam Backup for AWS to create image-level backups of EC2 instances, to perform restore operations from these backups, and to save EFS indexes to backup repositories, configure S3 gateway endpoints in AWS regions where worker instances are deployed for subnets to which worker instances must be connected. By default, Veeam Backup for AWS uses the default or the most appropriate network settings of AWS Regions to deploy worker instances. However, you can add specific worker configurations as described in section [Managing Worker Configurations](#).

For more information on AWS regions in which worker instances are deployed to perform specific operations, see [Worker Instance Locations](#).

To create a gateway endpoint for a subnet, do the following:

1. Log in to the **AWS Management Console** using credentials of an AWS account in which you want to create the endpoint.
2. In the **AWS services** section, navigate to **All Services > Networking & Content Delivery** and click **VPC**. The **VPC** console will open.
3. Navigate to **Virtual Private Cloud > Endpoints** and click **Create Endpoint**. The **Create endpoint** wizard will open.

4. Complete the **Create endpoint** wizard:
 - a. At the **Endpoint settings** step of the wizard, do the following:
 - i. [Optional] In the **Name tag** field, specify a name for the endpoint.
 - ii. In the **Service category** section, select **AWS services**.
 - b. At the **Services** step of the wizard, use the following filter *Type: Gateway* and select `com.amazonaws.<region>.s3`, where `<region>` is a name of an AWS Region in which worker instances will be deployed.
 - c. At the **VPC** step of the wizard, select a VPC to which the deployed worker instances will be connected.
 - d. At the **Route tables** step of the wizard, select the route tables to be used by the endpoint. AWS automatically will add a route that points traffic destined for the service to the endpoint network interface.
 - e. At the **Policy** step of the wizard, select **Full access** to allow full access to the service. Alternatively, select **Custom** and attach a VPC endpoint policy that will control permissions on resources available over the endpoint.
 - f. Click **Create Endpoint**.

For more information on gateway endpoints for Amazon S3, see [AWS Documentation](#).

IMPORTANT

When you create an S3 gateway endpoint, consider that a VPC and a service for which you create the endpoint must belong to the same AWS Region. That is, when you perform backup operations using endpoints, the processed source instances must reside in the region in which a repository where the backups will be stored is located; when you perform restore operations using endpoints, the instances must be restored to the region in which a repository where the backup files are stored is located.

This limitation is only region-specific-services and VPCs can belong to different AWS accounts.

Appendix D. Enabling Swap Partition

NOTE

This section provides instructions on steps performed in a third-party application. Keep in mind that the instructions may become outdated. For up-to-date instructions, see [AWS Documentation](#).

By enabling a swap partition on the EC2 instance where Veeam Backup for AWS is installed, you can prevent runtime issues on the backup appliance. The swap partition allows the system to allocate additional resources when the backup appliance runs out of physically allocated memory due to overload caused by multiple concurrent operations.

To enable the swap partition on the backup appliance, you must first create and attach an additional EBS volume to the EC2 instance running Veeam Backup for AWS, and then perform a number of configuration actions on the instance.

NOTE

When you deploy Veeam Backup for AWS on the EC2 instance of the *C5d* instance type, a number of [instance store volumes](#) is automatically attached to the instance, and the store volumes are partitioned with swap spaces, one of which equals to the amount of RAM allocated to the instance. However, if instance store volumes have already been manually configured before the product installation, the swap space formatting will not be created automatically, and you will have to create it manually as described in [AWS Documentation](#).

Creating and Attaching EBS Volume in AWS

To create a new EBS volume and attach it to the backup appliance, do the following:

1. Log in to the **AWS Management Console** using credentials of an AWS account where the backup appliance resides.
2. Navigate to **All Services > Compute** and click **EC2**.
3. In the **EC2** console, navigate to **Volumes** and click **Create Volume**.
4. Complete the **Create volume** wizard:
 - a. At the **Volume settings** section of the wizard, do the following:
 - i. From the **Volume** type drop-down list, select General Purpose SSD (gp3). For more information on EBS volume types, see [AWS Documentation](#).
 - ii. In the **Size (GB)** field, specify the size of the volume. For swap partition purposes, it is recommended that you create an EBS volume with a minimum size equal to the memory size of the EC2 instance.

For more information on EC2 instance memory sizes, see [AWS Documentation](#).
 - iii. In the **IOPS** field, specify the maximum number of input/output operations per second that the volume must provide. For swap partition purposes, 4000 IOPS is recommended.
 - iv. In the **Throughput** field, specify the throughput that the volume must provide. It is recommended that you specify the maximum throughput available for the selected volume size.

- v. From the **Availability Zone** drop-down list, select the availability zone in which the backup appliance resides.
- vi. From the **Snapshot ID** drop-down list, select the **Don't create volume from a snapshot** option.
- vii. If you want to encrypt the EBS volume, select the **Encrypt the volume** check box. You can either select a [default KMS key](#) from the **KMS key** drop-down list, which is automatically created by Amazon EBS in the specified AWS Region, or specify the amazon resource number (ARN) of the key in the **Specify a custom KMS key** window.

IMPORTANT

If you choose to encrypt the EBS volume, make sure that the EC2 instance type of the backup appliance supports Amazon EBS encryption. For more information, see [AWS Documentation](#).

For more information on KMS keys, see [AWS Documentation](#).

- b. At the **Tags** section of the wizard, you can specify AWS tags that will be assigned to the volume.
 - c. Click **Create volume**.
5. To attach the created EBS volume to the EC2 instance, select the volume from the **Volumes** list and click **Actions > Attach volume**.
 6. Complete the **Attach volume** wizard:
 - a. From the **Instance** drop-down list, select the EC2 instance running Veeam Backup for AWS.

NOTE

The backup appliance and the EBS volume that you want to attach must reside in the same availability zone.

- b. In the **Device name** field, specify a name for the volume that will be used by Amazon EC2. Note that the name must conform the available device name rules, and it will be changed later by the block device driver when mounting the volume. For more information, see [AWS Documentation](#).
- c. Click **Attach volume**.

Configuring Swap Partition on EC2 Instance

After you have attached the created volume to the backup appliance, you must perform a number of configuration actions to enable a swap partition:

1. Connect to the EC2 instance where Veeam Backup for AWS is installed. To do that, run the following `ssh` command in a terminal window:

```
ssh -i /path/EC2_instance.pem key ubuntu@<Public DNS hostname or IPv4 address of the EC2 instance>
```

2. To get a list of available volumes, run the following command:

```
sudo lsblk
```

You can identify the newly added volume by the absence of the mount point. Save the volume name for future reference.

3. To create a swap file system on the new volume, run the following command:

```
sudo mkswap /dev/<volume_name> -L "vbaws_swap"
```

4. To add the newly created file system to the */etc/fstab* file, do the following:

- a. Open the file:

```
sudo nano /etc/fstab.conf
```

- b. Add the following file system label:

```
LABEL=vbaws_swap swap swap defaults,nofail 0 0
```

- c. Save the changes.

5. To enable the swap partition, run the following command:

```
sudo swapon -all
```

6. To confirm that the swap partition is enabled, run the following command:

```
sudo swapon
```

7. To allow Veeam Backup for AWS to use swap space preference, do the following:

- a. Open the file:

```
sudo nano /etc/sysctl.d/99-sysctl.conf
```

- b. Add the following variable to the file and set its value to 1:

```
vm.swappiness = 1
```

- c. Save the changes.

8. Reload the */etc/sysctl.d/99-sysctl.conf* file to apply the changes without rebooting EC2 instance:

```
sudo sysctl -p /etc/sysctl.d/99-sysctl.conf
```

Appendix E. Uninstalling Backup Appliances Deployed from AWS Marketplace

Starting from version 8.0, you can deploy Veeam Backup for AWS from the Veeam Backup & Replication console only. However, if an appliance was previously deployed from the AWS Marketplace or is running Veeam Backup for AWS version 3.x (or earlier), use one of the following options to uninstall the solution:

- If you deployed a backup appliance from AWS Marketplace, you must [delete the CloudFormation stack](#) created while installing Veeam Backup for AWS. All resources included in the stack will be deleted automatically.
- If you deployed a backup appliance from the AMI, you must [manually delete AWS resources](#) created while installing Veeam Backup for AWS.

IMPORTANT

When you deploy Veeam Backup for AWS from the Veeam Backup & Replication console, the CloudFormation stack is not created and AWS resources cannot be managed as a single unit. Keep in mind that these resources are not automatically deleted from AWS when you remove the backup appliance from Veeam Backup & Replication. To learn how to manually delete resources created during Veeam Backup for AWS installation, see [Removing Appliances](#).

Note that backed-up data will not be removed automatically after you uninstall the solution. You can keep this data in your AWS environment and import it to a new backup appliance:

- To import cloud-native snapshots, rescan AWS Regions where the snapshots are stored. The snapshots will be automatically imported to the configuration database.
- To import image-level backups, assign the Amazon S3 bucket where the backups are stored to a new backup repository as described in section [Adding Backup Repositories Using Console](#).

If you do not want to keep the backed-up data, remove it manually as described in section [Managing Backed-Up Data](#). Alternatively, you can remove the data using the AWS Management Console:

1. Log in to the **AWS Management Console** using credentials of an AWS account where the data is stored.
2. Use the region selector in the upper-right corner of the page to select the AWS Region in which the backed-up data is stored.
3. Remove the backed-up data:
 - To remove backups, navigate to **Services > S3**. Select an Amazon S3 bucket where the backups are stored. Navigate to **Veeam > Backup**, select the backup repository folder, and click **Delete**.
 - To remove RDS cloud-native snapshots, navigate to **Services > RDS > Snapshots**, select the necessary Veeam snapshots, and click **Delete**.
 - To remove EC2 cloud-native snapshots, navigate to **Services > EC2 > Snapshots**, select the necessary Veeam snapshots, and click **Delete**.
 - To remove EFS cloud-native backups, navigate to **Services > EFS > Backups**, select the necessary Veeam backups, and click **Delete**.

- To remove FSx cloud-native backups, navigate to **Services > FSx > Backups**, select the necessary Veeam backups, and click **Delete**.
- To remove Redshift cloud-native backups, navigate to **Services > Redshift > Backups**, select the necessary Veeam backups, and click **Delete**.

Deleting CloudFormation Stack

When you deploy a backup appliance from AWS Marketplace, Veeam Backup for AWS is installed using an AWS CloudFormation stack. In AWS CloudFormation, a stack is a collection of AWS services and resources that you can manage as a single unit. To uninstall Veeam Backup for AWS, you must delete the CloudFormation stack from AWS. For more information on working with stacks, see [AWS Documentation](#).

To delete the Veeam Backup for AWS CloudFormation stack, perform the following steps:

1. Log in to the **AWS Management Console** using credentials of an AWS account where Veeam Backup for AWS is installed.
2. Use the region selector in the upper-right corner of the page to select the AWS Region in which the backup appliance resides.
3. Navigate to **Services > CloudFormation**.
4. From the **Stacks** list, select the CloudFormation stack created while installing Veeam Backup for AWS.
5. Click **Delete**.
6. In the confirmation window, click **Delete stack** to acknowledge deletion.

NOTE

Consider the following:

- After you acknowledge the operation, the Veeam Backup for AWS CloudFormation stack will acquire the *DELETE_IN_PROGRESS* state. When all AWS resources included in the stack are successfully deleted, the stack will acquire the *DELETE_COMPLETE* state. By default, deleted CloudFormation stacks are not displayed in the AWS Management Console. To learn how to view deleted stacks and to troubleshoot deletion issues, see [AWS Documentation](#).
- If a backup appliance managed by the Veeam Backup & Replication server has been upgraded from the Veeam Backup & Replication console, you will encounter the issue while deleting the CloudFormation stack of this appliance – you will not be able to delete it on the first try. To work around the issue, retry deleting stuck and choose the **Force delete this entire stack** option.

Deleting AWS Resources

When you deploy a backup appliance from the Amazon Machine Image (AMI), Veeam Backup for AWS creates a number of resources while operating in AWS, and these resources are not removed from infrastructure automatically when you delete the backup appliance. To uninstall Veeam Backup for AWS, you must locate and delete the following resources from your infrastructure:

- AWS::IAM::InstanceProfile
- AWS::DLM::LifecyclePolicy
- AWS::CloudWatch::Alarm
- AWS::EC2::SecurityGroup
- AWS::IAM::Role
- AWS::EC2::Instance

To delete a resource, do the following:

1. Log in to the **AWS Management Console** using credentials of an AWS account where Veeam Backup for AWS is installed.
2. Use the region selector in the upper-right corner of the page to select the AWS Region in which the backup appliance resides.
3. Navigate to AWS service to which the AWS resource belong.
4. Select the AWS resource that you want to remove, and click **Delete**.

Appendix F. Installing and Removing License

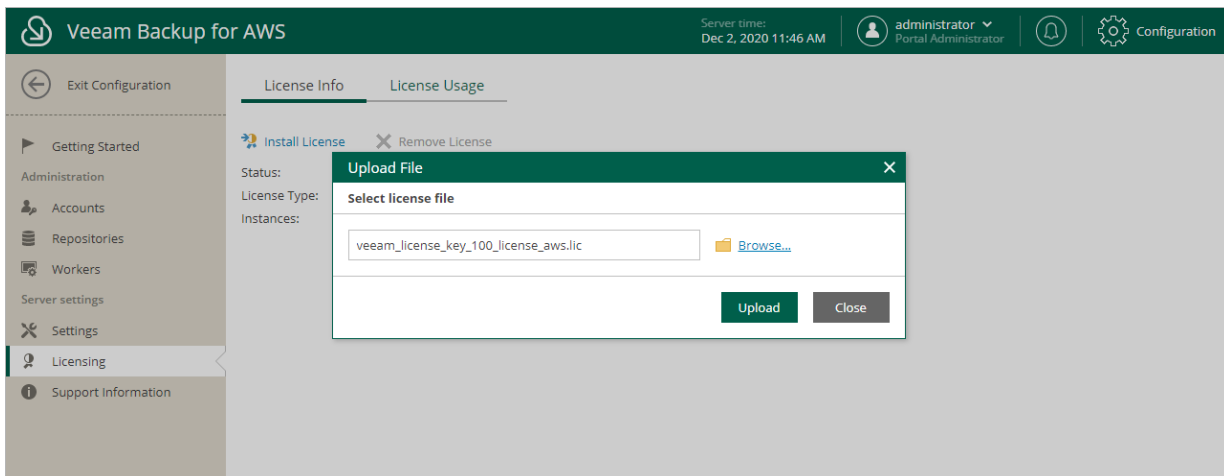
NOTE

This section applies only to the *BYOL* edition of Veeam Backup for AWS.

Installing License

To install or update a license installed on the backup appliance, do the following:

1. Switch to the **Configuration** page.
2. Navigate to **Licensing > License Info**.
3. Click **Install License**.
4. In the **Upload file** window, click **Browse** to browse to a license file, and then click **Upload**.

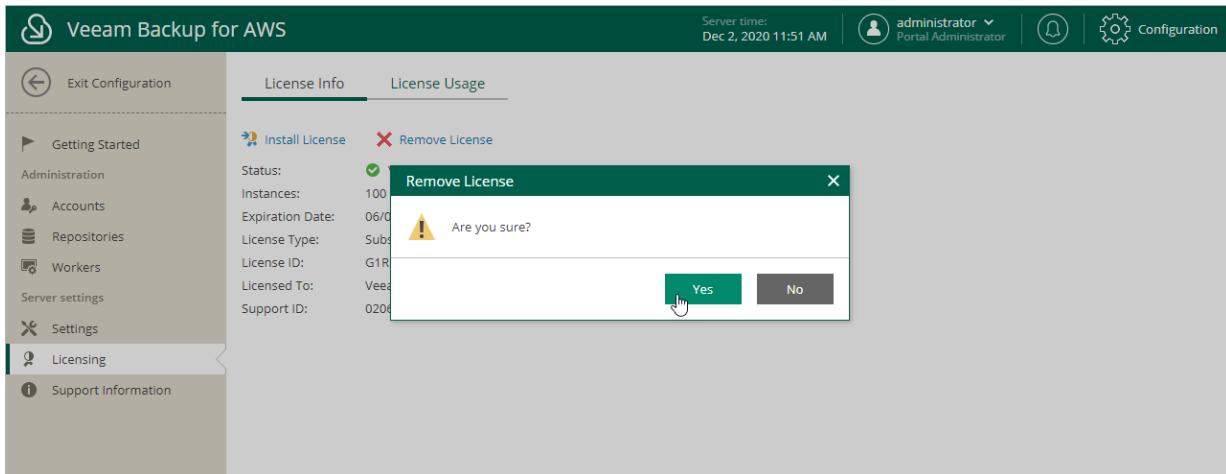


Removing License

To remove a license installed on the backup appliance if you no longer need it:

1. On the **License Info** tab, click **Remove License**.

2. In the **Remove License** window, click **Yes** to confirm that you want to remove the license.



After you remove the license, Veeam Backup for AWS will automatically switch back to the *Free* edition. In this case, according to the FIFO (first-in first-out) queue, only the first 10 instances registered in the configuration database will remain protected. You can revoke license units from these instances as described in section [Revoking License Units](#).