



Veeam Backup for AWS

Version 7.0

User Guide

May, 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	8
ABOUT THIS DOCUMENT	9
OVERVIEW	10
Integration with Veeam Backup & Replication	12
Solution Architecture	13
Backup Server	14
AWS Plug -In for Veeam Backup & Replication	15
Backup Appliances	16
Backup Repositories	18
Worker Instances	19
Additional Repositories and Tape Devices	24
Gateway Servers	25
Protecting EC2 Instances	26
EC2 Backup	27
EC2 Restore	36
Protecting RDS Resources	39
RDS Backup	40
RDS Restore	44
Protecting DynamoDB Tables	46
DynamoDB Backup	48
DynamoDB Restore	50
Protecting EFS File Systems	51
EFS Backup	52
EFS Restore	55
Protecting VPC Configurations	56
VPC Configuration Backup	57
Exporting VPC Configuration	59
VPC Configuration Restore	60
Retention Policies	62
Immutability	63
Private Network Deployment	65
Backup Appliances in Private Environment	66
Worker Instances in Private Environment	73
Data Encryption	97
Backup Repository Encryption	98
AWS KMS Encryption	99
PLANNING AND PREPARATION	113

System Requirements	114
Ports.....	116
AWS Services.....	118
Plug-In Permissions	120
IAM Permissions	133
Service IAM Permissions	134
Repository IAM Permissions.....	162
Backup IAM Permissions	164
Restore IAM Permissions	186
Full List of IAM Permissions	206
IAM Permissions Changelog	218
Considerations and Limitations	220
Sizing and Scalability Guidelines	224
Backup Appliance.....	225
Object Storage.....	230
Backup Policies	231
Worker Instances	233
DEPLOYMENT	234
Deploying Plug-In.....	235
Installing Plug-In	236
Installing and Uninstalling Plug-In in Unattended Mode.....	237
Upgrading Plug-In	239
Uninstalling Plug-In	240
Deploying Backup Appliance	241
Deploying Appliance from Console	242
Deploying Appliance from AWS Marketplace	253
LICENSING	284
Licensing of Managed Backup Appliances	285
Licensing of Standalone Backup Appliances	287
Installing and Removing License	288
Viewing License Information	290
Revoking License Units	293
ACCESSING VEEAM BACKUP FOR AWS	295
Accessing Web UI from Console	296
Accessing Web UI from Workstation	297
CONFIGURING VEEAM BACKUP FOR AWS.....	300
Managing Backup Appliances	301
Adding Appliances	302
Editing Appliance Settings.....	315
Rescanning Appliances	317

Removing Appliances	318
Managing IAM Roles	320
Adding IAM Roles	321
Editing IAM Role Settings	334
Checking IAM Role Permissions	336
Removing IAM Roles	340
Managing User Accounts	341
Adding User Accounts	343
Editing User Account Settings	344
Changing User Passwords	345
Configuring Multi-Factor Authentication	346
Managing Backup Repositories	348
Adding Backup Repositories Using Console	349
Adding Backup Repositories Using Web UI	362
Editing Backup Repository Settings	375
Rescanning Backup Repositories	378
Removing Backup Repositories	379
Managing Worker Instances	381
Managing Worker Configurations	382
Managing Worker Profiles	391
Adding Worker Tags	395
Configuring General Settings	396
Enabling Private Network Deployment	397
Configuring Global Retention Settings	399
Configuring Global Notification Settings	401
Replacing Security Certificates	405
Changing Time Zone	407
Configuring SSO Settings	408
Performing Configuration Backup and Restore	410
Performing Configuration Backup	411
Performing Configuration Restore	416
VIEWING AVAILABLE RESOURCES	434
Adding Resources to Policy	436
PERFORMING BACKUP	437
Performing Backup Using Console	439
Creating Backup Policies	440
Editing Backup Policy Settings	441
Enabling and Disabling Backup Policies	442
Starting and Stopping Backup Policies	443
Deleting Backup Policies	444

Creating Backup Copy Jobs	445
Copying Backups to Tapes	446
Performing Backup Using Web UI	447
Performing EC2 Backup	448
Performing RDS Backup	491
Performing DynamoDB Backup	528
Performing EFS Backup	559
Performing VPC Configuration Backup	588
Managing EC2, RDS, DynamoDB and EFS Backup Policies	599
MANAGING BACKED-UP DATA.....	604
Managing Backed-Up Data Using Console	605
Managing Backed-Up Data Using Web UI	608
EC2 Data	609
RDS Data	615
DynamoDB Data	618
EFS Data.....	620
VPC Configuration Data	622
PERFORMING RESTORE.....	641
EC2 Restore.....	642
EC2 Restore Using Console	643
EC2 Restore Using Web UI	662
RDS Restore	702
RDS Restore Using Console	703
RDS Restore Using Web UI	730
DynamoDB Restore.....	759
DynamoDB Restore Using Console	760
DynamoDB Restore Using Web UI	761
EFS Restore	773
EFS Restore Using Console	774
EFS Restore Using Web UI	786
VPC Configuration Restore	812
Performing VPC Configuration Restore Using Console	813
VPC Configuration Restore Using Web UI	814
Instant Recovery.....	830
Exporting Disks	832
Publishing Disks	833
Restoring to Microsoft Azure	834
Restoring to Google Cloud	836
Restoring to Nutanix AHV	837
REVIEWING DASHBOARD.....	839

VIEWING SESSION STATISTICS	841
COLLECTING OBJECT PROPERTIES.....	843
UPDATING VEEAM BACKUP FOR AWS	844
Updating Appliances Using Console	845
Upgrading to 7.0 from Version 6.0 or Earlier	847
Updating Appliances Using Web UI	849
Upgrading Appliances	850
Checking for Updates	851
Installing Updates	852
Updating IAM Roles	856
Viewing Updates History	858
Configuring Web Proxy.....	859
GETTING TECHNICAL SUPPORT.....	860
APPENDICES.....	863
Appendix A. Creating IAM Roles in AWS	864
Appendix B. Creating IAM Policies in AWS	865
Appendix C. Configuring Endpoints in AWS	866
Appendix D. Enabling Swap Partition	872

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 Amazon Elastic Compute Cloud (EC2), Amazon Relational Database Service (RDS), Amazon DynamoDB and Amazon Elastic File System (EFS) environments. Veeam Backup for AWS also allows you to back up and restore Amazon Virtual Private Cloud (VPC) configurations.

IMPORTANT

Veeam Backup for AWS is available only in AWS Global and AWS GovCloud (US) regions.

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

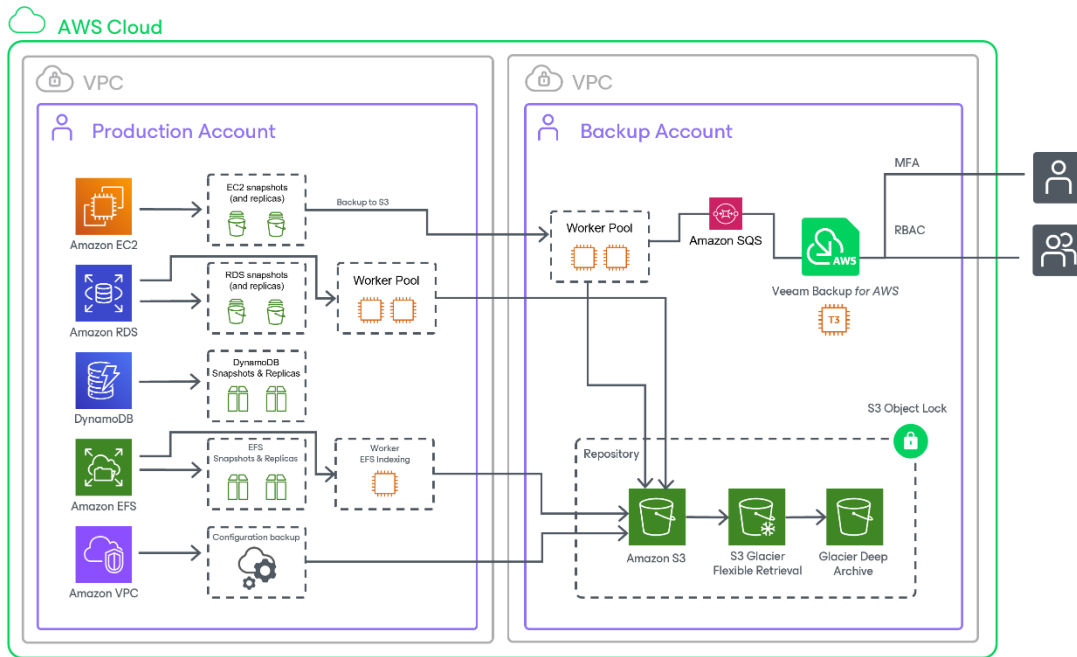
To recover backed-up data, you can perform the following operations:

- 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, PostgreSQL DB instances and Aurora DB clusters.
- Restore DynamoDB tables.
- Restore entire EFS 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

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.

Version 7.0 comes with 2 major features – the ability to create image-level backups of PostgreSQL DB instances and to back up DynamoDB tables – that are available only for those backup appliances managed by a Veeam Backup & Replication server. To unlock the full functionality, you must [install AWS Plug-in for Veeam Backup & Replication on the server](#) and [add your appliances](#) to the backup infrastructure.

IMPORTANT

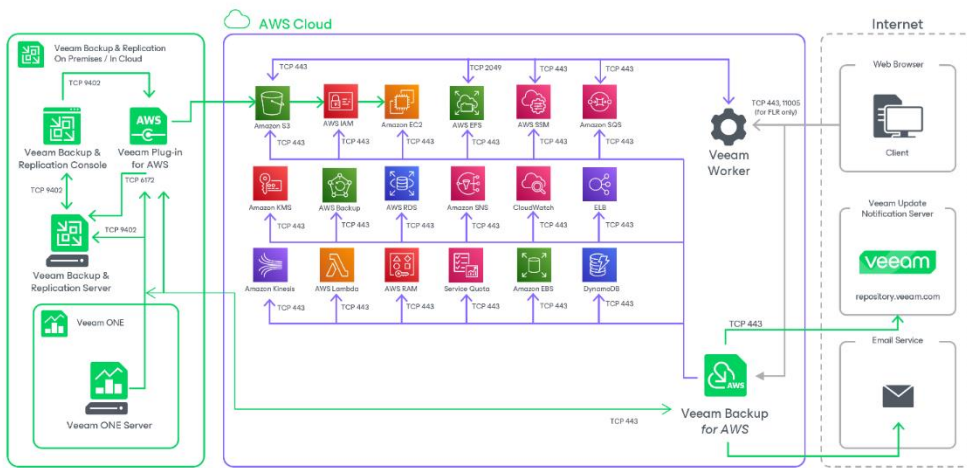
Consider the following:

- 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 RDS resources, 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 backup policies. Creating DynamoDB 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 RDS resources and DynamoDB tables.

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 6.0
- 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.
- **Web UI** – provides a web interface that allows user to access to the Veeam Backup for AWS functionality.
- **Updater service** – allows Veeam Backup for AWS to check, view and install product and package updates.
- **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 details, 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 Veeam Backup for AWS configuration backups.

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, the 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 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, AWS services and other Veeam Backup for AWS components.

Worker Instance Components

A worker instance uses the following components:

- **Veeam Data Mover** – the service that performs data processing tasks. During backup, the Veeam Data Mover retrieves data of an AWS protected resource (EC2, RDS). During restore, the 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 launched for file-level recovery.

Security Certificates for Worker Instances

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 during file-level recovery. A self-signed certificate is generated automatically on the worker instance when the restore session starts.

How Worker Instances Work

Veeam Backup for AWS automatically launches a worker instance in Amazon EC2 for the duration of a backup, restore or retention process and removes it immediately after the process is complete. Veeam Backup for AWS launches one worker instance per each AWS resource specified in a backup policy, restore or retention task.

Veeam Backup for AWS can launch worker instances in the following AWS accounts:

- The backup account is an AWS account to which the service IAM role specified to launch worker instances belongs. By default, Veeam Backup for AWS uses this account to launch worker instances for backup, restore and backup retention operations.
- Production accounts are the same AWS accounts where the processed resources belong. By default, Veeam Backup for AWS uses these accounts to launch worker instances for EFS indexing and for RDS backup and restore operations.

To minimize cross-region traffic charges, depending on the data protection and disaster recovery operation, Veeam Backup for AWS launches the worker instance in the following location:

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

Operation	Worker Instance Location	Possibility to Deploy Worker Instances in Production Accounts	Default Worker Instance Type
Restoring EC2 instances from image-level backups	AWS Region to which an EC2 instance is restored	Yes	<ul style="list-style-type: none"> 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 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
Creating RDS image-level backups	AWS Region and VPC in which a processed PostgreSQL DB instance resides	Yes	<ul style="list-style-type: none"> c5.large – if the total EBS volume size is less than 1024 GB

Operation	Worker Instance Location	Possibility to Deploy Worker Instances in Production Accounts	Default Worker Instance Type
Restoring PostgreSQL DB instances from image-level backups	AWS Region to which a DB instance is restored	Yes	<ul style="list-style-type: none"> c5.2xlarge – if the total storage size is less than 6 TB c5.4xlarge – if the total storage size is more than 6 TB
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	Yes	<ul style="list-style-type: none"> t3.medium
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

NOTE

For RDS image-level backup operations, performing EFS indexing, and restoring PostgreSQL DB instances from image-level backups, you can instruct Veeam Backup for AWS to deploy worker instances in production accounts only.

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

Required Ports

The following network ports must be open to ensure proper communication of components 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.
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:

- [Amazon Elastic Compute Cloud \(EC2\)](#)
- [AWS Systems Manager \(SSM\)](#), including access to the *ec2messages* and *ssmmessages* endpoints
- [Amazon Simple Queue Service \(SQS\)](#)
- [AWS Security Token Service \(STS\)](#)
- [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.](#)

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

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, 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](#).

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 launch 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. Launches 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 launch 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 [Enabling Data Encryption](#).

- d. When the backup session completes, removes the worker instance from Amazon EC2.

5. If you enable the [backup archiving mechanism](#), Veeam Backup for AWS performs the following operations:

- a. Launches 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. When the archive session completes, removes the worker instance from Amazon EC2.

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 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 details, 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. For details on how to remove them, 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](#).

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 total size of backups that must be deleted is more than 50 GB, launches 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 launch 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 launched, 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 launch 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](#).

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, launches 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 launch 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 launched, 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 launch 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 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. Launches 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. Launches 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. Launches 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 launches 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 launches 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

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. You can also instruct Veeam Backup for AWS to create image-level backups of the processed PostgreSQL DB instances. For more information on how RDS resources backup works, see [RDS Backup](#).

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 launch 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.

IMPORTANT

Snapshot replication is not supported for Aurora multi-master clusters.

3. If you enable image-level backup for the backup policy, Veeam Backup for AWS performs the following operations:
 - a. Launches a worker instance in an AWS Region in which the processed DB instance resides in an AWS account where 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. When the backup session completes, removes the worker instance from Amazon EC2.
5. If you enable the [backup archiving mechanism](#), Veeam Backup for AWS performs the following operations:
 - a. Launches a worker instance in an AWS Region where a backup repository storing backed-up data resides in AWS account to which the service IAM role used to launch worker instances belongs – that is, the backup account.
 - b. Retrieves data from the backup repository and transfers it to the archive backup repository.
 - c. When the archive session completes, removes the worker instance from Amazon EC2.

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 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 details, 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.

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. For details on 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 details, see [RDS Backup Retention](#).

Archive Backup Chain

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

1. During the first archive session, Veeam Backup for AWS detects backed-up data that is stored in the backups existing in the regular backup chain, creates a full archive backup with all the data, and copies this backup to the archive repository. The full archive backup becomes a starting point in the archive chain.
2. During subsequent archive sessions, Veeam Backup for AWS continues to create full archive backups.

The period of time during which RDS backups are kept in the archive backup chain is defined by retention policy settings. 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

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

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, snapshot replica or an AWS snapshot.
- [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, snapshot replica or an AWS snapshot, 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. Launches a worker instance in an AWS Region in which DB instance that will host the restored databases resides in an AWS account where 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 from Amazon RDS.

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

Protecting DynamoDB Tables

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 version 7.0 allows you to back up and restore the following table properties:

Property	Description	Possibility to Change Property During Restore to New Location
Table name	The name of a DynamoDB table.	Yes
Partition key	The first attribute of the primary key.	No
Sort key	The 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
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

Property	Description	Possibility to Change Property During Restore to New Location
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)	An 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:

- DynamoDB global table feature
- CloudWatch alarms
- 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
- Backup and restore of keys in your tables and indexes identified by Amazon CloudWatch Contributor Insights

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 details, 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. For details on 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 [Performing DynamoDB Backup](#).

Protecting EFS File Systems

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

EFS indexing allows you 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 launch workers while performing indexing of the processed EFS file systems](#).
5. [\[Optional\] Configure global retention settings for obsolete session records](#).
6. [\[Optional\] Configure email notification settings for automated delivery of backup policy results and daily reports](#).
7. [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. Launches a worker instance in an AWS Region in which the processed file system resides in an AWS account where the file system belong – 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

Veeam Backup for AWS encrypts and compresses data saved to backup repositories. For more information on data encryption, see [Enabling 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 details, 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 instance 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 Recovery](#).

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 details, 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. For details on 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 [Performing Entire File System 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 [Performing File-Level Recovery](#).

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 Veeam Backup for AWS database.

To back up VPC configurations of AWS Regions added to a backup policy, Veeam Backup for AWS uses permissions of an IAM role specified in the backup policy settings. The VPC configuration data is collected for the selected AWS Regions in the AWS account to which the specified IAM role belongs.

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 configure the VPC Configuration Backup policy to copy backup files to a backup repository, Veeam Backup for AWS launches the Veeam Data Mover service on the backup appliance to copy the restore point to the target backup repository specified in the backup policy settings. In the repository, for each AWS account in which VPC configuration data has been backed up, Veeam Backup for AWS creates an individual folder with VPC configuration backup files.

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](#).

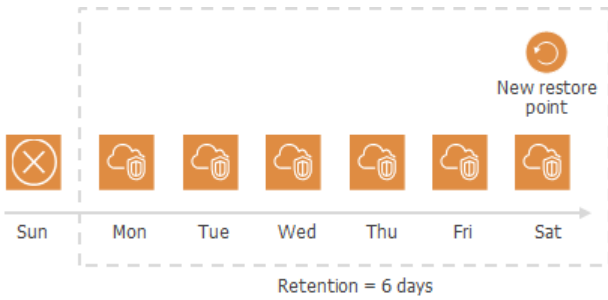
VPC Configuration Backup Retention

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 details, 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 details, see sections [EC2 Backup Retention](#), [RDS Backup Retention](#), [DynamoDB Backup Retention](#), [EFS 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 details, see [Configuring Global Retention Settings](#).

Immutability

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 10 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 10 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 10 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 $<10 - N>$ 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 11th day a new block generation period is initiated. Veeam Backup for AWS creates a new incremental backup and adds 10 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 10 days to the retention period specified in the backup policy settings. Thus, the retention period of the backup will be prolonged to 15 days, and the immutability expiration date will become March 16.
2. On March 2, Veeam Backup for AWS will create a new incremental backup at 7:00 AM and add 9 days to the retention period specified in the backup policy settings. Thus, the retention period of the incremental backup will be prolonged to 14 days, and the retention date will become March 16.
3. On March 3-10, Veeam Backup for AWS will continue creating incremental backups and extending their retention period so that the retention date will still remain March 16.
4. On March 11, 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 15 days, and the immutability expiration date will become March 26.

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](#)

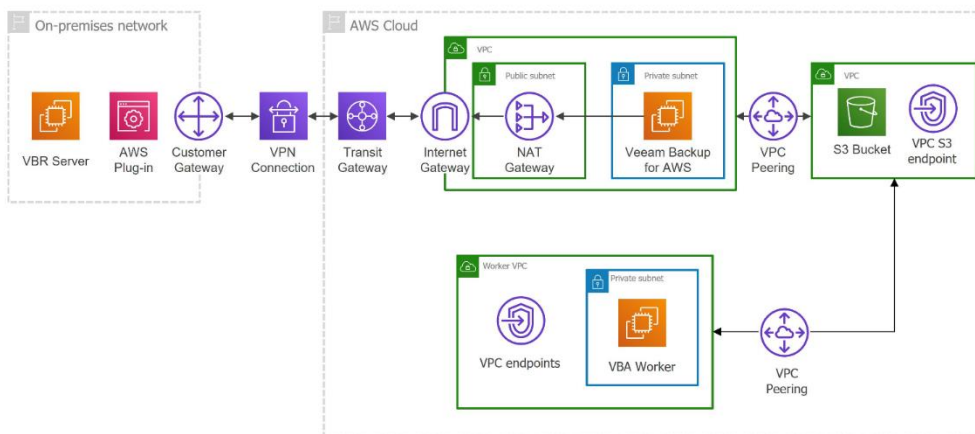
Backup Appliances in Private Environment

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 [using a CloudFormation template](#), you have an option to connect it either to an existing or to a new private VPC:

- If you choose to connect the appliance to a new private VPC, the VPC and two subnets (public and private) will be automatically created in the AWS Region in which the appliance resides; also, an internet gateway will be attached to the VPC to allow the appliance to access the internet. The appliance will be connected to the private subnet and will access the required [AWS services](#) through a route to a NAT gateway that will be created in the public subnet.
- If you choose to connect the appliance to an existing VPC, you will have to manually configure access both to the AWS services and the internet in the way that suites your security concerns best.

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](#).



In both cases, 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 Notification Server (repository.veeam.com)	HTTPS	443

From	To	Protocol	Port
	Ubuntu Security Update Repository (security.ubuntu.com)	HTTP	80
	DotNetCore Repository (packages.microsoft.com)	HTTPS	443
	PostgreSQL Apt Repository (apt.postgresql.org)	HTTP	80
	PostgreSQL Website* (postgresql.org)	HTTPS	443

*Required to download the file <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 launches 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 launches 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 launches 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 launch 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 [Enabling 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

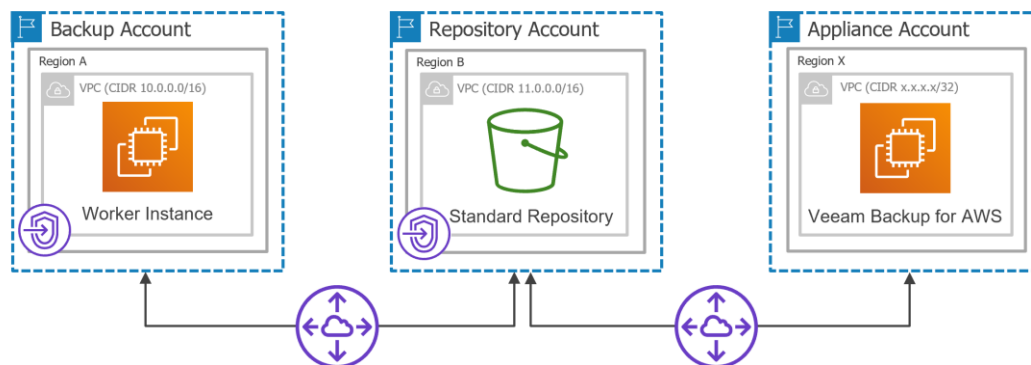
S3 gateway endpoints are not supported when using the private network deployment functionality.

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 (that is, the AWS account to which the service IAM role used to launch worker instances belongs), 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 [Managing Worker Configurations](#).



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>
Repository VPC	

Destination	Target
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 (that is, the AWS account to which the service IAM role used to launch worker instances belongs), 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>
Archive repository VPC (13.0.0.0/16)	<i>pcx-vvvv</i>
Archive Repository VPC	
Archive repository VPC (13.0.0.0/16)	<i>Local</i>

Destination	Target
Worker instance VPC (12.0.0.0/16)	<i>pcx-vvvv</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-kkkk</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-vvvv</i>
Standard repository VPC (11.0.0.0/16)	<i>pcx-kkkk</i>

4. 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

Security Groups Associated with S3 Interface Endpoint	From	Protocol	Port	Notes
	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>.ec2messages com.amazonaws.<region>.ssm com.amazonaws.<region>.sqs com.amazonaws.<region>.ebs 	<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
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
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>.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 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
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

Operation	Worker Instance Location	Possibility to Deploy Worker Instances in Production Accounts	Interface Endpoints	S3 Interface Endpoints
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>.ec2 messages 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>.ec2 messages 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	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>.ec2 messages 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 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 • com.amazonaws.<region>.sts 	<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 launched.

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

To launch 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 launched, 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 launched, 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**.

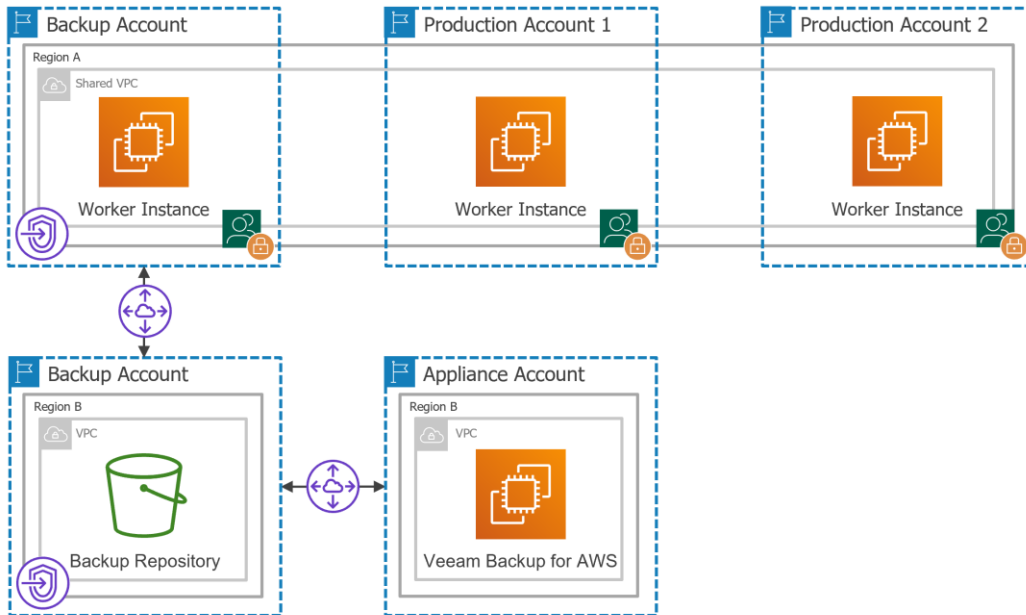
Configuring Private Networks for Production Accounts

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

To launch 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 launched, 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 launched, 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 file systems and DynamoDB tables.

For data encryption, Veeam Backup for AWS uses the 256-bit Advanced Encryption Standard (AES). For more information about AES, see [Advanced Encryption Standard \(AES\)](#).

NOTE

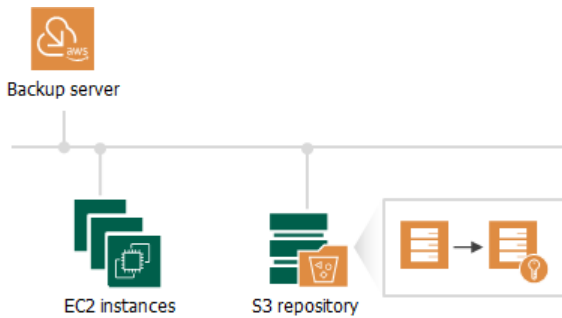
Sensitive customer data (credentials of user accounts required to connect to virtual servers and other systems, cloud credentials, and so on) is stored in the configuration database in the encrypted format.

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 file systems and DynamoDB tables 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 DynamoDB tables to a new 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:

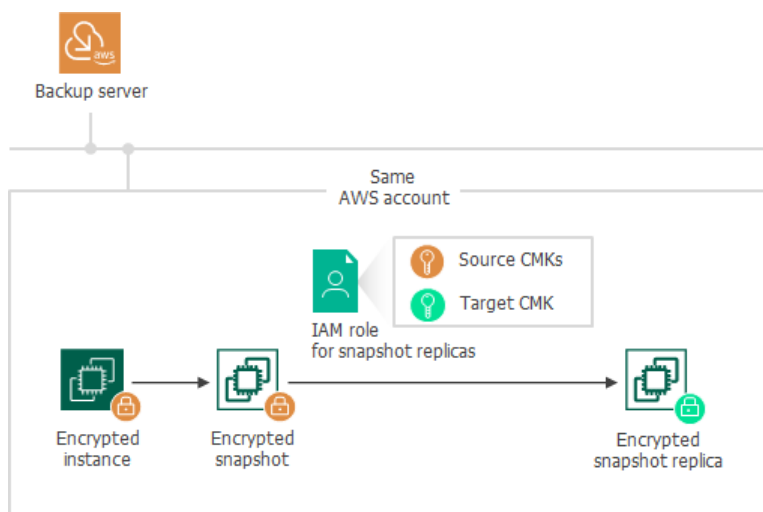
1. Takes an encrypted cloud-native snapshot of the instance.
2. Copies the created snapshot to the target AWS Region.

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 \(aws/eks alias\)](#), 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](#).

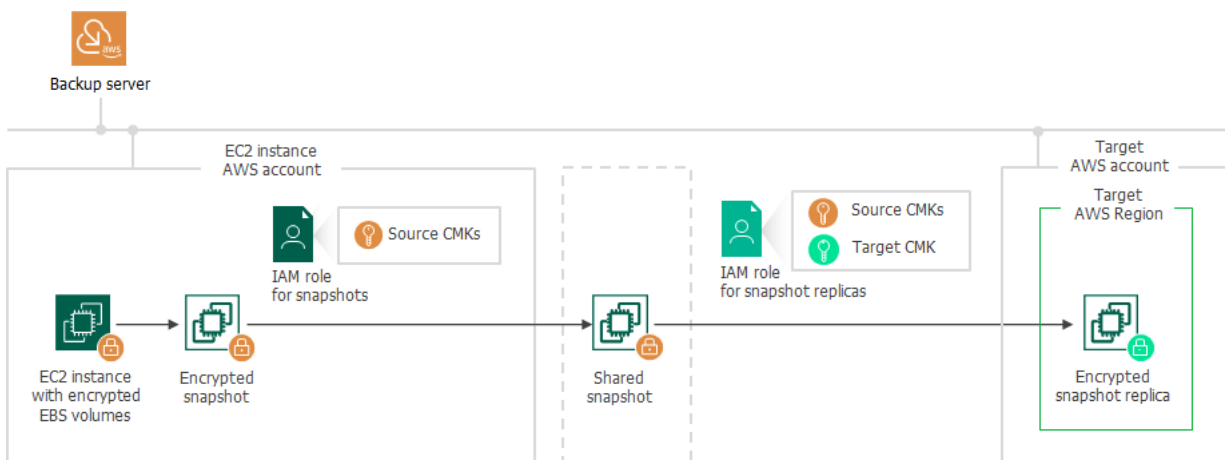
3. Copies the shared snapshot to the target AWS Region in the target AWS account.

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](#).

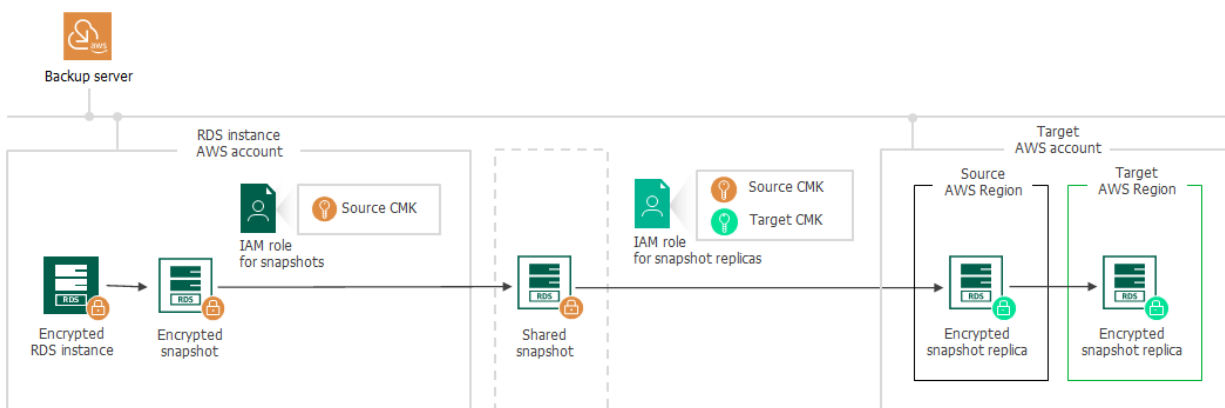
3. In the target AWS account, copies the shared encrypted snapshot to the same AWS Region to which the RDS instance belongs in the source AWS account. Then, if the target AWS Region differs from the source AWS Region, copies the shared snapshot to the target AWS Region.

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.](#)

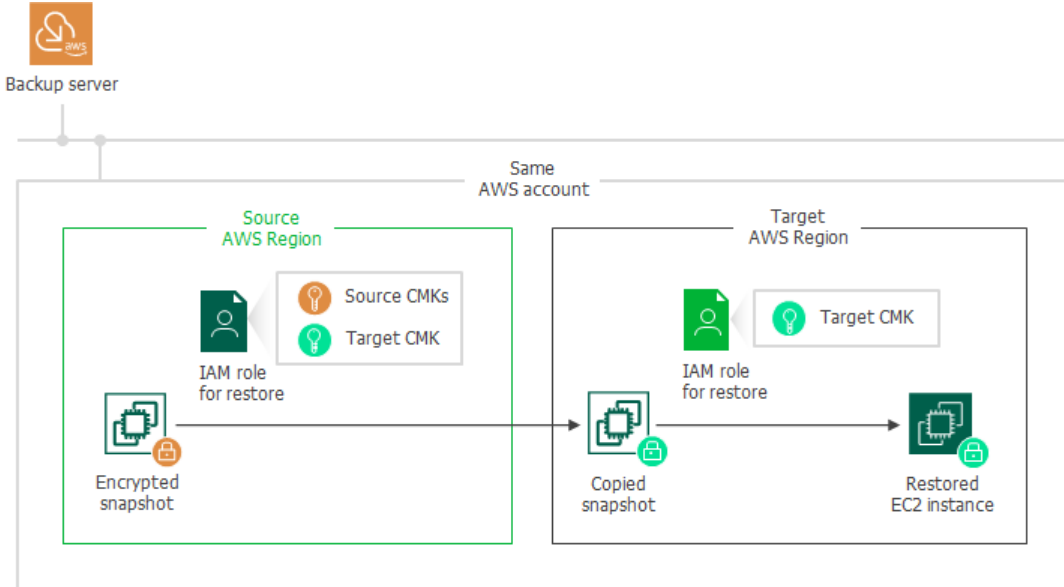
Restoring EC2 instance to Another AWS Region in Same AWS Account

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

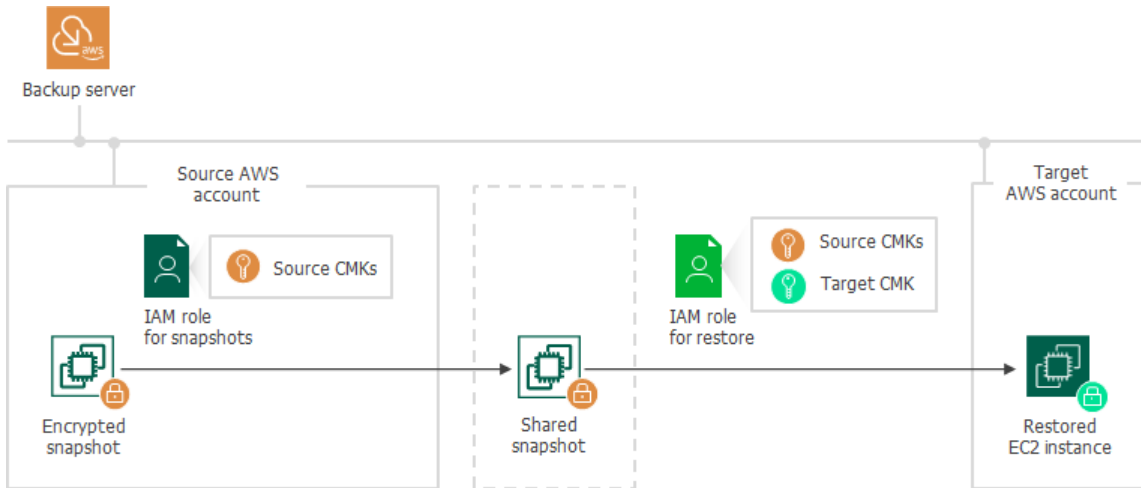
Due to AWS limitations, cloud-native snapshots encrypted with the [default key for EBS encryption \(aws/ebs alias\)](#) cannot be shared with other AWS accounts. Thus, if the cloud-native snapshot is encrypted with the default key for EBS encryption, Veeam Backup for AWS will not be able to share the 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

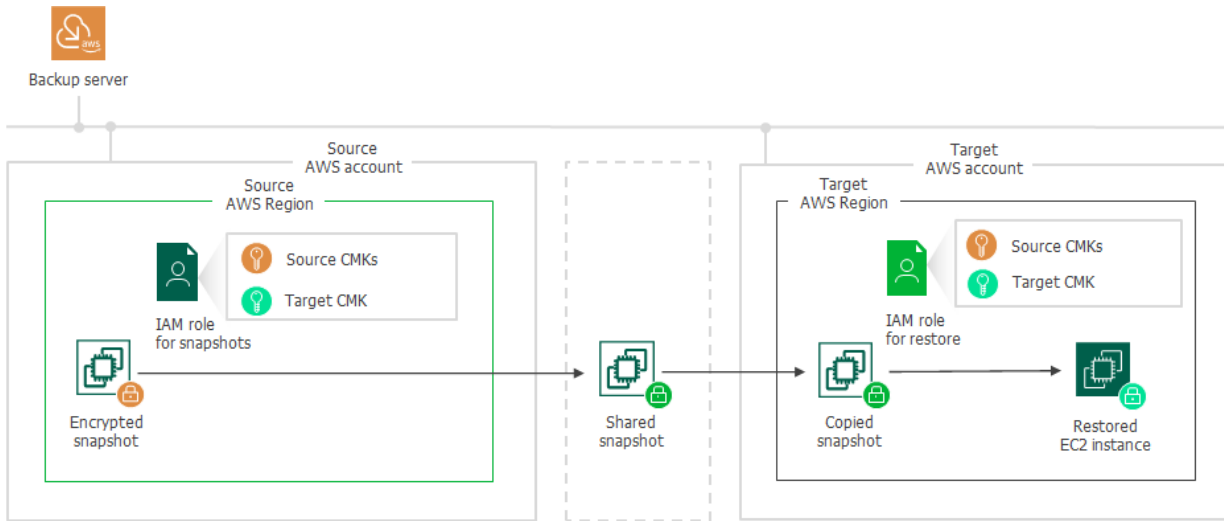
Due to AWS limitations, cloud-native snapshots encrypted with the [default key for EBS encryption \(aws/ebs alias\)](#) cannot be shared with other AWS accounts. Thus, if the cloud-native snapshot is encrypted with the default key for EBS encryption, Veeam Backup for AWS will not be able to share the 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).



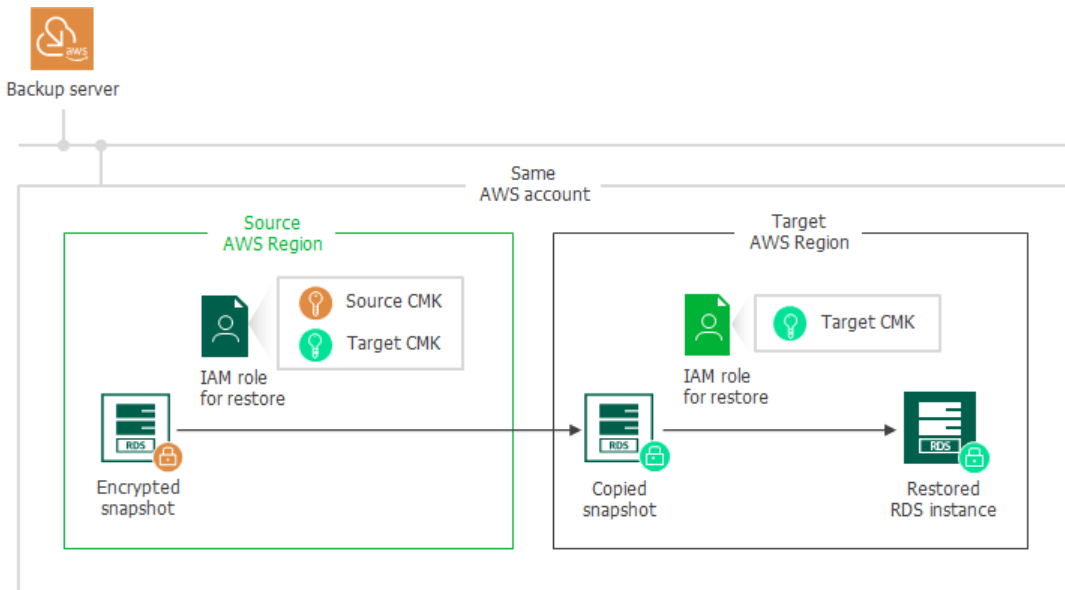
Restoring RDS Instance to Another AWS Region in Same AWS Account

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](#).

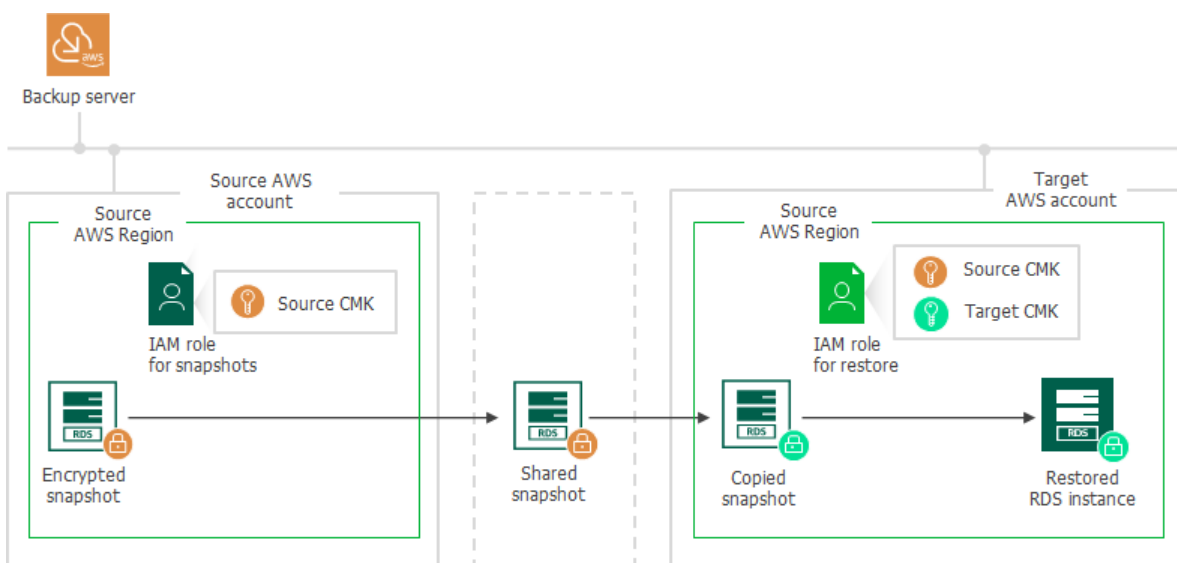
2. In the target AWS account, copies the shared snapshot to the same AWS Region where the snapshot resides in the source AWS account, and re-encrypts the snapshot with the KMS keys that you specified to encrypt the restored RDS instance.

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).

3. 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:

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 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

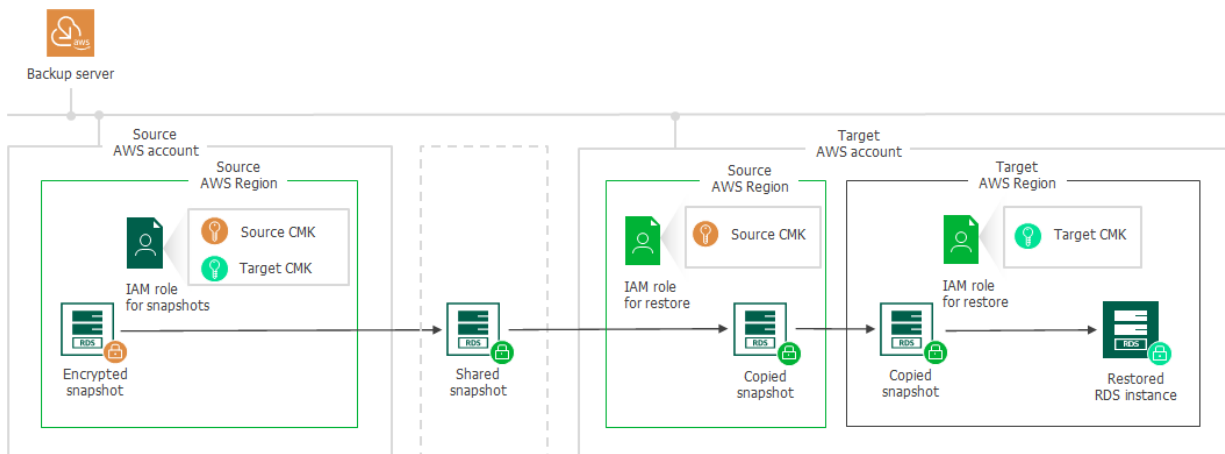
Due to AWS limitations, cloud-native snapshots encrypted with the [default encryption key \(aws/rds alias\)](#) cannot be shared with other AWS accounts. Thus, if the cloud-native snapshot is encrypted with the default encryption key, Veeam Backup for AWS will not be able to share the snapshot and the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

2. 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).

3. 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.
5. 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 launched in the same AWS account or not:

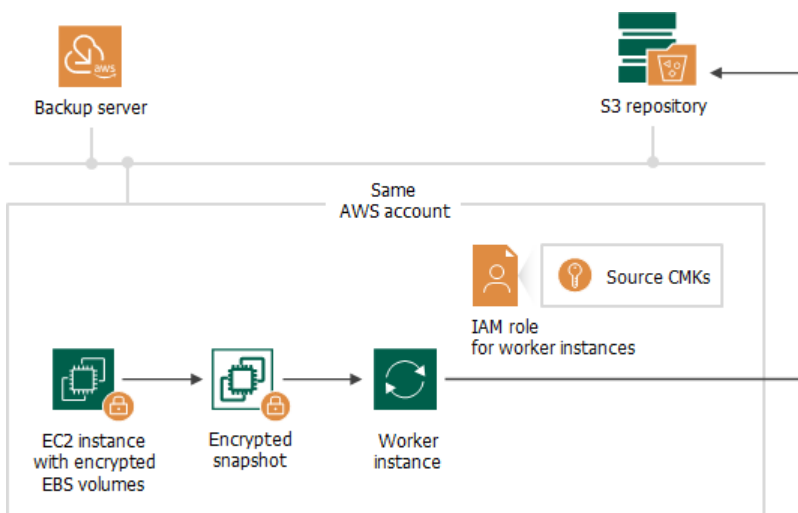
- [Creating the image-level backup in the same AWS account where the worker instance is launched.](#)
- [Creating the image-level backup in an AWS account other than the AWS account where the worker instance is launched.](#)

Creating Image-Level Backup in Same AWS Account

If a worker instance is launched in the same 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. 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 launch worker instances, as described in section [Configuring Worker Instance Settings](#). 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 launched 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 launched.

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

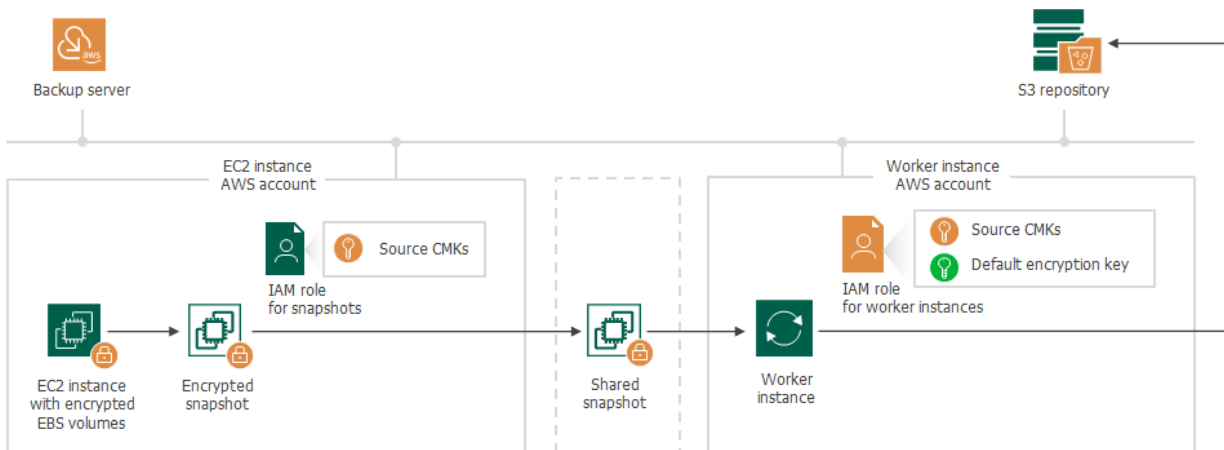
If EBS volumes of the EC2 instance are encrypted with the [default key for EBS encryption \(aws/eks alias\)](#), Veeam Backup for AWS will not be able to share the snapshot with another AWS account and the backup process will fail to complete successfully. To work around the issue, enable the worker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

To access the data, Veeam Backup for AWS uses an IAM role specified to launch worker instances, as described in section [Configuring Worker Instance Settings](#). 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 launched.



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 launched 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 launched.](#)
- [Performing restore from the image-level backup to an AWS account other than the AWS account where the worker instance is launched.](#)

NOTE

Consider the following:

- An AWS account to which an IAM role specified for launching 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 launches worker instances in a target AWS Region specified in the restore settings.

Restore to Same AWS Account

If a worker instance is launched 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 launch 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 launched 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 launch worker instances, as described in section [Configuring Worker Instance Settings](#). 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

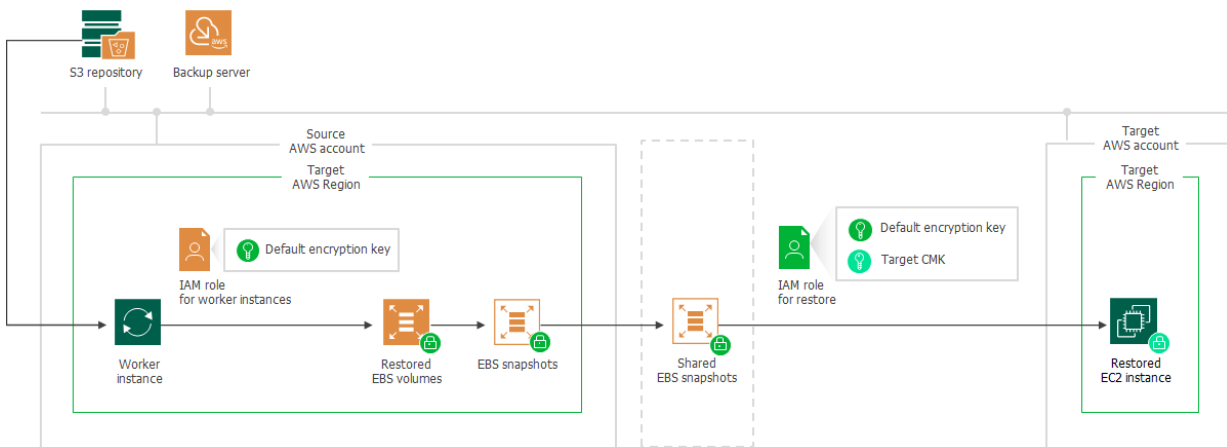
Due to AWS limitations, cloud-native snapshots encrypted with the [default key for EBS encryption \(aws/ebs alias\)](#) cannot be shared with other AWS accounts. Thus, if the default encryption key specified for the target AWS Region in the source AWS account is the default key for EBS encryption, Veeam 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 6.0.24
- Microsoft ASP.NET Core Shared Framework 6.0.24

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 12.7.0 supports integration with Veeam Backup & Replication version 12.1.

Veeam Backup for AWS

AWS Plug-in for Veeam Backup & Replication version 12.7.0 supports integration with Veeam Backup for AWS version 7.x.

Supported Applications

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

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

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
12.1.0.2131	12.7.0.1255	7.0.0, 7.0.1	Ubuntu 22.04 LTS
12.0.0.1420	12.2.6.5	6.1.0, 6.1.1, 6.1.2	
	12.1.6.93		
12.0.0.1420	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	
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/HTTP	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/HTTP	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/HTTP	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/HTTP	443	Required to download information on available product updates.
	Ubuntu Security Update Repository (security.ubuntu.com)	TCP/HTTP	80	Required to get OS security updates.
	DotNetCore Repository (packages.microsoft.com)	TCP/HTTP	443	Required to get DotNet 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/HTTPS	443	Required to download the file https://www.postgresql.org/media/keys/ACCC4CF8.asc .
	AWS services	TCP/HTTPS	443	Required to perform data protection and disaster recovery operations.
Worker instances	AWS services	TCP/HTTPS	443	Required to perform data protection and disaster recovery operations.
AWS Plug-in for Veeam Backup & Replication	Backup appliance, AWS services	TCP/HTTPS	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.

NOTE

When you deploy a backup appliance from the Veeam Backup & Replication console, Veeam Backup & Replication automatically creates security groups for the required ports to allow communication between the backup server and the appliance components.

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 [Installing Veeam Backup for AWS Using CloudFormation Template](#) and [Installing Veeam Backup for AWS from AMI](#).
- Security groups associated with worker instances. For more information, see [Configuring Worker Instance Settings](#).

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 AWS services.

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

AWS Plug-in for Veeam Backup & Replication must have outbound internet access to the following AWS services:

- [Amazon Elastic Compute Cloud \(EC2\)](#)
- [Amazon Simple Storage Service \(S3\)](#)
- [AWS Identity and Access Management \(IAM\)](#)

AWS Services Required For Backup Appliance

The backup appliance must have outbound internet access to the following AWS services:

- [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 Elastic File System \(EFS\)](#)
- [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 \(ELB\)](#)

- [Amazon DynamoDB](#)

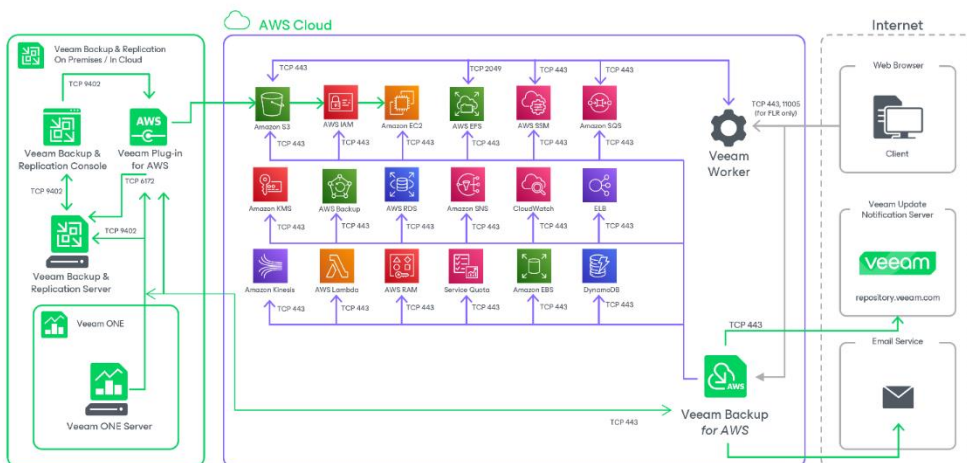
AWS Services Required For Worker Instances

Worker instances must have outbound internet access to the following AWS services:

- [Amazon Elastic Compute Cloud \(EC2\)](#)
- [AWS Systems Manager \(SSM\)](#), including access to the *ec2messages* and *ssmmessages* endpoints
- [Amazon Simple Queue Service \(SQS\)](#)
- [AWS Security Token Service \(STS\)](#)
- [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:
 - Full list of permissions

```

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

```

```

    "ec2:StartInstances",
    "ec2:ModifyVpcAttribute",
    "ec2:ReleaseAddress",
    "ec2:TerminateInstances",
    "iam:AddRoleToInstanceProfile",
    "iam:AttachRolePolicy",
    "iam:CreateInstanceProfile",
    "iam:CreatePolicy",
    "iam:CreateRole",
    "iam:CreatePolicyVersion",
    "iam:CreateServiceLinkedRole",
    "iam>DeleteInstanceProfile",
    "iam>DeleteRolePolicy",
    "iam>DeleteRole",
    "iam>DeletePolicy",
    "iam>DeletePolicyVersion",
    "iam:DetachRolePolicy",
    "iam:GetInstanceProfile",
    "iam:GetPolicy",
    "iam:GetRole",
    "iam:GetPolicyVersion",
    "iam:GetAccountSummary",
    "iam:ListAttachedRolePolicies",
    "iam:ListPolicyVersions",
    "iam:ListInstanceProfilesForRole",
    "iam:ListRolePolicies",
    "iam:PassRole",
    "iam:PutRolePolicy",
    "iam:SimulatePrincipalPolicy",
    "iam:UpdateAssumeRolePolicy",
    "kms:Decrypt",
    "kms:DescribeKey",
    "kms:Encrypt",
    "kms:ListAliases",
    "kms:ListKeys",
    "s3:CreateBucket",
    "s3>DeleteObject",
    "s3>DeleteObjectVersion",
    "s3:GetBucketLocation",
    "s3:GetObject",
    "s3:GetObjectRetention",
    "s3:GetObjectVersion",
    "s3:GetBucketObjectLockConfiguration",
    "s3:GetBucketVersioning",
    "s3:ListAllMyBuckets",
    "s3:ListBucketVersions",
    "s3:ListBucket",
    "s3:PutObject",
    "s3:PutObjectRetention",
    "ssm:GetCommandInvocation",
    "ssm:SendCommand",
    "sts:GetCallerIdentity",
    "servicequotas:ListServiceQuotas"
  ],

```

```
    "Resource": "*"
  }
]
```

- › List of permissions to deploy a backup appliance

```

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

```

```

        "ec2:StopInstances",
        "ec2:StartInstances",
        "iam:AddRoleToInstanceProfile",
        "iam:AttachRolePolicy",
        "iam:CreateInstanceProfile",
        "iam:CreatePolicy",
        "iam:CreatePolicyVersion",
        "iam:CreateRole",
        "iam:CreateServiceLinkedRole",
        "iam>DeleteInstanceProfile",
        "iam>DeleteRolePolicy",
        "iam>DeleteRole",
        "iam>DeletePolicy",
        "iam>DeletePolicyVersion",
        "iam:DetachRolePolicy",
        "iam:GetInstanceProfile",
        "iam:GetPolicy",
        "iam:GetRole",
        "iam:GetAccountSummary",
        "iam:GetPolicyVersion",
        "iam:PassRole",
        "iam:PutRolePolicy",
        "iam:SimulatePrincipalPolicy",
        "iam:ListAttachedRolePolicies",
        "iam:ListPolicyVersions",
        "iam:RemoveRoleFromInstanceProfile",
        "iam:UpdateAssumeRolePolicy",
        "ssm:GetCommandInvocation",
        "ssm:SendCommand",
        "sts:GetCallerIdentity",
        "servicequotas:ListServiceQuotas"
    ],
    "Resource": "*"
}
]
}

```

> List of permissions to connect a 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 7.0

```
{
  "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

Veeam Backup & Replication does not check permissions of permissions the *Default Backup Restore* IAM role created on the backup appliance during upgrade to version 7.0. To update permissions of the role, 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": "*"
    }
  ]
}

```

IMPORTANT

Note that the following permissions are only required to remove created resources during appliance deployment: `ec2:DeleteSubnet`, `ec2:DeleteSecurityGroup`, `ec2:DetachInternetGateway`, `ec2:DeleteInternetGateway`, `ec2:DeleteVpc`. If for any security reasons you do not want these permissions to be added, you will need to remove the resources manually in the AWS Management Console in case of a deployment failure or the removal of backup appliances from the backup infrastructure.

- 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 one-time 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 one-time 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 and one-time access keys of IAM users 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 create additional IAM roles to perform operations in this or in other AWS accounts. To learn how to create IAM roles in the AWS Management Console, see [Appendix A. Creating IAM Roles in AWS](#).

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

Service IAM Permissions

You can instruct Veeam Backup for AWS to launch worker instances in the backup account or in production accounts. For more information, see [Managing Worker Configurations](#).

Depending on the type of the account in which you plan to launch worker instances, 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](#).

Service IAM Role in Backup Account

The service IAM role is used to launch worker instances in the backup account to perform backup and restore operations, and to create IAM roles that are attached to the launched 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",
        "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](#).

Service IAM Roles in Production Accounts

Veeam Backup for AWS launches worker instances in production accounts to perform the following operations:

- To index EFS file systems.
- [Applies if enabled in the backup policy or restore settings] To create EC2 image-level backups and to perform restore from EC2 image-level backups.
- [Applies if enabled in the backup policy settings] To create RDS image-level backups and to perform restore from RDS image-level backups.

To launch worker instances in production accounts, Veeam Backup for AWS uses permissions of [backup IAM roles](#) and [restore IAM roles](#). However, the backup and restore IAM roles cannot be used to automatically create IAM roles that will be attached to the launched worker instances for communication with Veeam Backup for AWS. That is why you must either create worker IAM roles manually in the AWS Management Console or instruct Veeam Backup for AWS to do it:

- To create worker IAM roles manually in the AWS Management Console, follow the instructions provided in section [Appendix A. Creating IAM Roles in AWS](#), and assign them permissions listed in section [Indexing Worker IAM Role Permissions](#) or [Worker IAM Role Permissions](#).
- To instruct Veeam Backup for AWS to create worker IAM roles automatically, follow the instructions provided in section [Adding IAM Roles](#).

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 launch worker instances. That is why this role must be assigned permissions listed in section [FLR Worker IAM Role Permissions](#).

Worker Configuration IAM Role Permissions

By default, Veeam Backup for AWS automatically chooses the most appropriate network settings of AWS Regions in production accounts to launch worker instances when performing EFS indexing and RDS backup and restore operations, as well as the default network settings of AWS Regions to launch worker instances 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. When creating new worker configurations, Veeam Backup for AWS uses permissions of worker configuration IAM roles to list network settings available in AWS Regions of production AWS accounts. That is why if you add specific worker configurations that will be used to launch worker instances in production accounts, consider that IAM roles specified in the [worker configuration settings](#) 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"
    }
  ]
}
```

Indexing Worker IAM Role Permissions

When performing [EFS indexing operations](#), Veeam Backup for AWS launches worker instances in the same AWS account to which file systems processed by backup policies belong – production account. That is why Veeam Backup for AWS requires the following IAM role permissions to deploy worker instances when performing EFS indexing operations.

Backup and Restore Permissions

IAM roles require the following permissions to deploy worker instances in production accounts:

- 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 configure trust relationships, see [Before You Begin](#).

- The IAM roles must be granted the following permissions:
 - IAM role permissions specified in 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",
        "ec2:RunInstances",
        "elasticfilesystem:Backup",
        "elasticfilesystem:DescribeAccessPoints",
        "elasticfilesystem:DescribeBackupPolicy",
        "elasticfilesystem:DescribeFileSystemPolicy",
        "elasticfilesystem:DescribeFileSystems",
        "elasticfilesystem:DescribeLifecycleConfiguration",
        "elasticfilesystem:DescribeMountTargets",
        "elasticfilesystem:DescribeMountTargetSecurityGroups",
        "elasticfilesystem:DescribeTags",
        "elasticfilesystem:ListTagsForResource",
        "events>DeleteRule",
        "events:DescribeRule",
        "events>ListTargetsByRule",
        "events:PutRule",
        "events:PutTargets",
        "events:RemoveTargets",
        "iam:GetInstanceProfile",
        "iam:GetContextKeysForPrincipalPolicy",
        "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:SetQueueAttributes",
        "ssm:GetCommandInvocation",
        "ssm:GetParameter",
        "ssm:SendCommand"
    ],
    "Resource": "*"
},
{
    "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"
        }
    }
}
]
}

```

- IAM role permissions specified for manual backup operations

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "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",
        "backup-storage:MountCapsule",
        "ec2:DescribeAvailabilityZones",
        "ec2:DescribeNetworkInterfaceAttribute",
        "ec2:DescribeRegions",
        "elasticfilesystem:Backup",
        "elasticfilesystem:DescribeAccessPoints",
        "elasticfilesystem:DescribeBackupPolicy",
        "elasticfilesystem:DescribeFileSystemPolicy",
        "elasticfilesystem:DescribeFileSystems",
        "elasticfilesystem:DescribeLifecycleConfiguration",
        "elasticfilesystem:DescribeMountTargets",
        "elasticfilesystem:DescribeMountTargetSecurityGroups",
        "elasticfilesystem:DescribeTags",
        "elasticfilesystem:ListTagsForResource",
        "iam:GetContextKeysForPrincipalPolicy",
        "iam:GetRole",
        "iam:PassRole",
        "iam:SimulatePrincipalPolicy",
        "kms:DescribeKey"
      ],
      "Resource": "*"
    },
    {
      "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"
      }
    }
  }
]
}

```

For more information, see [Creating EFS Backups Manually](#).

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

Communication Requirements and Permissions

IAM roles require the following permissions to communicate with worker instances in production accounts:

- 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.

To allow the backup appliance 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": {
        "AWS": "<Role ARN>"
      }
    }
  ]
}

```

Where `<Role ARN>` is the ARN either of the [Impersonation IAM role](#) attached to the backup appliance or of an AWS account to which the backup appliance belongs.

To learn how to configure trust relationships for a role and how to find the ARN of the Impersonation IAM role, 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 configure trust relationships, see [Before You Begin](#).

- 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:ListAccountAliases",
        "iam:ListInstanceProfilesForRole",
        "iam:SimulatePrincipalPolicy",
        "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",
        "sts:AssumeRole"
      ],
      "Resource": "*",
      "Effect": "Allow"
    }
  ]
}

```

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

Worker IAM Role Permissions

You can instruct Veeam Backup for AWS to launch worker instances in production accounts in the following cases:

- When performing image-level backup, entire instance and volume-level restore operations for EC2 instances.

To do that, enable worker deployment in production accounts in [backup policy settings](#), [instance restore settings](#) or [volume-level restore settings](#), and specify IAM roles that will be attached to the worker instances to allow Veeam Backup for AWS to communicate with these instances.

- When performing image-level backup and database restore operations for RDS resources.

To do that, specify IAM roles that will be attached to the worker instances to allow Veeam Backup for AWS to communicate with these instances in [backup policy settings](#) and [database restore settings](#).

Backup and Restore Permissions

IAM roles require the following permissions to deploy worker instances in production accounts:

- › IAM role permissions specified in backup policy settings

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": [
        "ebs:ListChangedBlocks",
        "ebs:ListSnapshotBlocks",
        "ec2:AttachVolume",
        "ec2:AuthorizeSecurityGroupEgress",
        "ec2:AuthorizeSecurityGroupIngress",
        "ec2:CreateKeyPair",
        "ec2:CreateVolume",
        "ec2:CreateSecurityGroup",
        "ec2>DeleteSecurityGroup",
        "ec2:CopySnapshot",
        "ec2:CreateSnapshot",
        "ec2:CreateSnapshots",
        "ec2:CreateTags",
        "ec2:GetEbsDefaultKmsKeyId",
        "ec2>DeleteKeyPair",
        "ec2>DeleteVolume",
        "ec2:DescribeAccountAttributes",
        "ec2:DescribeKeyPairs",
        "ec2:DescribeSecurityGroups",
        "ec2:DetachVolume",
        "ec2:DetachVolume",
        "ec2>DeleteKeyPair",
        "ec2>DeleteSnapshot",
        "ec2>DeleteTags",
        "ec2:DescribeAddresses",
        "ec2:DescribeAvailabilityZones",
        "ec2:DescribeConversionTasks",
        "ec2:DescribeInternetGateways",
        "ec2:DescribeImages",
        "ec2:DescribeInstanceAttribute",
        "ec2:DescribeInstances",
        "ec2:DescribeInstanceTypes",
        "ec2:DescribeNetworkInterfaces",
        "ec2:DescribeRouteTables",
        "ec2:DescribeRegions",
        "ec2:DescribeSecurityGroups",
        "ec2:DescribeSnapshotAttribute",
        "ec2:DescribeSnapshots",
        "ec2:DescribeSubnets",
        "ec2:DescribeTags",
        "ec2:DescribeVolumeAttribute",
        "ec2:DescribeVolumes",
        "ec2:DescribeVpcs",
        "ec2:DescribeVpcEndpoints",
        "ec2:ModifySnapshotAttribute",
        "ec2:ModifyInstanceAttribute",
        "ec2:RevokeSecurityGroupEgress",
        "ec2:RevokeSecurityGroupIngress",
        "ec2:RunInstances",

```

```

        "ec2:StartInstances",
        "ec2:TerminateInstances",
        "events:DeleteRule",
        "events:DescribeRule",
        "events:ListTargetsByRule",
        "events:PutRule",
        "events:PutTargets",
        "events:RemoveTargets",
        "iam:GetRole",
        "iam:GetInstanceProfile",
        "iam:GetContextKeysForPrincipalPolicy",
        "iam:ListAccountAliases",
        "iam:ListInstanceProfiles",
        "iam:ListInstanceProfilesForRole",
        "iam:PassRole",
        "iam:SimulatePrincipalPolicy",
        "kms:CreateGrant",
        "kms:DescribeKey",
        "kms:GetKeyPolicy",
        "kms:ListAliases",
        "kms:ListKeys",
        "kms:ReEncryptFrom",
        "kms:ReEncryptTo",
        "rds:ModifyDBInstance",
        "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:DescribeInstanceInformation",
        "ssm:GetParameter",
        "ssm:GetCommandInvocation",
        "ssm:SendCommand"
    ],
    "Resource": "*",
    "Effect": "Allow"
}
]
}

```

› IAM role permissions specified for restore operations


```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": [
        "ec2:AllocateAddress",
        "ec2:AssignPrivateIpAddresses",
        "ec2:AssociateAddress",
        "ec2:AttachNetworkInterface",
        "ec2:AttachVolume",
        "ec2:AuthorizeSecurityGroupEgress",
        "ec2:AuthorizeSecurityGroupIngress",
        "ec2:CopySnapshot",
        "ec2:CreateKeyPair",
        "ec2:CreateSecurityGroup",
        "ec2:CreateNetworkInterface",
        "ec2:CreateTags",
        "ec2:CreateVolume",
        "ec2>DeleteKeyPair",
        "ec2>DeleteNetworkInterface",
        "ec2>DeleteSnapshot",
        "ec2>DeleteSecurityGroup",
        "ec2>DeleteTags",
        "ec2>DeleteVolume",
        "ec2:DeregisterImage",
        "ec2:DescribeAccountAttributes",
        "ec2:DescribeAddresses",
        "ec2:DescribeAvailabilityZones",
        "ec2:DescribeImages",
        "ec2:DescribeInstanceAttribute",
        "ec2:DescribeInternetGateways",
        "ec2:DescribeInstances",
        "ec2:DescribeInstanceStatus",
        "ec2:DescribeKeyPairs",
        "ec2:DescribeNetworkInterfaces",
        "ec2:DescribeRegions",
        "ec2:DescribeRouteTables",
        "ec2:DescribeSecurityGroups",
        "ec2:DescribeSnapshots",
        "ec2:DescribeSubnets",
        "ec2:DescribeTags",
        "ec2:DescribeVolumes",
        "ec2:DescribeVpcEndpoints",
        "ec2:DescribeVpcs",
        "ec2:DescribeVpcEndpoints",
        "ec2:DetachVolume",
        "ec2:DisassociateAddress",
        "ec2:GetEbsDefaultKmsKeyId",
        "ec2:ImportImage",
        "ec2:ModifyInstanceAttribute",
        "ec2:ModifyNetworkInterfaceAttribute",
        "ec2:ModifySnapshotAttribute",
        "ec2:ModifyVolume",
        "ec2:RevokeSecurityGroupEgress",

```

```

        "ec2:RevokeSecurityGroupIngress",
        "ec2:RunInstances",
        "ec2:StartInstances",
        "ec2:StopInstances",
        "ec2:TerminateInstances",
        "events:DeleteRule",
        "events:DescribeRule",
        "events:ListTargetsByRule",
        "events:PutRule",
        "events:PutTargets",
        "events:RemoveTargets",
        "iam:AddRoleToInstanceProfile",
        "iam:AttachRolePolicy",
        "iam:CreateInstanceProfile",
        "iam>DeleteInstanceProfile",
        "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",
        "kms:CreateGrant",
        "kms:DescribeKey",
        "kms:GetKeyPolicy",
        "kms:ListAliases",
        "kms:ListKeys",
        "kms:ReEncryptFrom",
        "kms:ReEncryptTo",
        "kms:GenerateDataKeyWithoutPlaintext",
        "rds:ModifyDBInstance",
        "s3:GetBucketLocation",
        "servicequotas:ListServiceQuotas"
        "sqs:CreateQueue",
        "sqs>DeleteMessage",
        "sqs>DeleteQueue",
        "sqs:ListQueues",
        "sqs:ReceiveMessage",
        "sqs:SendMessage",
        "ssm:GetCommandInvocation",
        "ssm:GetParameter",
        "ssm:SendCommand"
    ],
    "Resource": "*",
    "Effect": "Allow"
}
]
}

```

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

Communication Requirements and Permissions

IAM roles require the following permissions to communicate with worker instances in production accounts:

- 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.

To allow the backup appliance 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": {
        "AWS": "<Role ARN>"
      }
    }
  ]
}
```

Where `<Role ARN>` is the ARN either of the [Impersonation IAM role](#) attached to the backup appliance or of an AWS account to which the backup appliance belongs.

To learn how to configure trust relationships for a role and how to find the ARN of the Impersonation IAM role, 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 configure trust relationships, see [Before You Begin](#).

- 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](#).

FLR Worker IAM Role Permissions

You can instruct Veeam Backup for AWS to launch worker instances in production accounts when performing file-level recovery operations for EC2 instances. To do that, enable worker deployment in the production account in the [file-level recovery settings](#), and specify an IAM role that will be used to launch worker instances, and then attached to these instances and used by Veeam Backup for AWS to communicate with them.

IAM Role Requirements and Permissions

To allow Veeam Backup for AWS to launch worker instances, 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.

To allow the backup appliance 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": {
        "AWS": "<Role ARN>"
      }
    }
  ]
}
```

Where `<Role ARN>` is the ARN either of the [Impersonation IAM role](#) attached to the backup appliance or of an AWS account to which the backup appliance belongs.

To learn how to configure trust relationships for a role and how to find the ARN of the Impersonation IAM role, 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 configure trust relationships, see [Before You Begin](#).

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",
        "sqs:CreateQueue",
        "sqs:DeleteMessage",
        "sqs:DeleteQueue",
        "sqs:ListQueues",
        "sqs:ReceiveMessage",
        "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 an Amazon S3 bucket 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 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 configure trust relationships, see [Before You Begin](#).

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](#)
- [EFS 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 session.
- 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 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: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 session.
- 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 settings](#) must be granted the following permissions:

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "ec2:DescribeRegions",
        "ec2:DescribeAvailabilityZones",
        "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 Deploy Worker Instances in Production Account

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

```
{
  "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",
        "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 session.
- To create backups of DynamoDB tables.
- To create backup copies, and so on.

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

1. 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 configure trust relationships, see [Before You Begin](#).

2. 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",
        "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](#).

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 session.
- 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 settings must meet the following requirements:

1. 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 configure trust relationships, see [Before You Begin](#).

2. 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: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 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 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",
      "Resource": "*"
    }
  ]
}
```

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 session.
- 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](#)
- [EFS 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:DescribeInstanceStatus",
        "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",
        "iam:GetContextKeysForPrincipalPolicy",
        "iam:ListAccountAliases",
        "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",
        "sqs:CreateQueue",
        "sqs>DeleteMessage",
        "sqs>DeleteQueue",
        "sqs:ListQueues",
        "sqs:ReceiveMessage",
        "sqs:SendMessage",
        "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: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: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",
        "sqs:CreateQueue",
        "sqs>DeleteMessage",
        "sqs>DeleteQueue",
        "sqs:ListQueues",
        "sqs:ReceiveMessage",
        "sqs:SendMessage",
        "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 must be granted specific permissions.

IAM Role Permissions

IAM roles specified in the [restore settings](#) must meet the following requirements:

1. 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 configure trust relationships, see [Before You Begin](#).

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](#).

EFS Restore IAM Permissions

To perform EFS restore operations, IAM roles and IAM users must be granted specific permissions.

IAM Role Permissions

IAM roles specified in the [restore settings](#) must meet the following requirements:

1. 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 configure trust relationships, see [Before You Begin](#).

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: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](#).

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 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 configure trust relationships, see [Before You Begin](#).

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: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",
        "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: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: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",
"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: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: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",
"rds:AddTagsToResource",
"rds:CopyDBClusterSnapshot",
"rds:CopyDBSnapshot",
"rds>CreateDBClusterSnapshot",
"rds>CreateDBSnapshot",
"rds>CreateDbInstance",
"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",
"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",
"sts:AssumeRole"
```



```
    ],  
    "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 6.a to version 7.0, consider that additional permissions must be granted to the IAM roles:

- For Veeam Backup for AWS to be able to use the Standard accelerated mode when performing restore from backups stored in repositories of the S3 Glacier Flexible Retrieval or S3 Glacier Deep Archive storage class, IAM roles [specified in the repository settings](#) must meet the following requirements:
 - a. 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 configure trust relationships, see [Before You Begin](#).

- b. The IAM roles must be granted the following permissions:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "s3:CreateJob",
        "s3:DescribeJob",
        "iam:PassRole"
      ],
      "Resource": "*"
    }
  ]
}
```

If you plan to enable the [private network deployment functionality](#), the IAM roles must be granted the following additional permission:

```
"ec2:DescribeVpcEndpoints"
```

- For Veeam Backup for AWS to be able to perform EC2 file-level recovery from cloud-native snapshots with [product codes](#), the IAM role [specified in the worker settings](#) to deploy worker instances in the backup account, or the IAM role [specified in the restore settings](#) to deploy worker instances in production accounts must be granted the following additional permission:

```
"ec2:DescribeSnapshotAttribute"
```

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](#).

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 format.

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 with only S3 Object Lock enabled is not supported. It is recommended that S3 Object Lock and S3 Versioning are either both enabled or both disabled for a bucket.
- 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.
- You cannot change immutability settings for the repository since these settings are based on the immutability settings of the selected Amazon S3 bucket, which are configured in the AWS Management Console upon bucket creation and cannot be modified afterward. For more information, see [AWS Documentation](#).
- 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.

Backup

When protecting AWS resources, consider the following:

- Veeam Backup for AWS supports backup of the following PostgreSQL versions on Linux machines: PostgreSQL 15, PostgreSQL 14, PostgreSQL 13, PostgreSQL 12.
- 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 does not support backup and restore of RDS [Multi-AZ DB clusters](#).
- Snapshot replication is not supported for Aurora multi-master clusters.
- For Veeam Backup for AWS to be able to create RDS image-level backups, 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](#).
- 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](#).
- 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 of Standalone Backup Appliances](#).
- 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.

Restore

When restoring AWS resources, 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 RDS resources with gp3 storage volumes is not supported. For more information on General Purpose gp3 storage volumes, see [AWS Documentation](#).
- 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.
- 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](#).
- The [AWS Backup](#) service does not support copying DynamoDB backups stored in a cold storage tier to another AWS Region. This 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 where the source tables belong.
- You can change the Time to Live (TTL) setting for DynamoDB tables only an hour after the restore operation completes.

- Veeam Backup for AWS supports restore of EFS file systems only to the same AWS account where the source file systems belong.
- 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 (7.0.0.615 or later).

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	95 MB	1 MB
Snapshots and snapshot replicas	110 MB	1 MB
Snapshots and backups	150 MB	3 MB
Snapshots, snapshot replicas and backups	150 MB	3 MB
RDS Backup Policy		
Snapshots only	100 MB	1 MB

Backup Policy Configuration	RAM Utilization (Default)	Additional RAM (per Workload)
Snapshots and snapshot replicas	125 MB	1 MB
Snapshots and backups	160 MB	3 MB
Snapshots, snapshot replicas and backups	160 MB	3 MB
EFS Backup Policy		
Snapshots only	90 MB	3 MB
Snapshots and backup copies	110 MB	3 MB
Snapshots and indexing	125 MB	3 MB
Snapshots, backup copies and indexing	140 MB	3 MB
DynamoDB Backup Policy		
Snapshots only	90 MB	3 MB
Snapshots and backup copies	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	$4 \text{ GB} - 1.5 \text{ GB} - 4 \text{ GB} * 0.05$ = 2.3 GB

Total RAM	Number of Backup Policies	Workloads per Backup Policy	Utilized RAM ¹	Free RAM ²
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
16 vCPU	> 1,600
RDS CPU	
2 vCPU	< 300
4 vCPU	< 800

Amount of vCPUs	Number of Snapshots Taken Simultaneously
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 launched 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 launched. 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.

Profile	Instance Type	Case
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 details 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_12.7.0.1255.zip` file and launch the `AWSPugin_12.7.0.1255.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 [Installing Plug-In](#) (steps 1-4).
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_12.7.0.1255.exe /silent /accepteula /acceptlicensingpolicy /acceptthirdpartylicenses /acceptrequiredsoftware /uninstall"
/repair	No	Replaces missing files, firewall rules and registry keys. Example: "AWSPlugin_12.7.0.1255.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 [Upgrading 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, RDS resources, backups of EFS file systems and backups of 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 comes as an image of a Linux-based EC2 instance that you can either deploy from AWS Marketplace or from the Veeam Backup & Replication console.

Intended Audience

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.

Deploying Appliance from Console

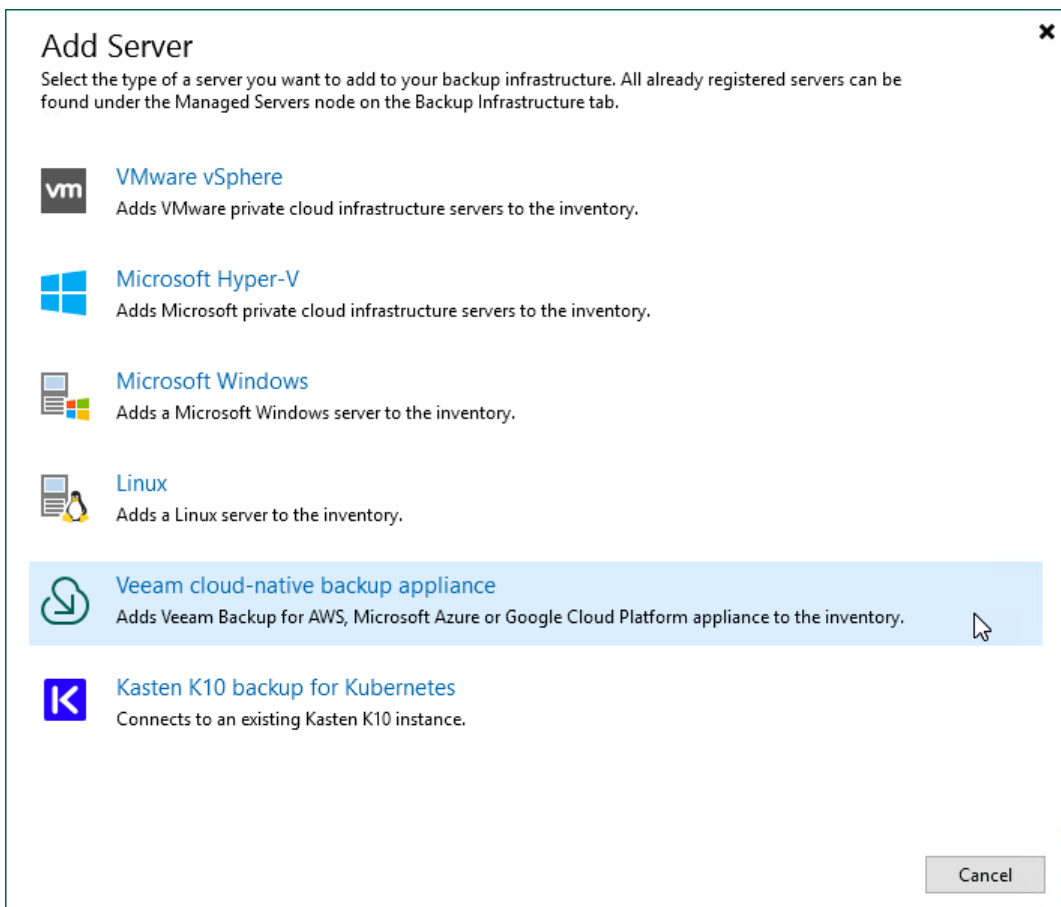
To deploy a new backup appliance from the Veeam Backup & Replication console, do the following:

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

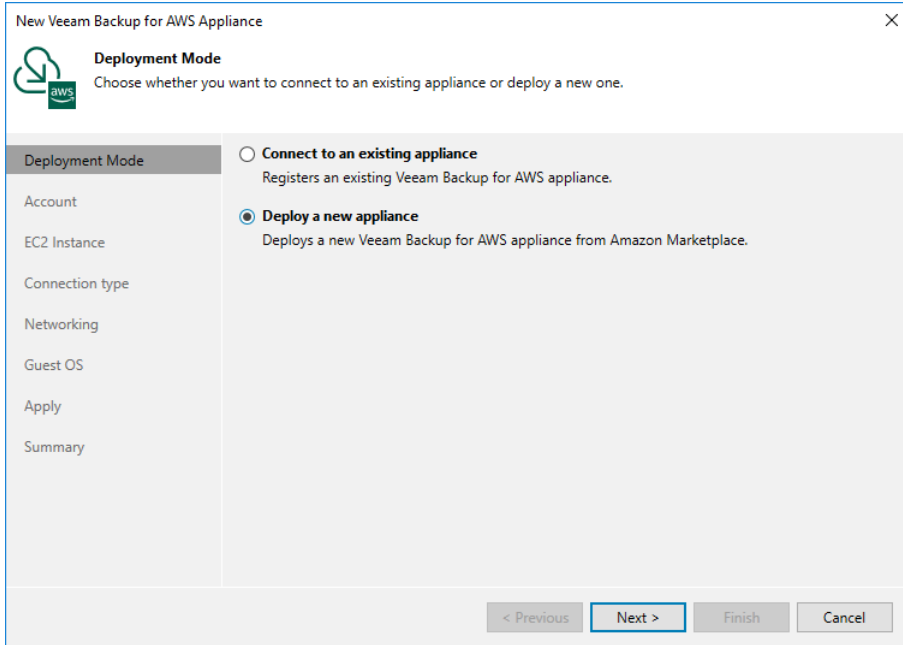
To launch the **New Veeam Backup for AWS Appliance** wizard, do one of the following:

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. 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.

IMPORTANT

An AWS account in which the backup appliance will be deployed must have the *Veeam Backup for AWS FREE Trial & BYOL Edition* subscription in AWS Marketplace. To learn how to subscribe to *Veeam Backup for AWS FREE Trial & BYOL Edition*, see [Installing Veeam Backup for AWS Using CloudFormation Template](#) (the steps 1-4).

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 region availability, Veeam Backup & Replication establishes a temporary test connection to 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.

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](#).

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 the subtitle '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', 'Networking', 'Guest OS', 'Apply', and 'Summary'. The main content area contains three dropdown menus: 'AWS account:' with a key icon and a placeholder 'XXXXXXXXXXXXXXXX (last edited: less than a day ago)', 'AWS region:' with 'Global' selected, and 'Data center:' with 'EU (Paris) (eu-west-3)' selected. There is an 'Add...' button next to the AWS account dropdown and a 'Manage accounts' link below it. A warning icon and text at the bottom state: 'If you have never installed Veeam appliance before, please go to the following link and accept the license agreement. Otherwise the automated deployment process will fail.' followed by the URL 'https://aws.amazon.com/marketplace/pp?sku=d6ajvyz5ma9pyupj1vh21pnt'. At the bottom of the window are four buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'.

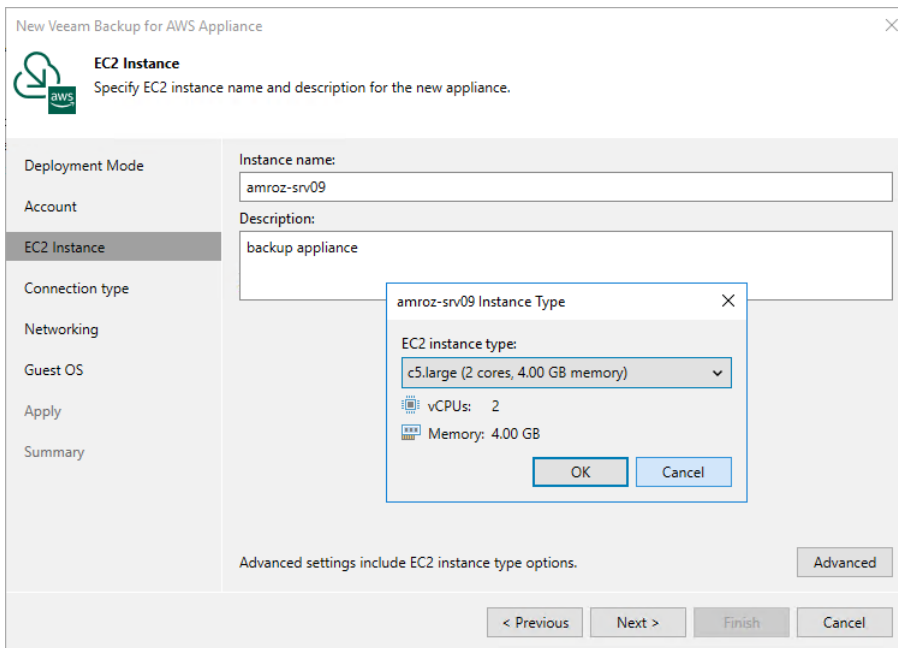
Step 4. Specify EC2 Instance Name and Description

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 displays the following configuration:

- Deployment Mode:** Amazon VPC: amroz-srv VPC
- Account:** Specify Amazon Virtual Private Cloud (VPC) to use.
- EC2 Instance:** Subnet: subnet-08f829b6981cb8585 172.28.0.0/20 (eu-west-3b)
- Connection type:** Choose an IP address range for the selected VPC.
- Networking:** Security group: amroz-srv-VcbSecurityGroup-U56HOSRE8YT9
- Guest OS:** Specify Amazon security group to use.
- Apply:** Backup server public IP address: 62.44.21.21/32
- 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'. The 'Next >' button is highlighted with a red border.

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.

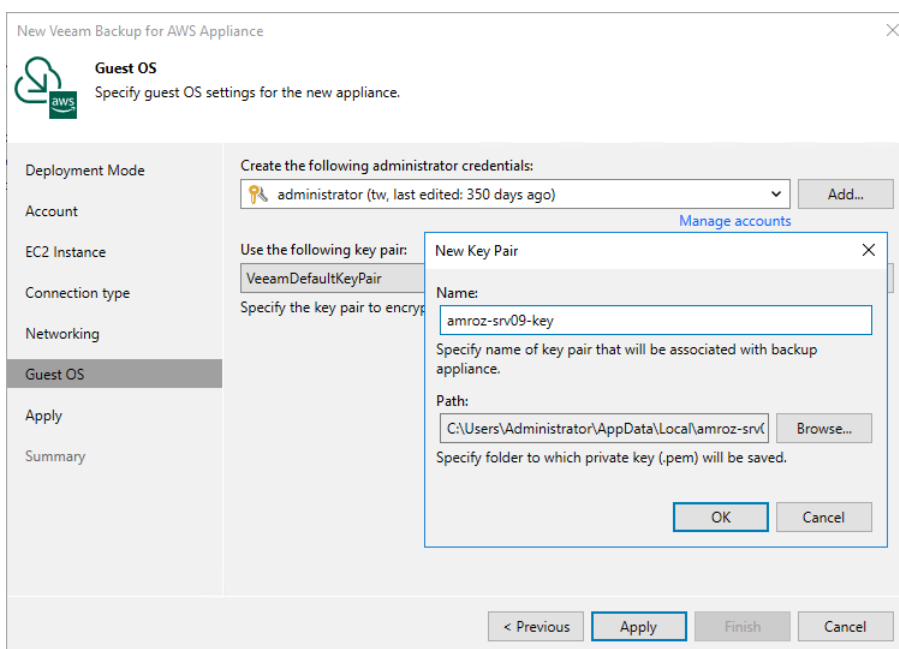
For a user to be displayed in the **Create the following administrator credentials** drop-down list, it must be added to the Credentials Manager. If you have not added a user to the Credential Manager beforehand, you can do it without closing the **New Veeam Backup for AWS Appliance** wizard. To add a new user, click either the **Manage accounts** link or the **Add** button, and specify a 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.

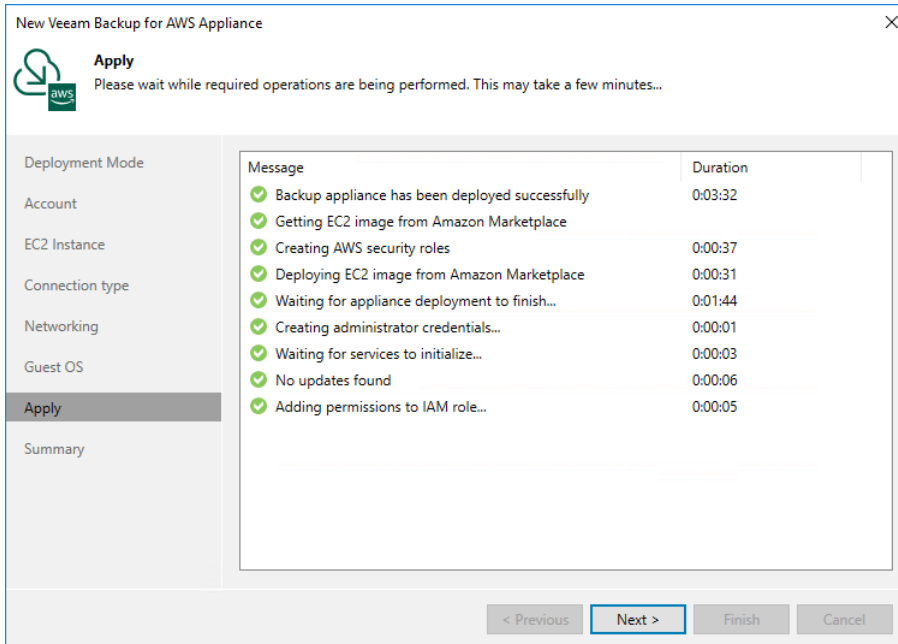
2. From the **Use the following key pair** drop-down list, select a key pair that will be used to authenticate against the backup appliance.

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 a key pair beforehand, you can do it without closing the **Guest OS** wizard. To do that, click **Add** and, in the **New Key Pair** window, specify a name for the private key and the path to a folder where the private key will be located. By default, Veeam Backup & Replication creates a key of *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**.



The screenshot shows a window titled "New Veeam Backup for AWS Appliance" with a close button (X) in the top right corner. The window is in the "Apply" step, indicated by a green checkmark icon and the text "Apply" and "Please wait while required operations are being performed. This may take a few minutes...".

On the left side, there is a navigation pane with the following items: Deployment Mode, Account, EC2 Instance, Connection type, Networking, Guest OS, **Apply** (highlighted), and Summary.

The main area displays a table of progress messages:

Message	Duration
✓ Backup appliance has been deployed successfully	0:03:32
✓ Getting EC2 image from Amazon Marketplace	
✓ Creating AWS security roles	0:00:37
✓ Deploying EC2 image from Amazon Marketplace	0:00:31
✓ Waiting for appliance deployment to finish...	0:01:44
✓ Creating administrator credentials...	0:00:01
✓ Waiting for services to initialize...	0:00:03
✓ No updates found	0:00:06
✓ Adding permissions to IAM role...	0:00:05

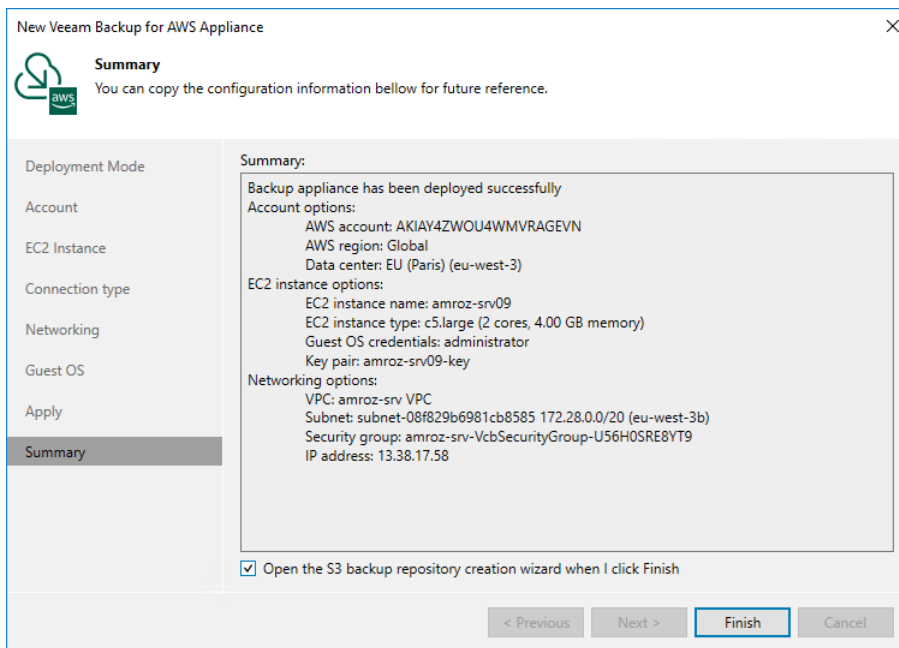
At the bottom of the window, there are four buttons: "< Previous", "Next >" (highlighted with a blue border), "Finish", and "Cancel".

Step 9. Finish Working with Wizard

At the **Summary** step of the wizard, review summary information and click **Finish**. After the backup appliance is deployed, you will be able to configure its settings in the Veeam Backup for AWS Web UI as described in section [Configuration](#).

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](#).



Deploying Appliance from AWS Marketplace

To deploy Veeam Backup for AWS from AWS Marketplace, you can use one the following installation options:

- [Installing Veeam Backup for AWS using a CloudFormation template](#) (recommended) – allows you to deploy Veeam Backup for AWS with most of the backup appliance settings configured out of the box.
- [Launching Veeam Backup for AWS from an Amazon Machine Image \(AMI\)](#) – allows you to configure the backup appliance settings manually when deploying Veeam Backup for AWS.

NOTE

The deployment of Veeam Backup for AWS environment takes approximately 15 minutes.

Installing Veeam Backup for AWS Using CloudFormation Template

Veeam Backup for AWS is installed on a single EC2 instance. The EC2 instance is created during the product installation.

NOTE

When you install the solution using a CloudFormation template, Veeam Backup for AWS automatically creates 2 IAM roles required for the backup appliance configuration and performing backup and disaster recovery operations. These roles have wide scopes of permissions and capabilities. After the deployment completes, you can either limit permissions assigned to the IAM roles or remove the roles and replace them with custom IAM roles created manually. However, this scenario is not preferred. If you want to create the required IAM roles manually, it is recommended that you use the [installing Veeam Backup for AWS from an AMI](#) option.

For more information on the created IAM roles and permissions that must be assigned to them, see [Required IAM Permissions](#).

To install Veeam Backup for AWS using a CloudFormation template:

1. Log in to [AWS Marketplace](#) using credentials of an AWS account in which you plan to install Veeam Backup for AWS.

IMPORTANT

Do not use the root user for login when deploying Veeam Backup for AWS. Deployment or operation of Veeam Backup for AWS does not require the use of root privileges for the AWS account.

You can install Veeam Backup for AWS in the production site – in the AWS account where resources that you plan to back up reside. It is recommended, however, that you use a separate AWS account for Veeam Backup for AWS installation. In this case, if a disaster strikes in the production site, you will still be able to access Veeam Backup for AWS and perform recovery operations.

2. Open the Veeam Backup for AWS overview page for the necessary product edition:
 - [Veeam Backup for AWS Free Edition](#)
 - [Veeam Backup for AWS Paid Edition](#)

- [Veeam Backup for AWS BYOL Edition](#)

For more information on product editions, see [Licensing of Standalone Backup Appliances](#).

3. Click **Continue to Subscribe**.

The screenshot shows the AWS Marketplace product page for Veeam Backup for AWS Paid Edition. The page header includes the AWS Marketplace logo, a search bar, and navigation links such as 'About', 'Categories', 'Delivery Methods', 'Solutions', 'AWS IQ', 'Resources', and 'Your Saved List'. The product title is 'Veeam Backup for AWS Paid Edition' by Veeam, with the latest version being 7.0.0.615. The description states it is 'Automated, secure AWS backup and recovery' for Linux/Unix. The page features a 'Continue to Subscribe' button and a 'Save to List' button. The 'Product Overview' section describes the product as a native, policy-based protection for reliable recovery from accidental deletion, ransomware, and other data loss scenarios. It highlights features like immutable backups using Amazon S3 Object Lock, enterprise scalability, and support for various AWS services including EC2, EBS, RDS, EFS, VPC, S3, and Glacier. A 'Highlights' box lists three key benefits: Relentless security, Fast and reliable recovery, and Zero compromise.

aws marketplace Search Hello, assumed-role/AWSRes...

About Categories Delivery Methods Solutions AWS IQ Resources Your Saved List
Become a Channel Partner Sell in AWS Marketplace Amazon Web Services Home Help

veeam **Veeam Backup for AWS Paid Edition** Continue to Subscribe
By: Veeam Latest Version: 7.0.0.615 Save to List
Automated, secure AWS backup and recovery
Linux/Unix ★★★★★ 5 AWS reviews | 122 external reviews ⓘ

Overview Pricing Usage Support Reviews

Product Overview

Veeam Backup for AWS delivers native, policy-based protection for reliable recovery from accidental deletion, ransomware and other data loss scenarios. An API-first approach, immutable backups and full- and file-level restores ensure resilient protection that's easy and cost-optimized, freeing up time and resources for strategic IT priorities.

New in V6!

- Immutable backups: Data integrity through WORM states using Amazon S3 Object Lock
- Enterprise scalability: Performance and efficiency enhancements for large-scale AWS environments
- And much more!

Supported services:

- Amazon EC2
- Amazon EBS
- Amazon RDS
- Amazon EFS
- Amazon VPC
- Amazon S3, S3 Glacier, and Deep Archive

Highlights

- Relentless security: Secured access and management of data to overcome ransomware and other cyberthreats
- Fast, reliable recovery: Powerful recovery options that keep businesses productive with near-zero recovery time objectives (RTOs)
- Zero compromise: Zero-fuss backup that meets service level agreements (SLAs) and budgetary requirements across the hybrid cloud

4. On the **Subscribe to this software** page, read the product license agreement and click **Continue to Configuration**.

To view the license agreement, expand the details in the **Terms and Conditions** section and click **End User License Agreement**.

aws marketplace Search Hello, assumed-role/AWSRes...

About Categories Delivery Methods Solutions AWS IQ Resources Your Saved List

Become a Channel Partner Sell in AWS Marketplace Amazon Web Services Home Help

veeam Veeam Backup for AWS Paid Edition Continue to Configuration

Veeam Offer

You have subscribed to this software and agreed that your use of this software is subject to the pricing terms and the seller's [End User License Agreement \(EULA\)](#). You agreed that AWS may share information about this transaction (including your payment terms) with the respective seller, reseller or underlying provider, as applicable, in accordance with the [AWS Privacy Notice](#). AWS will issue invoices and collect payments from you on behalf of the seller through your AWS account. Your use of AWS services remains subject to the [AWS Customer Agreement](#) or other agreement with AWS governing your use of such services.

Product	Effective date	Expiration date	Action
Veeam Backup for AWS Paid Edition	12/5/2023	N/A	Hide Details

The following table shows pricing information for the listed software components. You're charged separately for your use of each component.

Veeam Backup for AWS Paid Edition	Additional taxes or fees may apply.
Veeam Backup for AWS Paid Edition	
Unit type	Cost/host/hour
Per 1 instance protected per hour	\$0.005
Per 10 instances protected per hour	\$0.045
Per 100 instances protected per hour	\$0.444
Per 1 Unit protected	\$1.00

[End User License Agreement](#)

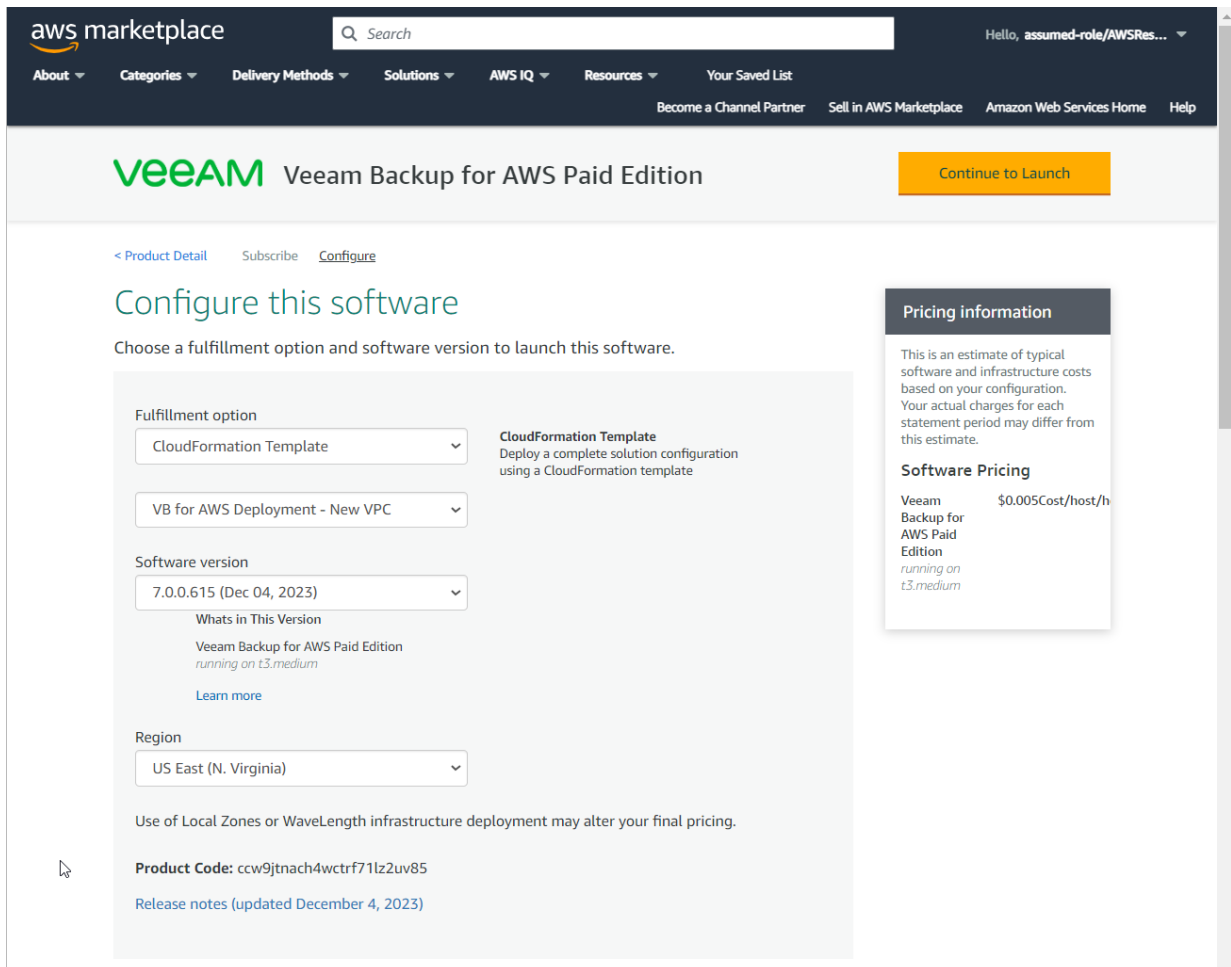
5. On the **Configure this software** page, configure installation settings:
 - a. From the **Fulfillment option** drop-down list, select *CloudFormation Template* and then choose whether you want to connect the EC2 instance running Veeam Backup for AWS to an existing Amazon VPC and subnet, or to create a new Amazon VPC and subnet for the instance.
 - *VB for AWS Deployment - Existing VPC* – select this option if you want to use an existing Amazon VPC and subnet.
 - *VB for AWS Deployment - New VPC* – select this option if you want to create a new Amazon VPC and public subnet. In this case, the VPC and public subnet will be automatically created in the AWS Region in which the appliance will reside; also, an internet gateway will be attached to the VPC.
 - *VB for AWS Deployment - Private VPC* – select this option if you want to create a new Amazon VPC and two subnets (public and private). In this case, the VPC and two subnets will be automatically created in the AWS Region in which the appliance will reside; also, an internet gateway will be attached to the VPC and a NAT gateway will be created in the public subnet.

For more information on Amazon VPCs and subnets, see [AWS Documentation](#).

- b. From the **Software Version** drop-down list, select the latest version of Veeam Backup for AWS.
- c. From the **Region** drop-down list, select an AWS Region in which the EC2 instance running Veeam Backup for AWS will reside.

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

6. Click **Continue to Launch**.



7. On the **Launch this software** page, do the following:

- a. In the **Configuration Details** section, review the product installation settings.
- b. From the **Choose Action** drop-down list, select *Launch CloudFormation*.
- c. Click **Launch**. The **Create stack** wizard will open.

Veeam Backup for AWS is installed using AWS CloudFormation stacks. In AWS CloudFormation, a stack is a collection of AWS services and resources that you can manage as a single unit. You can create a stack in an AWS account, use resources included in the stack to run an application, or delete a stack if you no longer need it. For more information on AWS CloudFormation stacks, see [AWS Documentation](#).

In the **Create stack** wizard, you will create a stack for Veeam Backup for AWS.

The screenshot shows the AWS Marketplace interface for launching Veeam Backup for AWS Paid Edition. The page includes a navigation bar with the AWS Marketplace logo, a search bar, and user information. Below the navigation bar, the product name 'veeam Veeam Backup for AWS Paid Edition' is displayed. A breadcrumb trail shows the path: < Product Detail > Subscribe > Configure > **Launch**. The main heading is 'Launch this software', followed by a sub-heading 'Review the launch configuration details and follow the instructions to launch this software.' The 'Configuration details' section lists: Fulfillment option: VB for AWS Deployment - New VPC, Veeam Backup for AWS Paid Edition, running on t3.medium; Software version: 7.0.0.615; Region: US East (N. Virginia). A 'Usage instructions' button is present. The 'Choose Action' section has a dropdown menu set to 'Launch CloudFormation' and a 'Launch' button.

aws marketplace Hello, assumed-role/AWSRes... ▾

About ▾ Categories ▾ Delivery Methods ▾ Solutions ▾ AWS IQ ▾ Resources ▾ Your Saved List

Become a Channel Partner Sell in AWS Marketplace Amazon Web Services Home Help

veeam Veeam Backup for AWS Paid Edition

< Product Detail Subscribe Configure **Launch**

Launch this software

Review the launch configuration details and follow the instructions to launch this software.

Configuration details

Fulfillment option	VB for AWS Deployment - New VPC Veeam Backup for AWS Paid Edition <i>running on t3.medium</i>
Software version	7.0.0.615
Region	US East (N. Virginia)

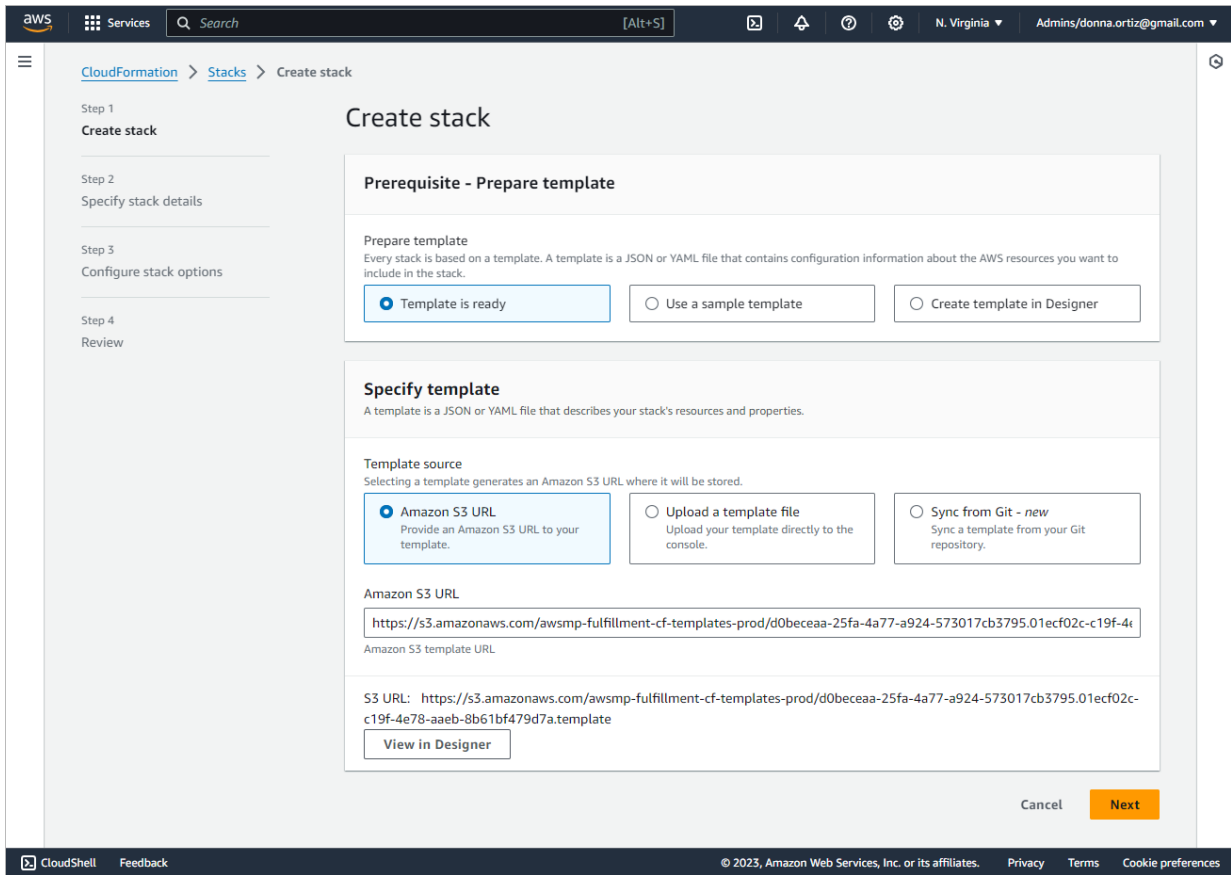
[Usage instructions](#)

Choose Action

▾ Choose this action to launch your configuration through the AWS CloudFormation console.

[Launch](#)

8. At the **Specify template** step of the wizard, the stack template settings are preconfigured by Veeam Backup for AWS and cannot be changed.



9. At the **Specify stack details** step of the wizard, configure the following stack settings:
- In the **Stack name** field, specify a name for the new stack.

Provide a stack name

Stack name

Stack name can include letters (A-Z and a-z), numbers (0-9), and dashes (-).

- In the **Instance Configuration** section, do the following:
 - Select the EC2 instance type for the backup appliance. The recommended EC2 instance type is *t3.medium*.

Veeam Backup for AWS will be deployed on the EC2 instance of the specified instance type with 2 gp3 volumes attached – the root volume with 16 GB of storage capacity and an additional EBS volume with 20 GB of storage capacity. The second volume is intended for storing Veeam Backup for AWS [configuration database](#).

TIP

It is recommended to encrypt your EBS volumes as described in [AWS Documentation](#).

To prevent runtime issues caused by multiple concurrent operations running on the backup appliance, you can later attach an additional EBS volume to the backup appliance and allow the system to allocate its resources in case of memory shortage. For more information, see [Appendix D. Enabling Swap Partition](#).

- ii. Select a key pair that will be used to authenticate against the backup appliance.

For a key pair to be displayed in the **Key pair for Veeam Backup for AWS server** list, it must be created in the Amazon EC2 console. To learn how to create key pairs, see [AWS Documentation](#).

Instance Configuration

Instance type for Veeam Backup for AWS server

Key pair for Veeam Backup for AWS server

Select one, or create a new one at AWS console

- c. In the **Network Configuration** section, do the following:

- i. Select *true* if you want to create an Elastic IP address for the backup appliance.

For more information on Elastic IP addresses, see [AWS Documentation](#).

- ii. Specify the IPv4 address ranges from which Veeam Backup for AWS Web UI will be accessible.

Make sure the IPv4 address of the local machine from which you plan to access Veeam Backup for AWS lies within the specified IPv4 range.

The IPv4 address ranges must be specified in the CIDR notation (for example, `12.23.34.0/24`). To let all IPv4 addresses access Veeam Backup for AWS, you can specify `0.0.0.0/0`. Note that allowing access from all IPv4 addresses is unsafe and thus not recommended in production environments.

Based on the specified IPv4 ranges, AWS CloudFormation will create a security group for Veeam Backup for AWS with an inbound rule for HTTPS traffic. By default, port 443 is open for the inbound HTTPS traffic. If you plan to change the security group for Veeam Backup for AWS upon the product installation, you will need to manually add inbound rules to the new security group and make sure this security group allows access to AWS services listed in the [AWS Services](#) section.

Network Configuration

Create elastic IP for Veeam Backup for AWS server?

By default a dynamic IP will be created (and it could change during reboots of this instance)

Allowed source IP addresses for connection to HTTPS

The IP address range in CIDR format (e.g. `12.23.34.0/24`) from which Veeam Backup for AWS Management portal will be accessible

- d. In the **VPC and Subnet** section, specify an Amazon VPC and subnet to which the backup appliance will be connected.

Depending on the option selected at [step 5a](#), you can either select an existing Amazon VPC and subnet, or specify IPv4 address ranges in the CIDR notation for the new Amazon VPC and subnet.

In case you have chosen the **Private VPC** option, you must specify IPv4 address ranges in the CIDR notation for the private subnet to which the appliance will be connected and for the public subnet in which a NAT gateway that will be created. For more information, see [Backup Appliances in Private Environment](#).

IMPORTANT

Consider the following:

- The specified Amazon VPC and subnet must have the outbound internet access to AWS services listed in the [AWS Services](#) section.
- The specified Amazon VPC and subnet must allow the inbound internet access from the local machine from which you plan to access Veem Backup for AWS.

To learn how to enable internet access for Amazon VPCs and subnets, see [AWS Documentation](#).

VPC and Subnet

VPC CIDR

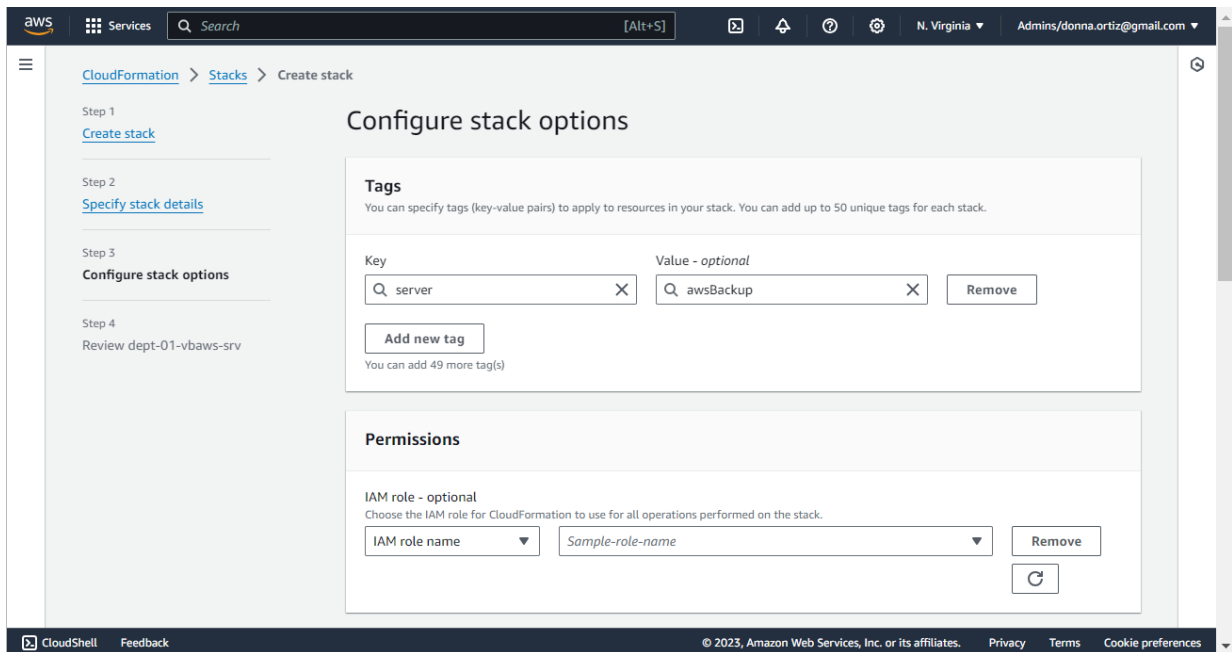
Specify the IPv4 address range as a Classless Inter-Domain Routing (CIDR) block; for example, 172.28.0.0/16. You cannot specify an IPv4 CIDR block larger than /16.

Subnet CIDR

Primary Public Subnet CIDR (Must be within VPC CIDR range).

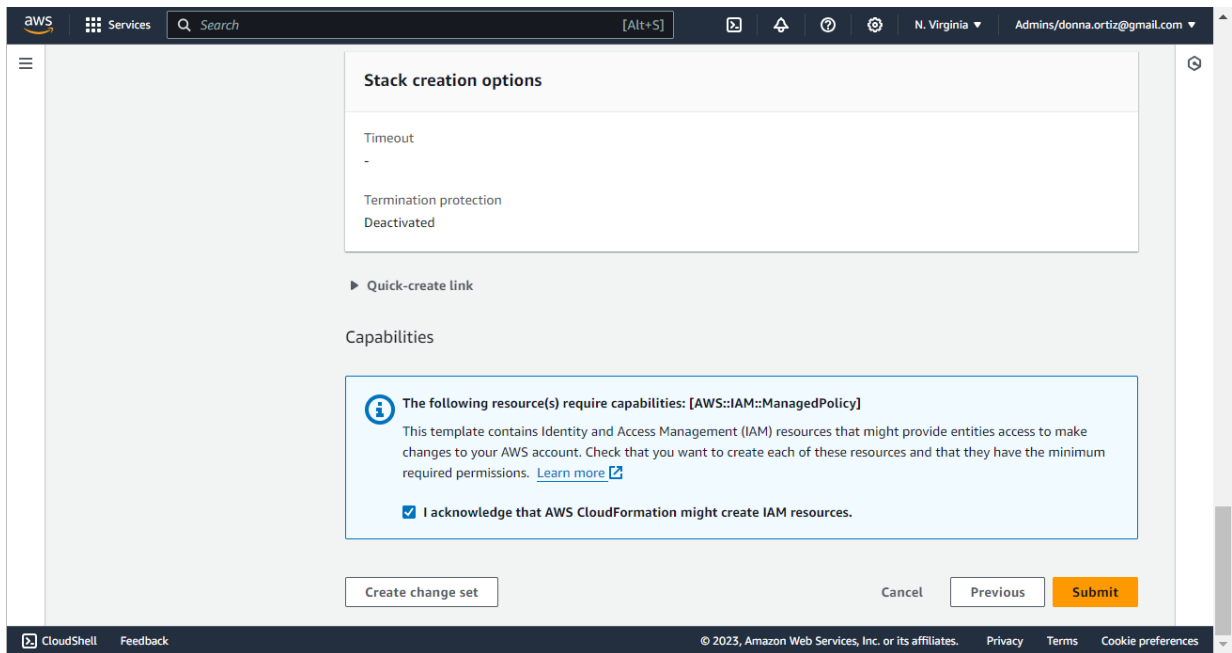
10. At the **Configure stack options** step of the wizard, specify AWS tags, IAM role permissions and other additional settings for the stack.

For more information on available stack options, see [AWS Documentation](#).



11. At the **Review** step of the wizard, do the following:
 - a. Review the configured settings.
 - b. Select the **I acknowledge that AWS CloudFormation might create IAM resources** check box.

c. Click **Submit**.



Right after installation, you must accept license agreements and create a default user. To learn how to do that, see [After You Install](#).

After You Install

After you install Veeam Backup for AWS, you must perform the following steps to start working on the backup appliance:

1. 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 and is available over HTTPS only. For more information, see [Accessing Veeam Backup for AWS](#).

IMPORTANT

Consider the following:

- If the backup device 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).
2. Read and accept the Veeam license agreement, Veeam licensing policy, 3rd party components and software license agreements. If you reject the agreements, you will not be able to continue installation.
 3. In the **Instance ID** field, specify the AWS ID of the EC2 instance running Veeam Backup for AWS to prove that you are the owner of this EC2 instance.

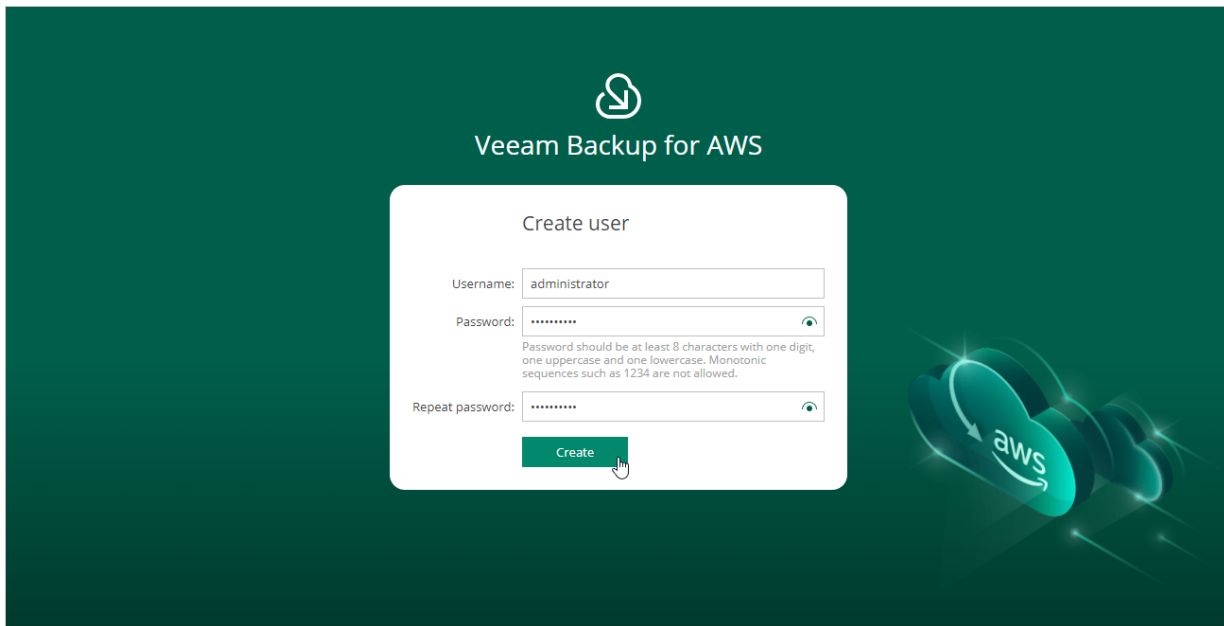
You can find the EC2 instance AWS ID in the in the [AWS Management Console](#).

4. Create a default user whose credentials you will use for your first login to Veeam Backup for AWS. A user name cannot be *admin*, can contain only lowercase Latin letters, numeric characters, underscores and dashes. You can use the dollar sign (\$) as the last character of the name. The maximum length of the name is 32 characters.

Veeam Backup for AWS will create the default user and display the welcome screen where you can [log in](#).

NOTE

To increase the security of the default user account, it is recommended that you enable multi-factor authentication (MFA) for the account after you first log in to Veeam Backup for AWS. To learn how to enable MFA, see [Configuring Multi-Factor Authentication](#).



Installing Veeam Backup for AWS from AMI

Veeam Backup for AWS is installed on a single EC2 instance. The EC2 instance is created during the product installation.

IMPORTANT

After you install Veeam Backup for AWS from the Amazon Machine Image (AMI), you will be asked to provide one-time access keys of an IAM user that Veeam Backup for AWS will use to create IAM roles required for the backup appliance configuration. If you do not want to provide the keys, you can create the required IAM roles manually before you begin the installation. For more information on the required IAM roles, see [Required IAM Permissions](#).

To install Veeam Backup for AWS from an AMI:

1. Log in to [AWS Marketplace](#) using credentials of an AWS account in which you plan to install Veeam Backup for AWS.

IMPORTANT

Do not use the root user for login when deploying Veeam Backup for AWS. Deployment or operation of Veeam Backup for AWS does not require the use of root privileges for the AWS account.

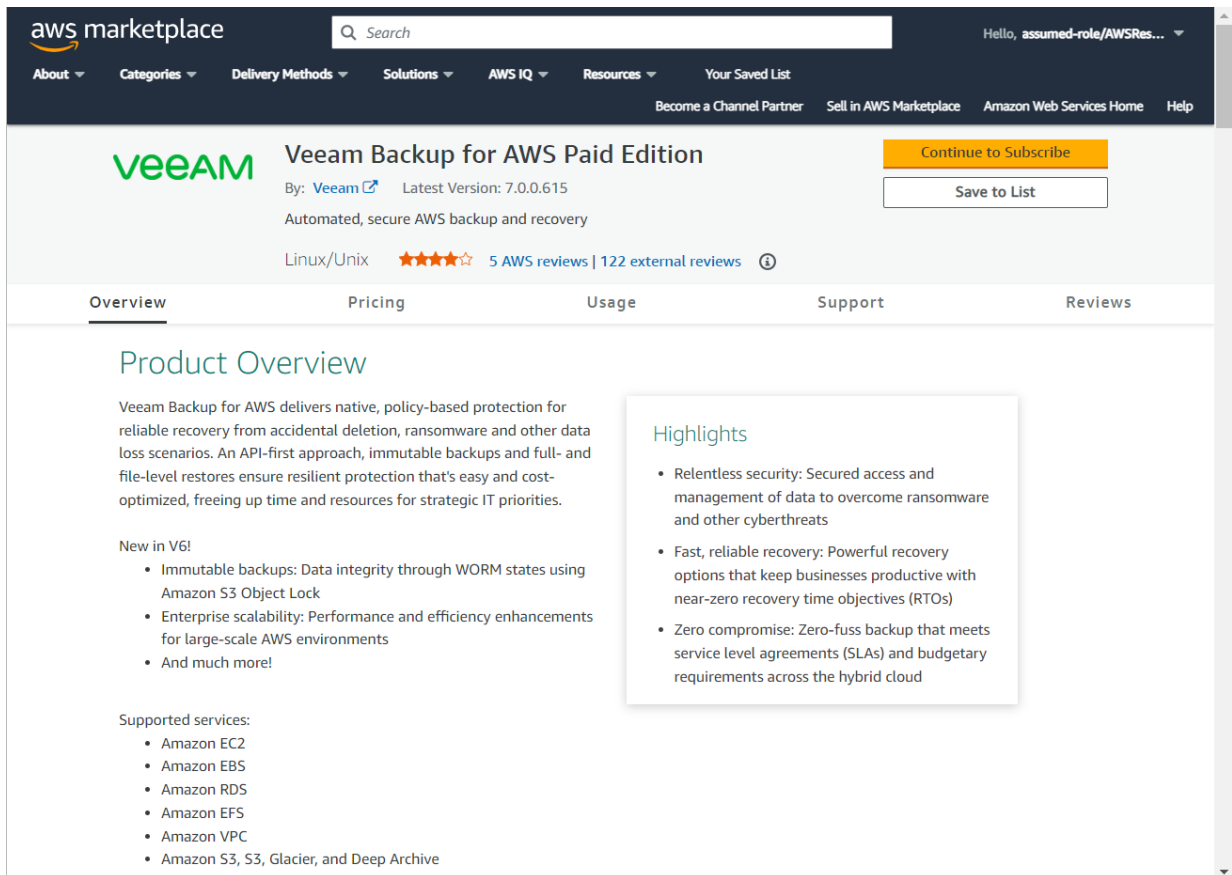
You can install Veeam Backup for AWS in the production site – in the AWS account where resources that you plan to back up reside. It is recommended, however, that you use a separate AWS account for Veeam Backup for AWS installation. In this case, if a disaster strikes in the production site, you will still be able to access Veeam Backup for AWS and perform recovery operations.

2. Open the Veeam Backup for AWS overview page for the necessary product edition:

- [Veeam Backup for AWS Free Edition](#)
- [Veeam Backup for AWS Paid Edition](#)
- [Veeam Backup for AWS BYOL Edition](#)

For more information on product editions, see [Licensing of Standalone Backup Appliances](#).

3. Click **Continue to Subscribe**.



The screenshot shows the AWS Marketplace page for Veeam Backup for AWS Paid Edition. The page header includes the AWS Marketplace logo, a search bar, and the user's name 'Hello, assumed-role/AWSRes...'. The main navigation bar contains links for 'About', 'Categories', 'Delivery Methods', 'Solutions', 'AWS IQ', 'Resources', and 'Your Saved List'. Below the navigation bar, there are links for 'Become a Channel Partner', 'Sell in AWS Marketplace', 'Amazon Web Services Home', and 'Help'. The product title 'Veeam Backup for AWS Paid Edition' is prominently displayed, along with the Veeam logo. Below the title, it says 'By: Veeam' and 'Latest Version: 7.0.0.615'. The description reads 'Automated, secure AWS backup and recovery'. There are also links for 'Linux/Unix', '5 AWS reviews | 122 external reviews', and an information icon. A yellow 'Continue to Subscribe' button and a 'Save to List' button are visible. The page has a tabbed interface with 'Overview', 'Pricing', 'Usage', 'Support', and 'Reviews'. The 'Overview' tab is selected, showing a 'Product Overview' section with a description of the product's capabilities. A 'Highlights' box on the right lists three key features: 'Relentless security', 'Fast, reliable recovery', and 'Zero compromise'. The 'Supported services' section lists various AWS services supported by the product.

aws marketplace Hello, assumed-role/AWSRes...
About Categories Delivery Methods Solutions AWS IQ Resources Your Saved List
Become a Channel Partner Sell in AWS Marketplace Amazon Web Services Home Help

veeam **Veeam Backup for AWS Paid Edition** [Continue to Subscribe](#)
By: [Veeam](#) Latest Version: 7.0.0.615
Automated, secure AWS backup and recovery
Linux/Unix ★★★★☆ [5 AWS reviews](#) | [122 external reviews](#) ⓘ

[Overview](#) [Pricing](#) [Usage](#) [Support](#) [Reviews](#)

Product Overview

Veeam Backup for AWS delivers native, policy-based protection for reliable recovery from accidental deletion, ransomware and other data loss scenarios. An API-first approach, immutable backups and full- and file-level restores ensure resilient protection that's easy and cost-optimized, freeing up time and resources for strategic IT priorities.

New in V6!

- Immutable backups: Data integrity through WORM states using Amazon S3 Object Lock
- Enterprise scalability: Performance and efficiency enhancements for large-scale AWS environments
- And much more!

Supported services:

- Amazon EC2
- Amazon EBS
- Amazon RDS
- Amazon EFS
- Amazon VPC
- Amazon S3, S3, Glacier, and Deep Archive

Highlights

- Relentless security: Secured access and management of data to overcome ransomware and other cyberthreats
- Fast, reliable recovery: Powerful recovery options that keep businesses productive with near-zero recovery time objectives (RTOs)
- Zero compromise: Zero-fuss backup that meets service level agreements (SLAs) and budgetary requirements across the hybrid cloud

4. On the **Subscribe to this software** page, read the product license agreement and click **Continue to Configuration**.

To view the license agreement, expand the details in the **Terms and Conditions** section and click **End User License Agreement**.

The screenshot shows the AWS Marketplace interface for Veeam Backup for AWS Paid Edition. At the top, there is a search bar and navigation links. The main heading is 'Veeam Backup for AWS Paid Edition' with a yellow 'Continue to Configuration' button. Below this, a 'Veeam Offer' section contains a paragraph of terms and conditions. A table lists the product details:

Product	Effective date	Expiration date	Action
Veeam Backup for AWS Paid Edition	12/5/2023	N/A	Hide Details

Below the table, a section titled 'Veeam Backup for AWS Paid Edition' provides pricing information. It includes a table of unit types and costs:

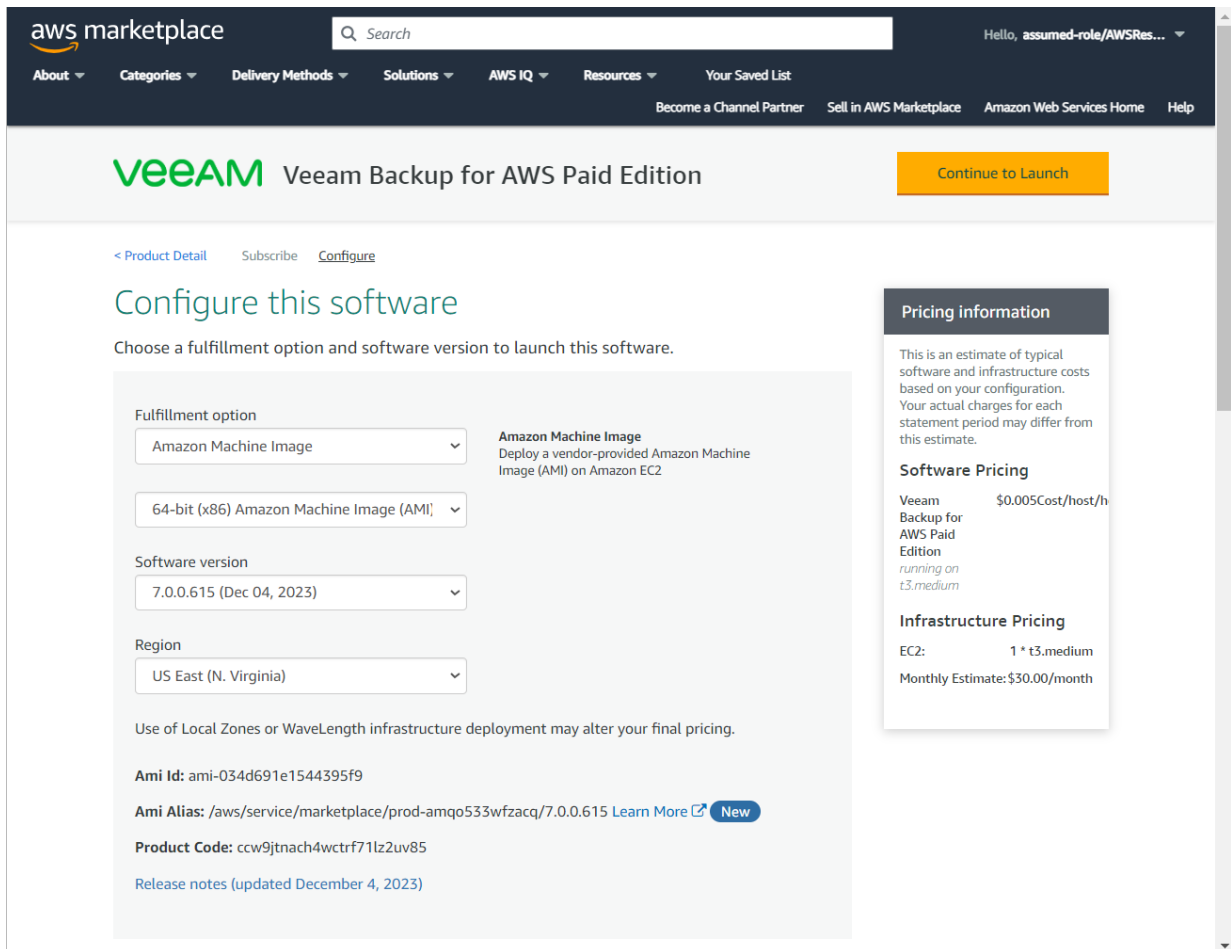
Unit type	Cost/host/hour
Per 1 instance protected per hour	\$0.005
Per 10 instances protected per hour	\$0.045
Per 100 instances protected per hour	\$0.444
Per 1 Unit protected	\$1.00

At the bottom of this section, there is a link to the 'End User License Agreement'.

5. On the **Configure this software** page, configure installation settings:
 - a. From the **Fulfillment option** drop-down list, select *Amazon Machine Image*.
 - b. From the **Software Version** drop-down list, select the latest version of Veeam Backup for AWS.
 - c. From the **Region** drop-down list, select an AWS Region in which the EC2 instance running Veeam Backup for AWS will reside.

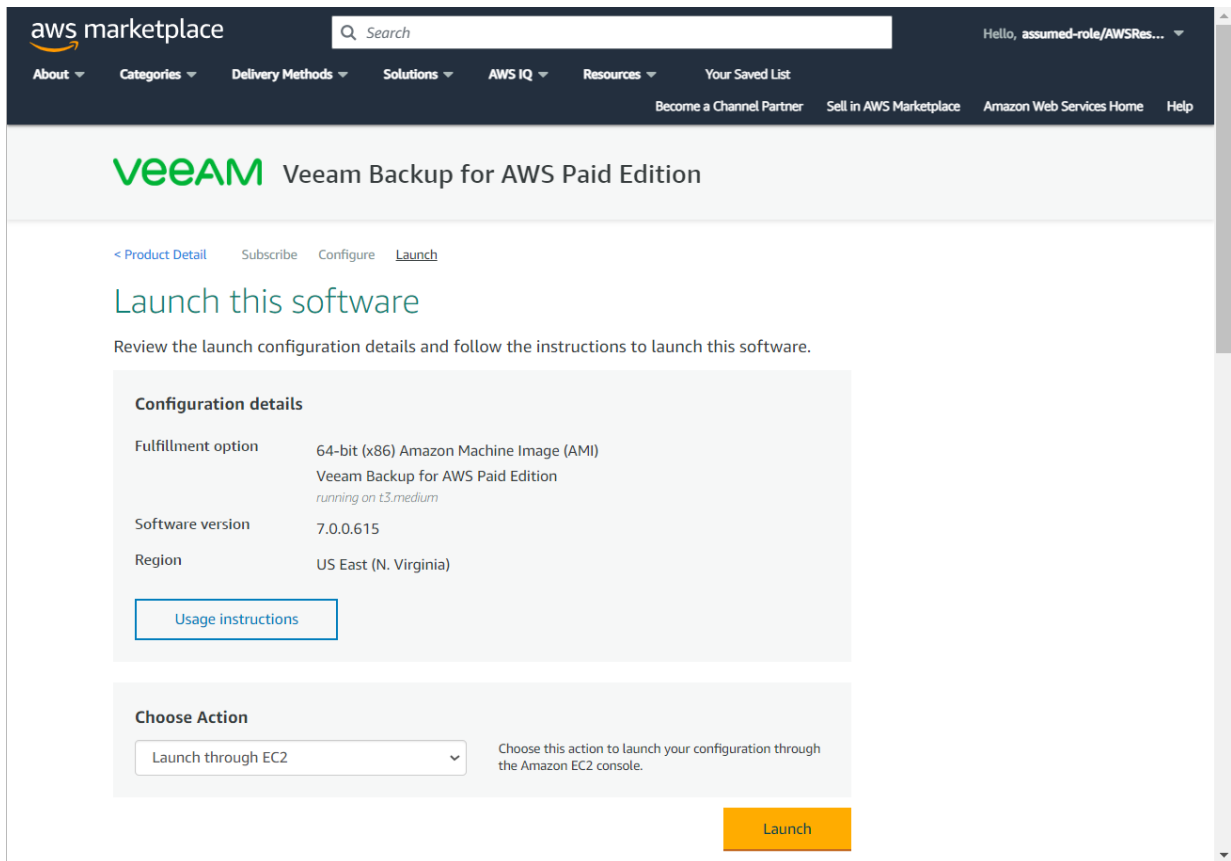
For more information on AWS Regions, see [AWS Documentation](#).

6. Click **Continue to Launch**.

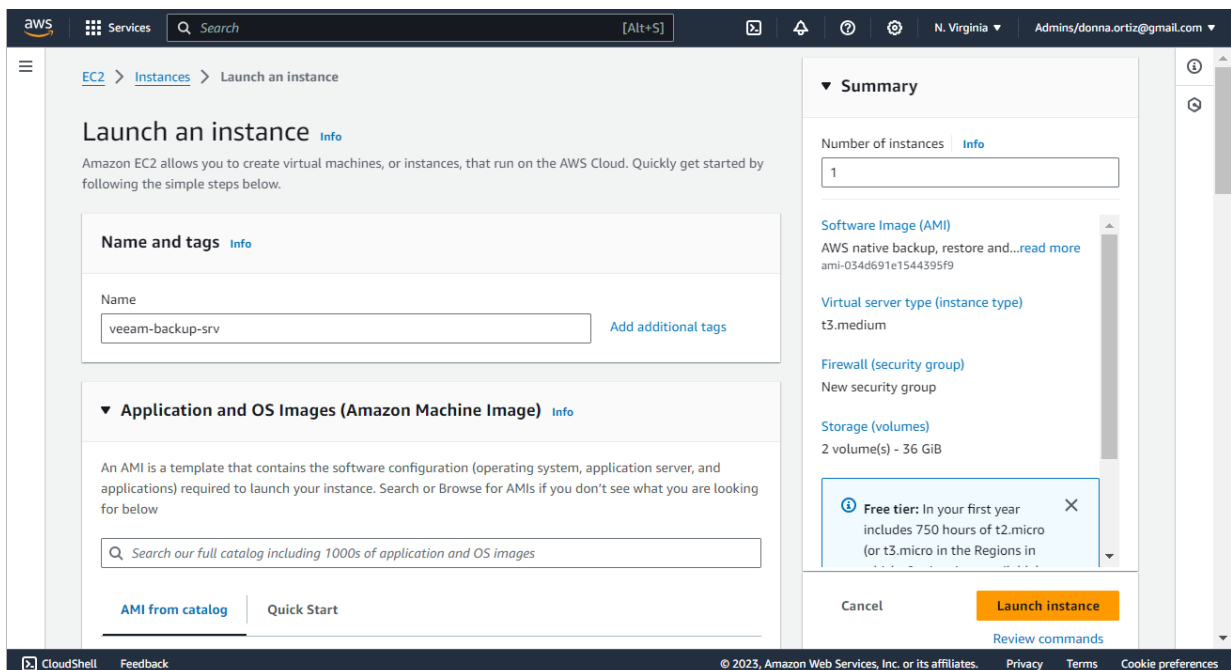


7. On the **Launch this software** page, do the following:
- a. In the **Configuration Details** section, review the product installation settings.
 - b. From the **Choose Action** drop-down list, select *Launch through EC2*.

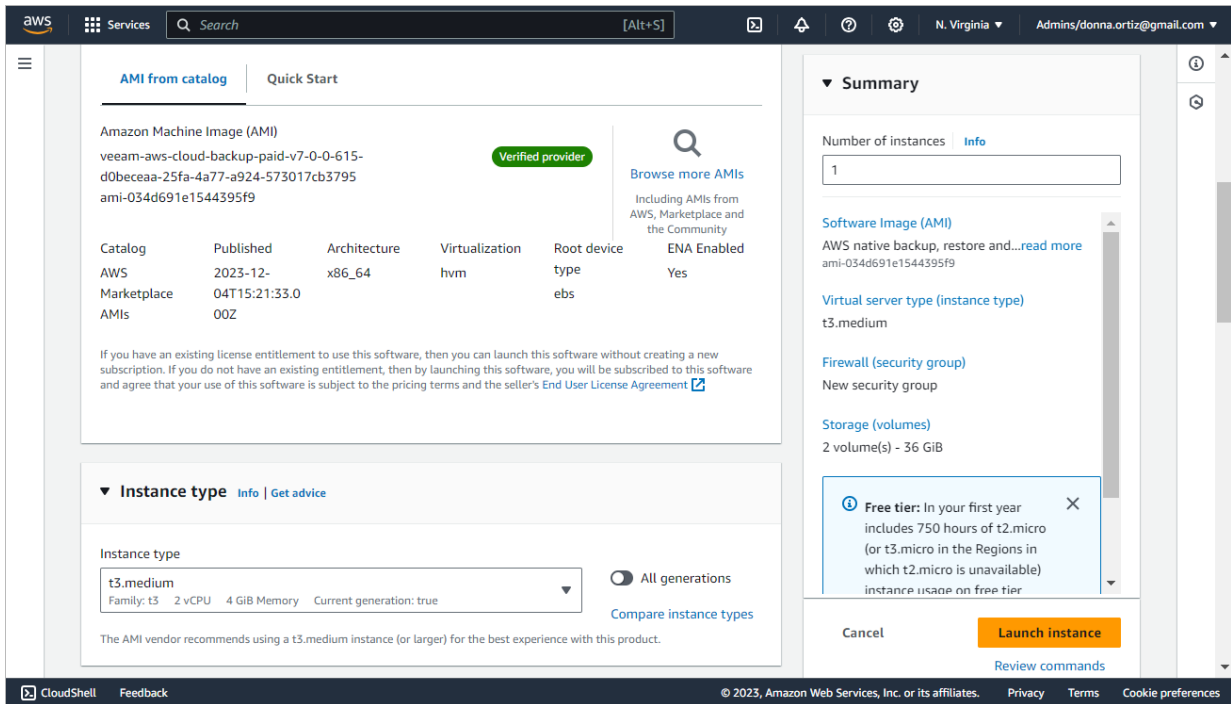
c. Click **Launch**. The **Launch an instance** wizard will open.



8. At the **Name and tags** step of the wizard, you can specify a name that will help you easily identify and locate the appliance and AWS tags that will be assigned to the instance.



9. At the **Instance type** step of the wizard, select an EC2 instance type for the backup appliance. The minimum recommended EC2 instance type is *t3.medium*.



10. At the **Key pair (login)** step of the wizard, specify a key pair that will be used to authenticate against the backup appliance. You can select an existing key pair or create a new one.

For a key pair to be displayed in the **Key pair name** drop-down list, it must be created in the Amazon EC2 console. To learn how to create key pairs, see [AWS Documentation](#).

11. At the **Network settings** step of the wizard, do the following:
- Click **Edit**.
 - In the **Network** and **Subnet** fields, specify an Amazon VPC and subnet to which the backup appliance will be connected. You can either select an existing Amazon VPC and subnet, or create a new subnet. For more information on Amazon VPCs and subnets, see [AWS Documentation](#).

IMPORTANT

Consider the following:

- The specified Amazon VPC and subnet must have the outbound internet access to AWS services listed in the [AWS Services](#) section.
- The specified Amazon VPC and subnet must allow the inbound internet access from the local machine that you plan to use to access Veeam Backup for AWS.

To learn how to enable internet access for Amazon VPCs and subnets, see [AWS Documentation](#).

- From the **Auto-assign Public IP** drop-down list, select **Enable**.

If you want the backup appliance to be deployed without a public IP address, you will have to manually configure access both to the AWS services and the internet in the way that suits your security concerns best. For more information, see [Backup Appliances in Private Environment](#).

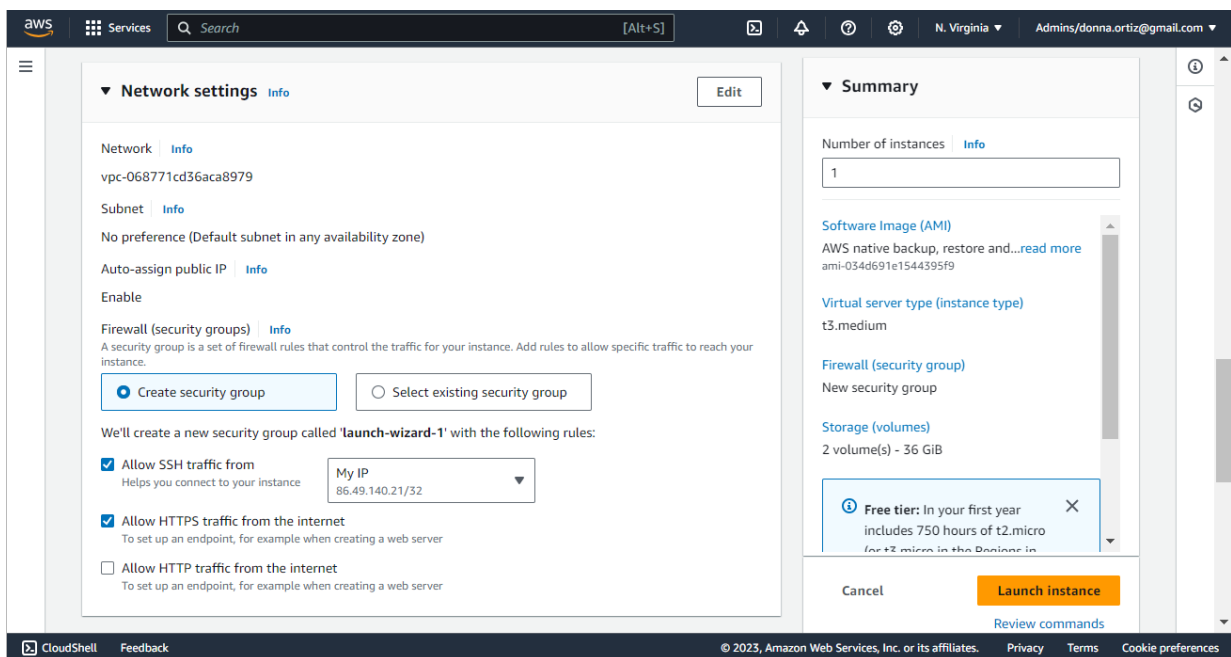
- d. Choose a security group that will control the inbound and outbound traffic for the backup appliance. You can either associate an existing security group with the backup appliance or create a new security group. If you choose an existing security group, make sure it allows access to AWS services listed in the [AWS Services](#) section.

If you choose to create a new security group, add a new inbound rule for the HTTPS traffic:

- i. In the **Inbound security groups rules** section, click **Add security group rule**. The **Security group rule 2** settings will appear.
- ii. Select *HTTPS* from the **Type** drop-down list.
- iii. Select *Custom* from the **Source type** drop-down list.
- iv. In the **Source** field, specify IPv4 address ranges from which Veeam Backup for AWS Web UI will be accessible.

Make sure the IPv4 address of the local machine from which you plan to access Veeam Backup for AWS lies within the specified IPv4 ranges.

IPv4 address ranges must be specified in the CIDR notation (for example, 12.23.34.0/24). To allow unrestricted access to the backup appliance, you can specify 0.0.0.0/0. However, the latter is not recommended since unrestricted access to Veeam Backup for AWS can violate your organization security policy.

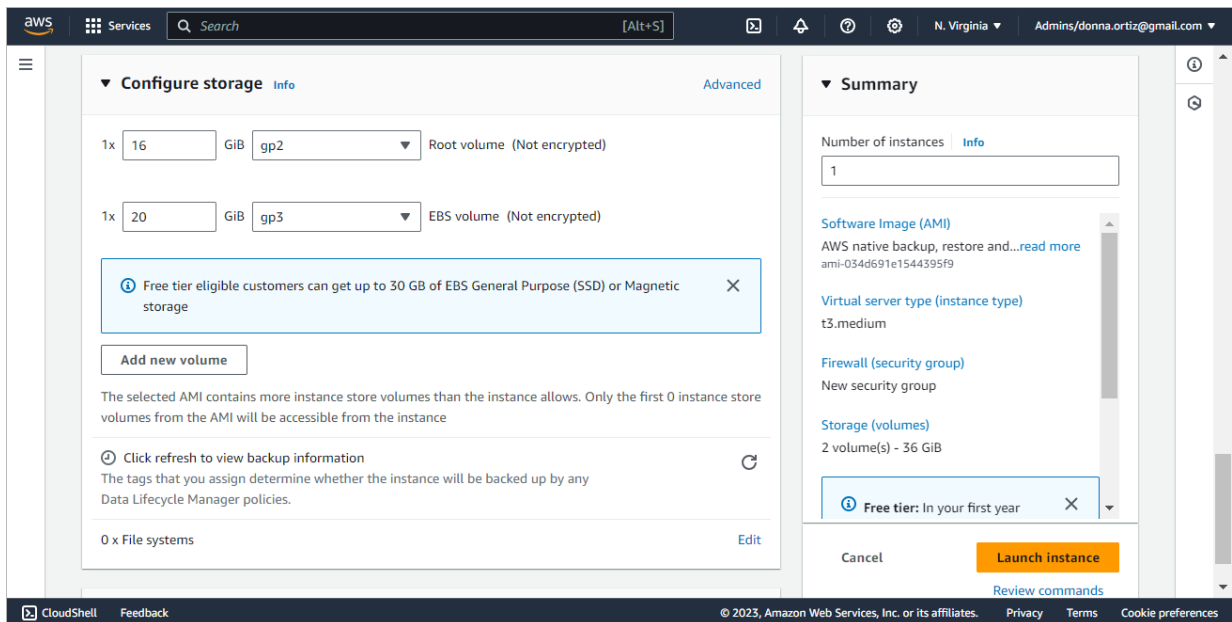


12. At the **Configure storage** step of the wizard, review the preconfigured storage settings and proceed to the next step. For technical reasons, it is not recommended to change these settings.

The EC2 instance will be created with 2 gp3 volumes attached – the root volume with 16 GB of storage capacity and an additional EBS volume with 20 GB of storage capacity. The second volume is intended for storing Veeam Backup for AWS [configuration database](#).

TIP

To prevent runtime issues caused by multiple concurrent operations running on the backup appliance, you can later attach an additional EBS volume to the backup appliance and allow the system to allocate its resources in case of memory shortage. For more information, see [Appendix D. Enabling Swap Partition](#).



13. At the **Advanced details** step of the wizard, do the following:
 - a. [Applies if you have created [IAM roles required for the product installation](#) beforehand] In the **IAM instance profile** field, specify the *Impersonation* IAM role that will be attached to the backup appliance. This role will allow Veeam Backup for AWS to assume IAM roles to perform backup and restore operations.
 - b. Enable access to the instance metadata to allow Veeam Backup for AWS to use the Instance Metadata Service (IMDS) to be able to configure and manage the running backup appliance. To do that, select *Enabled* from the **Metadata accessible** drop-down list.
 - c. Configure additional settings for the backup appliance to meet your organization requirements. To learn how to configure Amazon Linux instances, see [AWS Documentation](#).
14. In the **Summary** section, review the configured settings and click **Launch instance**.

Right after installation, you must perform a number of additional actions for the backup appliance configuration. For more information, see [After You Install](#).

Required IAM Permissions

When you [install the solution using CloudFormation Template](#), Veeam Backup for AWS creates 2 IAM roles:

- **Impersonation IAM role** – is attached to the backup appliance and is then used to assume other IAM roles added to Veeam Backup for AWS.
- **Default Backup Restore IAM role** – is automatically added to Veeam Backup for AWS and is assigned all the permissions required to perform operations within the initial AWS account. For example, the role is used to back up AWS resources within the account, to store backups in any Amazon S3 bucket within the account, and so on.

When you [install the solution from the AMI](#), you can either create these IAM roles manually, or instruct Veeam Backup for AWS to use one-time access keys for automatic creation of the required IAM roles.

Using One-Time Access Keys

If you choose to use one-time keys of an IAM user to create IAM roles automatically, no additional steps are required before or during Veeam Backup for AWS installation. However, after installation, you must instruct Veeam Backup for AWS to automatically create IAM roles required for the backup appliance configuration. To learn how to do that, see [After You Install](#).

The IAM user must have the following permissions:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "VisualEditor0",
      "Effect": "Allow",
      "Action": [
        "cloudwatch:DeleteAlarms",
        "cloudwatch:PutMetricAlarm",
        "ec2:AssociateIamInstanceProfile",
        "ec2:CreateTags",
        "ec2:DescribeIamInstanceProfileAssociations",
        "ec2:DescribeInstances",
        "ec2:DisassociateIamInstanceProfile",
        "iam:AddRoleToInstanceProfile",
        "iam:AttachRolePolicy",
        "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>ListInstanceProfiles",
        "iam>ListPolicyVersions",
        "iam:PassRole",
        "iam:PutRolePolicy",
        "iam:RemoveRoleFromInstanceProfile",
        "iam:SimulatePrincipalPolicy"
      ],
      "Resource": "*"
    }
  ]
}
```

Creating IAM Roles Manually

If you choose to create IAM roles manually, you must do this in the AWS Management Console before you start installing Veeam Backup for AWS. To learn how to create IAM roles, see [Appendix A. Creating IAM Roles in AWS](#).

The created IAM roles must have specific permissions:

- The *Impersonation* IAM role attached to the backup appliance operating in the *BYOL* or *Free* license edition must have the following permissions:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "sts:AssumeRole"
      ],
      "Resource": "*"
    }
  ]
}
```

- The *Impersonation* IAM role attached to the backup appliance operating in the *Paid* license edition must have the following permissions:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": [
        "aws-marketplace:MeterUsage"
      ],
      "Resource": "*",
      "Effect": "Allow"
    },
    {
      "Action": [
        "sts:AssumeRole"
      ],
      "Resource": "*",
      "Effect": "Allow"
    }
  ]
}
```

- The *Default Backup Restore* IAM role must meet the following requirements:
 -

- You must allow the *Impersonation* IAM role to assume the *Default Backup Restore* IAM role. To do that, 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": {
        "AWS": "<Role ARN>"
      }
    }
  ]
}
```

To learn how to configure trust relationships, see [Before You Begin](#).

- The *Default Backup Restore* IAM role must have permissions required to perform all operations available in Veeam Backup for AWS within the initial AWS account. For more information on the required permissions, see [Full List of IAM Permissions](#).

However, if you plan to use this role for specific operations or do not plan to use this role at all, you can assign the role granular permissions. For more information, see [IAM Permissions](#).

TIP

You will be able to add other IAM roles later, after Veeam Backup for AWS installation. For more information, see [Managing IAM Roles](#).

After You Install

To start working with Veeam Backup for AWS, you must perform the initial configuration of the backup appliance. To do that, 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 and is available over HTTPS only. For more information, see [Accessing Veeam Backup for AWS](#).

IMPORTANT

Consider the following:

- If the backup device 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).

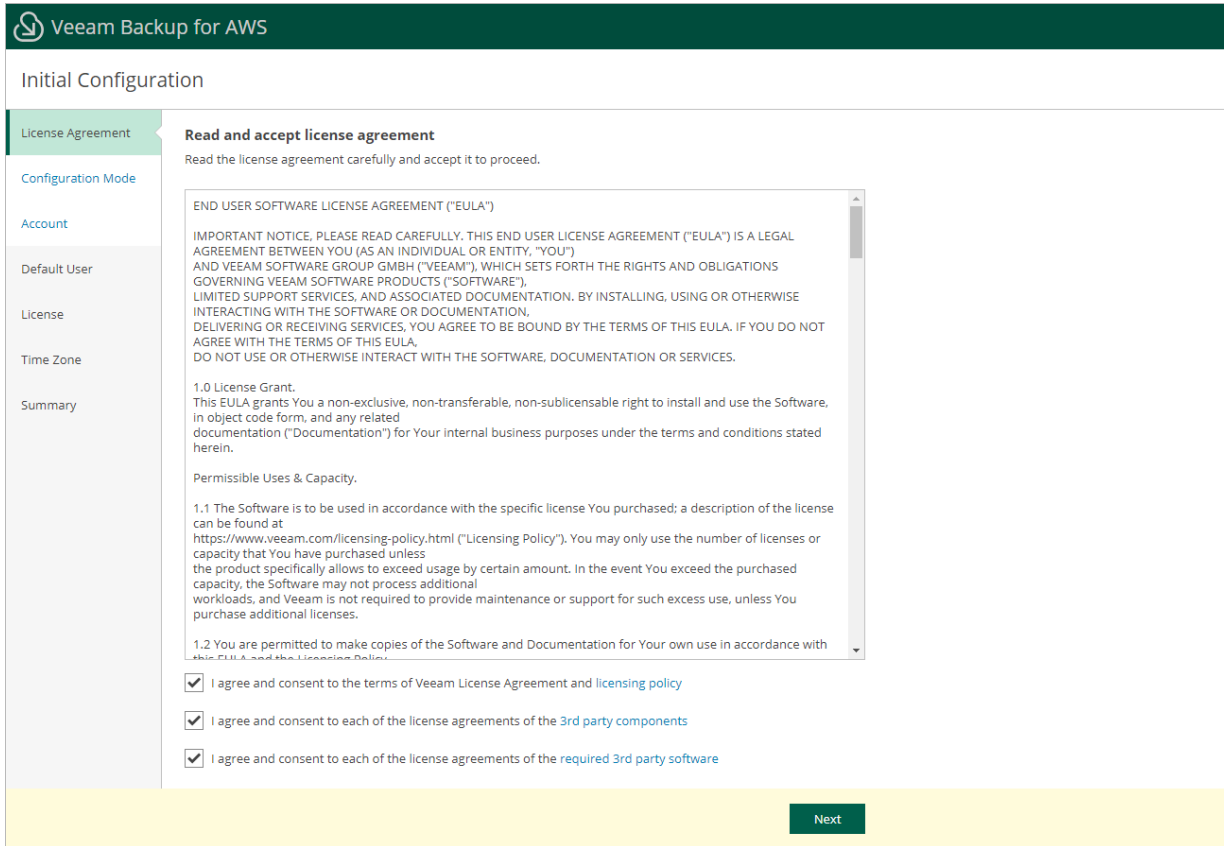
To configure backup appliance settings, complete the **Initial Configuration** wizard:

1. [Read and accept license agreements](#).
2. [Choose a configuration mode](#).
3. [Specify an IAM identity](#).
4. [Create the default user](#).

5. [Install a Veeam Backup for AWS license.](#)
6. [Specify a time zone.](#)
7. [Finish working with the wizard.](#)

Step 1. Accept License Agreement

At the **License Agreement** step of the wizard, read and accept the Veeam license agreement, Veeam licensing policy, 3rd party components and software license agreements. If you reject the agreements, you will not be able to continue installation.



The screenshot shows the 'Initial Configuration' wizard for Veeam Backup for AWS. The 'License Agreement' step is active, displaying the 'Read and accept license agreement' section. A scrollable text area contains the 'END USER SOFTWARE LICENSE AGREEMENT ("EULA")' with sections for '1.0 License Grant' and '1.1 Permissible Uses & Capacity'. Below the text area are three checked checkboxes for accepting the Veeam license agreement, 3rd party components, and required 3rd party software. A 'Next' button is located at the bottom right of the wizard.

Initial Configuration

License Agreement **Read and accept license agreement**
Read the license agreement carefully and accept it to proceed.

END USER SOFTWARE LICENSE AGREEMENT ("EULA")

IMPORTANT NOTICE, PLEASE READ CAREFULLY, THIS END USER LICENSE AGREEMENT ("EULA") IS A LEGAL AGREEMENT BETWEEN YOU (AS AN INDIVIDUAL OR ENTITY, "YOU") AND VEEAM SOFTWARE GROUP GMBH ("VEEAM"), WHICH SETS FORTH THE RIGHTS AND OBLIGATIONS GOVERNING VEEAM SOFTWARE PRODUCTS ("SOFTWARE"), LIMITED SUPPORT SERVICES, AND ASSOCIATED DOCUMENTATION. BY INSTALLING, USING OR OTHERWISE INTERACTING WITH THE SOFTWARE OR DOCUMENTATION, DELIVERING OR RECEIVING SERVICES, YOU AGREE TO BE BOUND BY THE TERMS OF THIS EULA. IF YOU DO NOT AGREE WITH THE TERMS OF THIS EULA, DO NOT USE OR OTHERWISE INTERACT WITH THE SOFTWARE, DOCUMENTATION OR SERVICES.

1.0 License Grant.
This EULA grants You a non-exclusive, non-transferable, non-sublicensable right to install and use the Software, in object code form, and any related documentation ("Documentation") for Your internal business purposes under the terms and conditions stated herein.

Permissible Uses & Capacity.

1.1 The Software is to be used in accordance with the specific license You purchased; a description of the license can be found at <https://www.veeam.com/licensing-policy.html> ("Licensing Policy"). You may only use the number of licenses or capacity that You have purchased unless the product specifically allows to exceed usage by certain amount. In the event You exceed the purchased capacity, the Software may not process additional workloads, and Veeam is not required to provide maintenance or support for such excess use, unless You purchase additional licenses.

1.2 You are permitted to make copies of the Software and Documentation for Your own use in accordance with this EULA and the Licensing Policy.

I agree and consent to the terms of Veeam License Agreement and [licensing policy](#)

I agree and consent to each of the license agreements of the [3rd party components](#)

I agree and consent to each of the license agreements of the [required 3rd party software](#)

Next

Step 2. Choose Configuration Mode

At the **Configuration Mode** step of the wizard, choose whether you want to instruct Veeam Backup for AWS to automatically create IAM roles required for the backup appliance configuration, or you want to specify an IAM role created manually.

IMPORTANT

Consider the following:

- If you select the *Automatic* configuration mode, Veeam Backup for AWS will create 2 IAM roles with wide scopes of permissions and capabilities. You can limit permissions assigned to the IAM roles later, or remove the roles and replace them with custom IAM roles created manually.
- If you select the *Manual* configuration mode, make sure you have created the required IAM role beforehand as described in section [Required IAM Permissions](#).

The screenshot shows the 'Initial Configuration' wizard for Veeam Backup for AWS. The 'Configuration Mode' step is active, showing two options: 'Automatic (recommended)' and 'Manual'. The 'Automatic' option is selected. The 'Manual' option includes a link to the 'User Guide'.

Veeam Backup for AWS

Initial Configuration

License Agreement

Configuration Mode

Account

Default User

License

Time Zone

Summary

Choose configuration mode

Choose the configuration mode that will be used to set up the Veeam Backup for AWS appliance.

Automatic (recommended)

Set up the appliance by providing temporary access keys. With this option selected, Veeam Backup for AWS will automatically create required IAM roles and a lifecycle policy used to protect the appliance data.

Manual

Set up the appliance by providing an IAM role. With this option selected, you must manually create the required IAM roles beforehand as described in the [User Guide](#).

Previous Next

Step 3. Specify IAM Identity

At the **Account** step of the wizard, do the following:

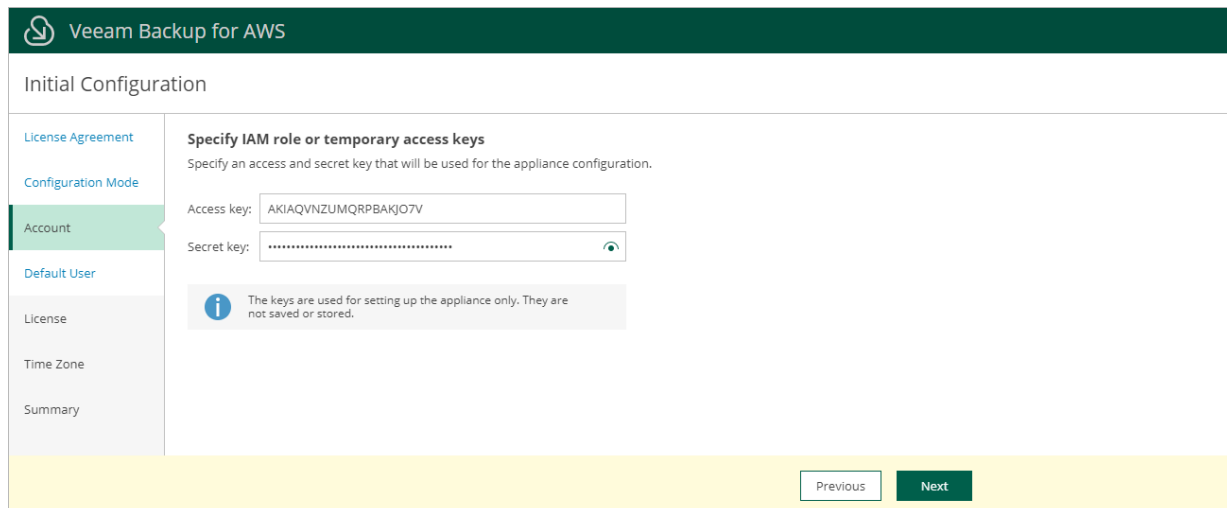
- If you have selected the *Automatic* option at the **Configuration Mode** step of the wizard, specify one-time access keys that will be used to create the *Impersonation* and *Default Backup Restore* IAM roles. For more information on the IAM roles, see [Required IAM Permissions](#).
- If you have selected the *Manual* option at the **Configuration Mode** step of the wizard, specify the *Default Backup Restore* IAM role that will be added to the Veeam Backup for AWS and used to perform operations.

Specifying One-Time Access Keys

To specify the access key ID and the secret access key of an IAM user, use the **Access key** and **Secret key** fields. Note that the IAM user must be authorized to create IAM roles. To learn what permissions the IAM user must have to create IAM roles, see [Required IAM Permissions](#).

NOTE

Veeam Backup for AWS does not store one-time access keys in the configuration database.



The screenshot shows the 'Initial Configuration' wizard for Veeam Backup for AWS. The 'Account' step is selected in the left-hand navigation pane. The main content area is titled 'Specify IAM role or temporary access keys' and includes the instruction: 'Specify an access and secret key that will be used for the appliance configuration.' There are two input fields: 'Access key:' with the value 'AKIAQVNZUMQRPBAKJO7V' and 'Secret key:' with a masked value. An information icon and text state: 'The keys are used for setting up the appliance only. They are not saved or stored.' At the bottom right, there are 'Previous' and 'Next' buttons.

Specifying IAM Role

To specify the *Default Backup Restore* IAM role, enter the IAM role name specified in AWS when creating the role. The IAM role must be created beforehand as described in section [Required IAM Permissions](#).

NOTE

If there is a path identifying the IAM role, you must specify the role name in the `PATH/NAME` format (for example, `dept_1/s3_role`). To learn how to add identifiers to IAM roles, see [AWS Documentation](#).

You can check whether the specified IAM role has permissions required to perform all Veeam Backup for AWS operations. To run the IAM role permission check, click **Check Permissions**. If some permissions of the IAM role are missing, Veeam Backup for AWS will display a warning, but you will still be able to proceed with the wizard without granting the missing permissions to the role. To learn how to grant permissions to IAM roles using the AWS Management Console, see [AWS Documentation](#).

TIP

You can grant permissions to this IAM role and add other IAM roles that will be used to perform backup and restore operations later, after the backup appliance configuration completes. For more information, see [Managing IAM Roles](#).

Veeam Backup for AWS

Initial Configuration

- License Agreement
- Configuration Mode
- Account**
- Default User
- License
- Time Zone
- Summary

Specify IAM role or temporary access keys

Specify an IAM role that will be added to Veeam Backup for AWS and used as the default role for operations. For more information on the required permissions, see the [User Guide](#).

IAM role:

[Check Permissions](#)

Warning: The specified IAM role is missing permissions. Use the JSON file below to assign all the required permissions to the role.

Required Permissions

Veeam Backup for AWS automatically creates a JSON file with the list of all required permissions. Add the permissions to an IAM policy attached to the specified IAM role and run the permission check again.

[Download](#)

[Previous](#) [Next](#)

Step 4. Create Default User

At the **Default User** step of the wizard, create the default user whose credentials you will use for your [first login to Veeam Backup for AWS](#).

Note that the specified user name cannot be *admin*, can contain only lowercase Latin letters, numeric characters, underscores and dashes. You can use the dollar sign (\$) as the last character of the name. The maximum length of the name is 32 characters.

NOTE

To increase the security of the default user, it is recommended that you enable multi-factor authentication (MFA) for the user account after you first log in to Veeam Backup for AWS. To learn how to enable MFA, see [Configuring Multi-Factor Authentication](#).

The screenshot shows the 'Initial Configuration' wizard for Veeam Backup for AWS. The 'Default User' step is selected in the left-hand navigation pane. The main content area is titled 'Specify credentials for default user' and includes the following fields and instructions:

- Name:** administrator
- Password:** [Redacted] (with a visibility toggle icon)
- Repeat password:** [Redacted] (with a visibility toggle icon)

Instructions for the password: Password must be 8 characters minimum with one digit, one uppercase and one lowercase. Monotonic sequences such as 1234 are not allowed.

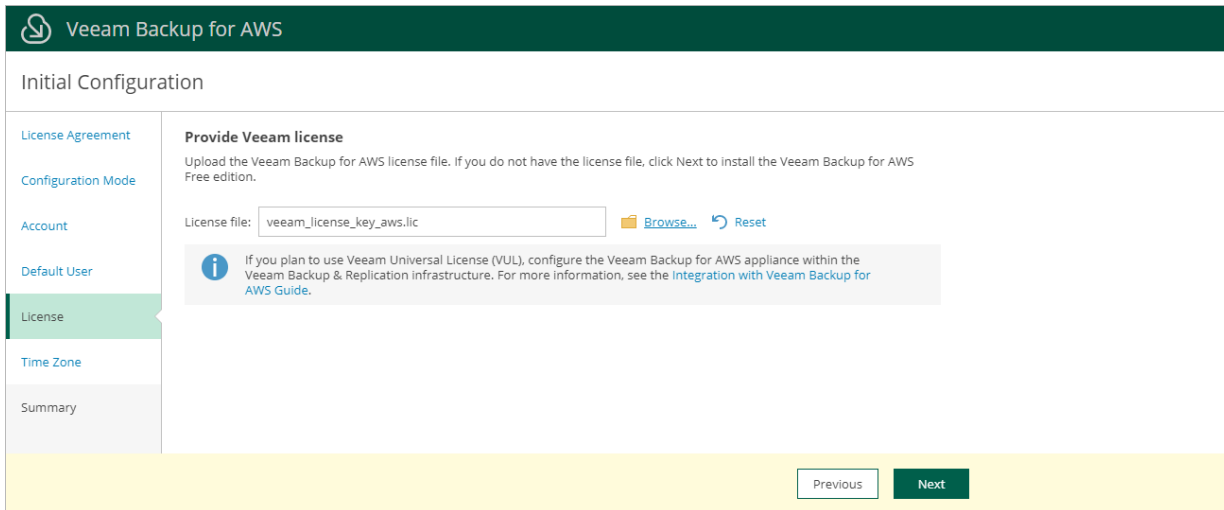
At the bottom right of the wizard, there are 'Previous' and 'Next' navigation buttons.

Step 5. Install License

At the **License** step of the wizard, browse to the license file supplied to you by Veeam. You will still be able to proceed with the wizard without providing a license – in this case, the *Free edition* of Veeam Backup for AWS will be installed.

TIP

You can install a valid license later, after the backup appliance configuration completes. For more information, see [Installing and Removing License](#).



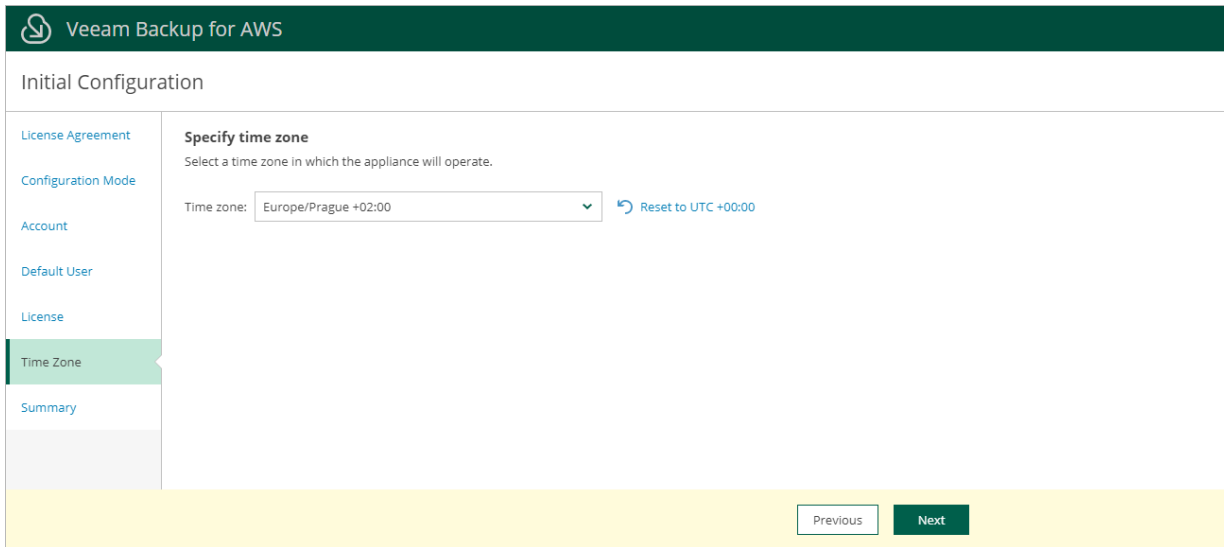
The screenshot shows the 'Initial Configuration' wizard for Veeam Backup for AWS, specifically the 'License' step. The left sidebar contains a navigation menu with the following items: License Agreement, Configuration Mode, Account, Default User, License (highlighted), Time Zone, and Summary. The main content area is titled 'Provide Veeam license' and includes the following text: 'Upload the Veeam Backup for AWS license file. If you do not have the license file, click Next to install the Veeam Backup for AWS Free edition.' Below this text is a 'License file:' label followed by a text input field containing 'veeam_license_key_aws.lic', a 'Browse...' button with a folder icon, and a 'Reset' button with a circular arrow icon. An information icon (i) is followed by a note: 'If you plan to use Veeam Universal License (VUL), configure the Veeam Backup for AWS appliance within the Veeam Backup & Replication infrastructure. For more information, see the [Integration with Veeam Backup for AWS Guide](#).' At the bottom right of the wizard, there are two buttons: 'Previous' and 'Next'.

Step 6. Specify Time Zone

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 at the **Time Zone** step of the wizard if required. For example, you may want the time on the backup appliance to match the time on the local machine from which you access Veeam Backup for AWS.

TIP

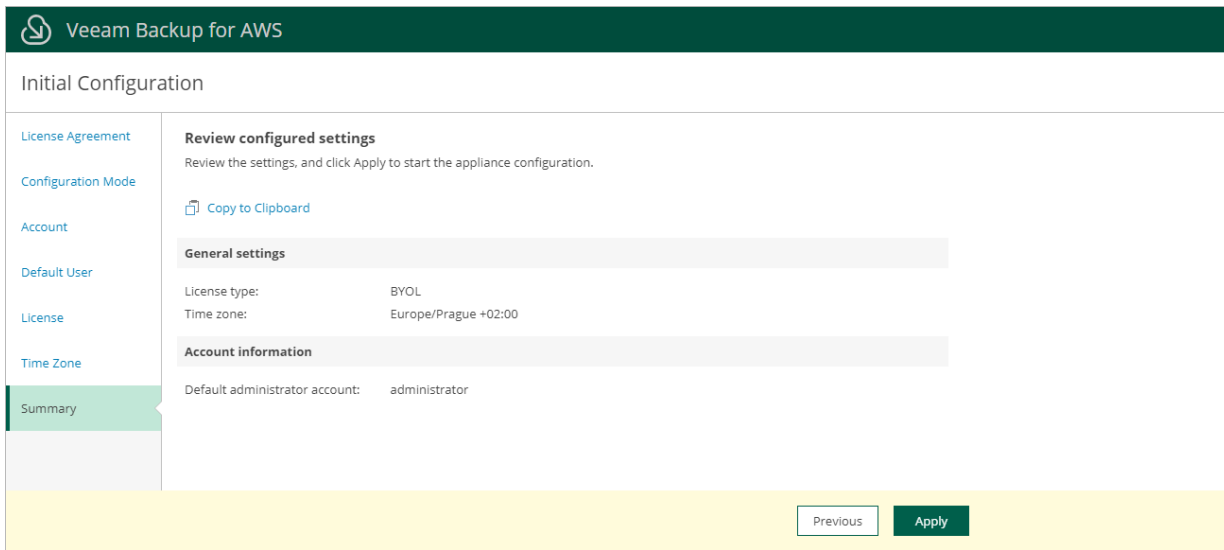
You can change time zone settings later, after the backup appliance configuration completes. For more information, see [Changing Time Zone](#).



The screenshot shows the 'Initial Configuration' wizard for Veeam Backup for AWS. The 'Time Zone' step is highlighted in the left sidebar. The main content area is titled 'Specify time zone' and includes the instruction 'Select a time zone in which the appliance will operate.' Below this, there is a 'Time zone:' label followed by a dropdown menu currently set to 'Europe/Prague +02:00'. To the right of the dropdown is a 'Reset to UTC +00:00' link with a circular arrow icon. At the bottom of the wizard, there are 'Previous' and 'Next' buttons.

Step 7. Finish Working with Wizard

At the **Summary** step of the wizard, review configuration settings and click **Finish**. After the initial configuration process completes, Veeam Backup for AWS will display the [welcome screen where you can log in](#).



Uninstalling Veeam Backup for AWS

Depending on the installation option you chose to deploy Veeam Backup for AWS, 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](#).

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**.

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

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](#).

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**.

Failure and Recovery

Even though by default, the EC2 Instance on which Veeam Backup for AWS is installed is backed up using snapshots, it is recommended that you regularly perform backup of the EC2 instance that hosts the product and perform backup of the configuration database that stores data collected from Veeam Backup for AWS. If the Veeam Backup for AWS fails for some reason, you can restore it from the backup. To learn how to perform configuration backup for backup appliances, see [Performing Configuration Backup and Restore](#).

For application related errors and issues, see the [Veeam Knowledge Base](#) or consider opening a [support ticket](#).

Licensing

This section describes how the solution is licensed, how to manage license workloads, and what licensing limitations and scenarios can apply.

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 of Managed Backup Appliances

If a backup appliance is managed by a Veeam Backup & Replication server, Veeam Backup for AWS uses the same license that is installed on the backup server. For more information on the Veeam Backup & Replication licensing, see the Veeam Backup & Replication User Guide, section [Licensing](#).

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 details 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](#).

Licensing Scenarios

When you add a backup appliance to the backup infrastructure, the following scenarios are applied:

- If you [connect to an existing backup appliance](#), the [BYOL license](#) installed on the appliance becomes invalid. 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.

- 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 backup 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. To back up more than 10 workloads, you must install a *BYOL* license on the backup appliance. To see how to install a new *BYOL* license, see [Installing and Removing License](#).

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.

Licensing of Standalone Backup Appliances

Veeam Backup for AWS is licensed per protected instance. An instance is defined as a single AWS resource – EC2 instance, RDS resource, DynamoDB table or EFS file system. 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.

NOTE

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](#).

Product Editions

Veeam Backup for AWS is available in 3 editions:

- **Free**

Veeam Backup for AWS operating in the *Free* edition allows you to protect up to 10 instances free of charge. Note that this edition does not support indexing of EFS file systems.

- **Paid**

Veeam Backup for AWS operating in the *Paid* edition allows you to protect an unlimited number of instances.

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](#).

- **BYOL (Bring Your Own License)**

Veeam Backup for AWS operating in the *BYOL* edition allows you to protect the number of instances equivalent to the number of units specified in your license.

Veeam Backup for AWS *BYOL* edition can be licensed using either the Veeam Universal License (VUL) or a separate product license that can be obtained by contacting a Veeam sales representative at [Sales Inquiry](#).

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.

For details on how to install the license on the backup appliance, see [Installing and Removing License](#).

Installing and Removing License

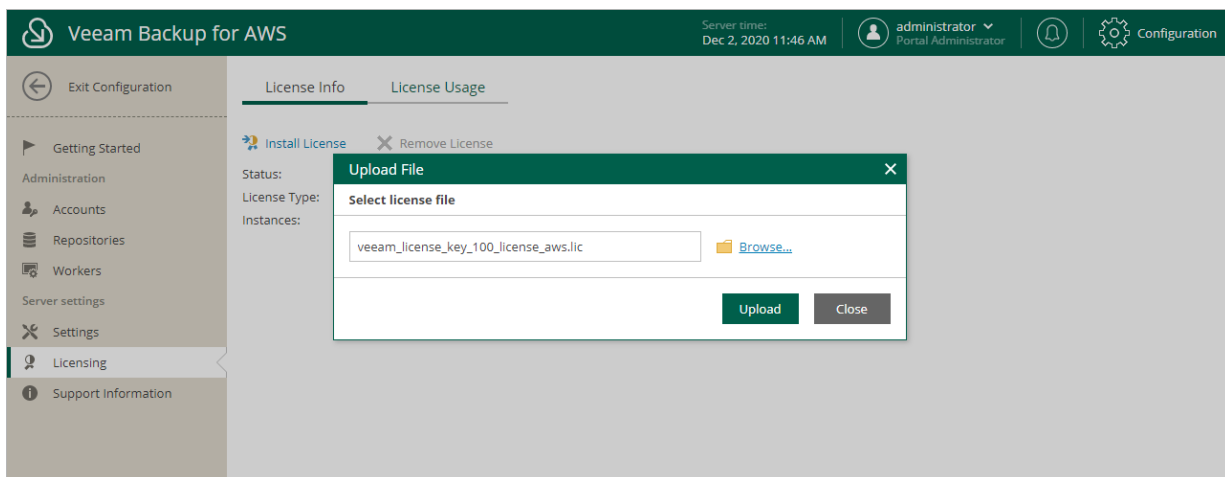
NOTE

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:

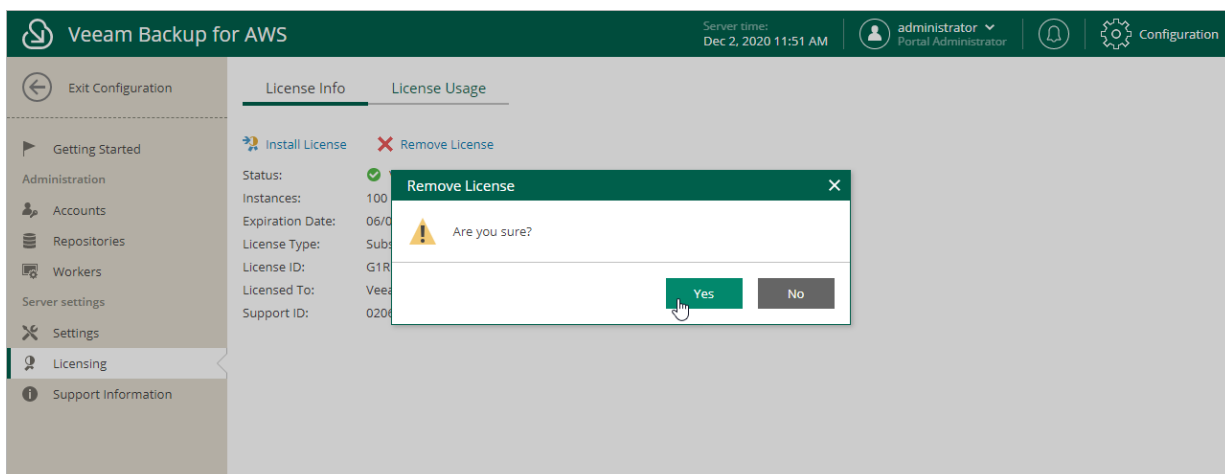
1. Switch to the **Configuration** page.
2. Navigate to **Licensing > License Info**.
3. Click **Install License**.
4. In the **Upload file** window, click **Browse** to browse to a license file, and then click **Upload**.



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](#).

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**.

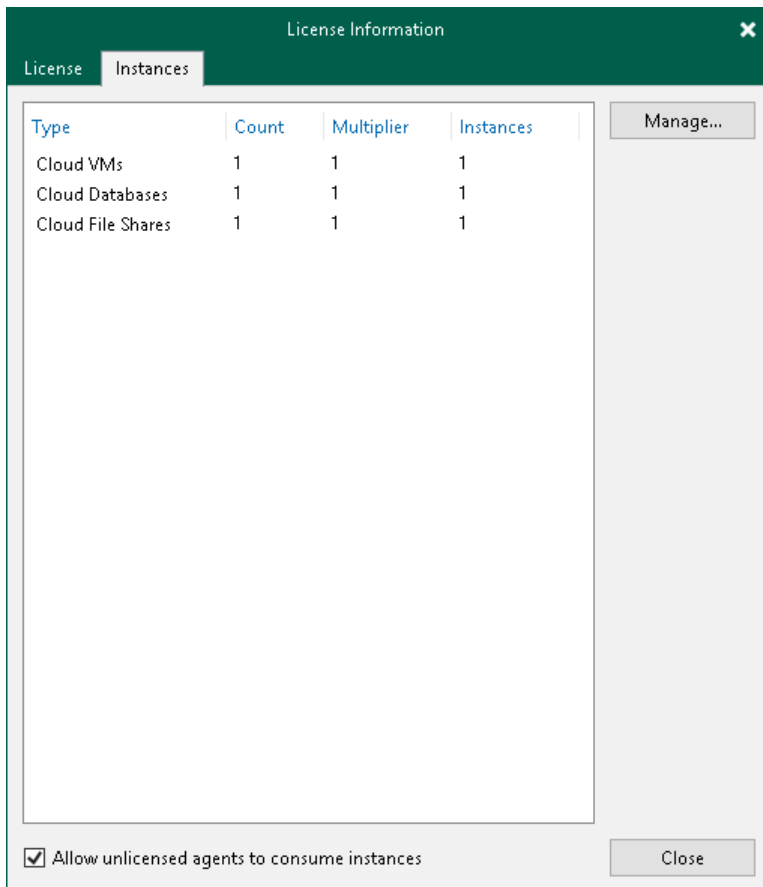
The **License** tab of the **License Information** window provides general information on the currently installed AWS Plug-in for Veeam Backup & Replication license:

- **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 shows the Veeam Backup for AWS configuration interface. 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 titled "License Info" and "License Usage". It features two buttons: "Install License" (with a plus icon) and "Remove License" (with a minus icon). Below these buttons, the following license details are displayed:

- 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

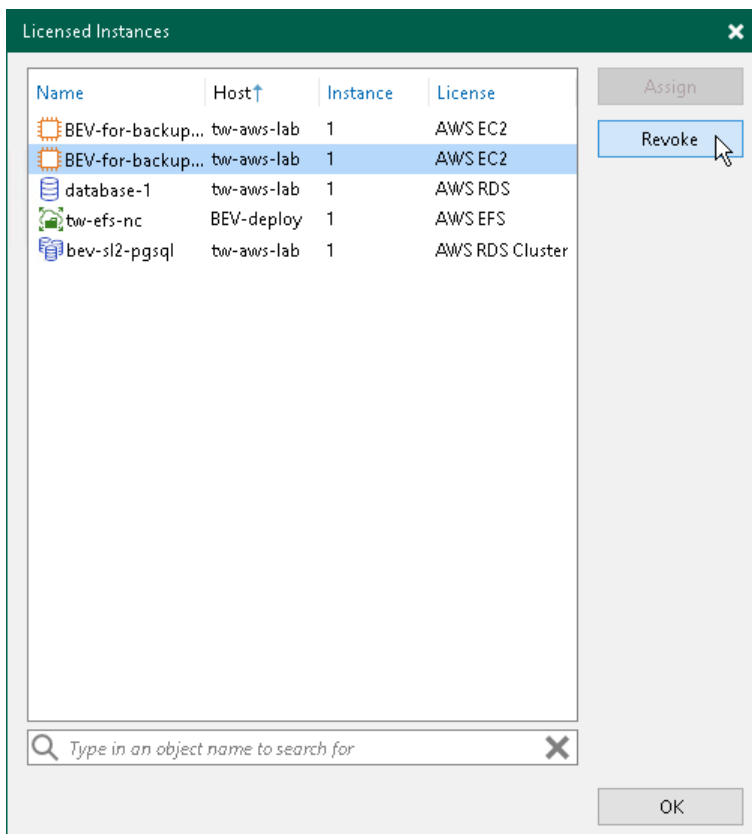
Revoking License Units

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.



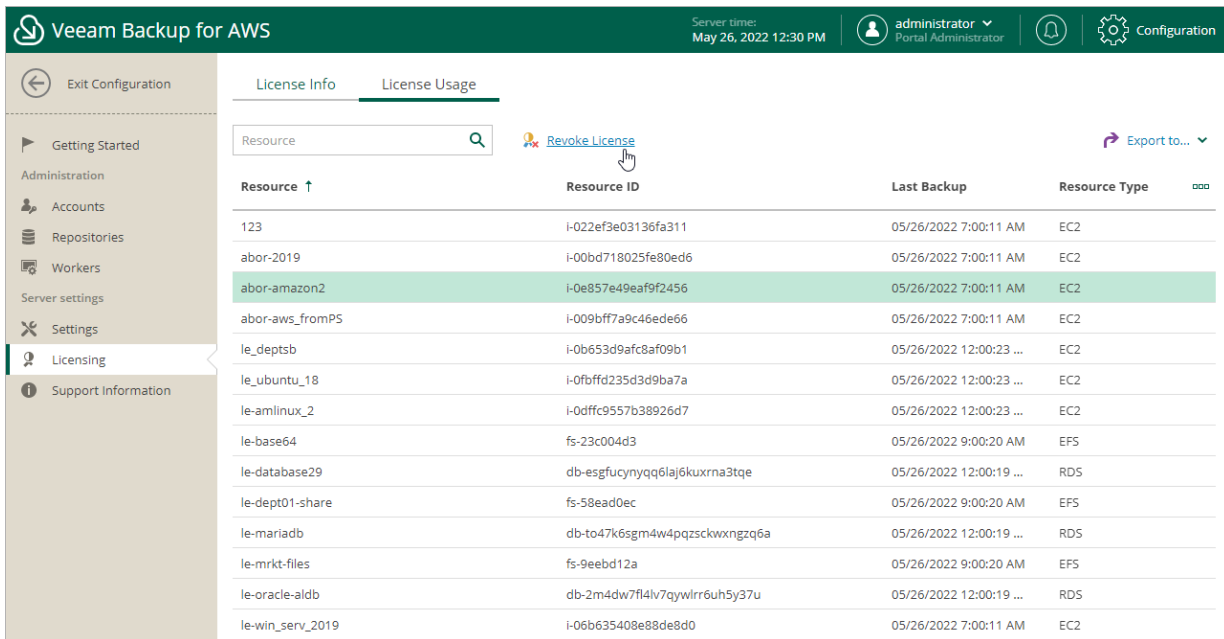
Revoking License Units Using Veeam Backup for AWS Web UI

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. 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 'License Info' and 'License Usage' tabs. The 'License Usage' tab is active, showing a search bar, a 'Revoke License' button, and an 'Export to...' dropdown. Below is a table with columns: 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](#).

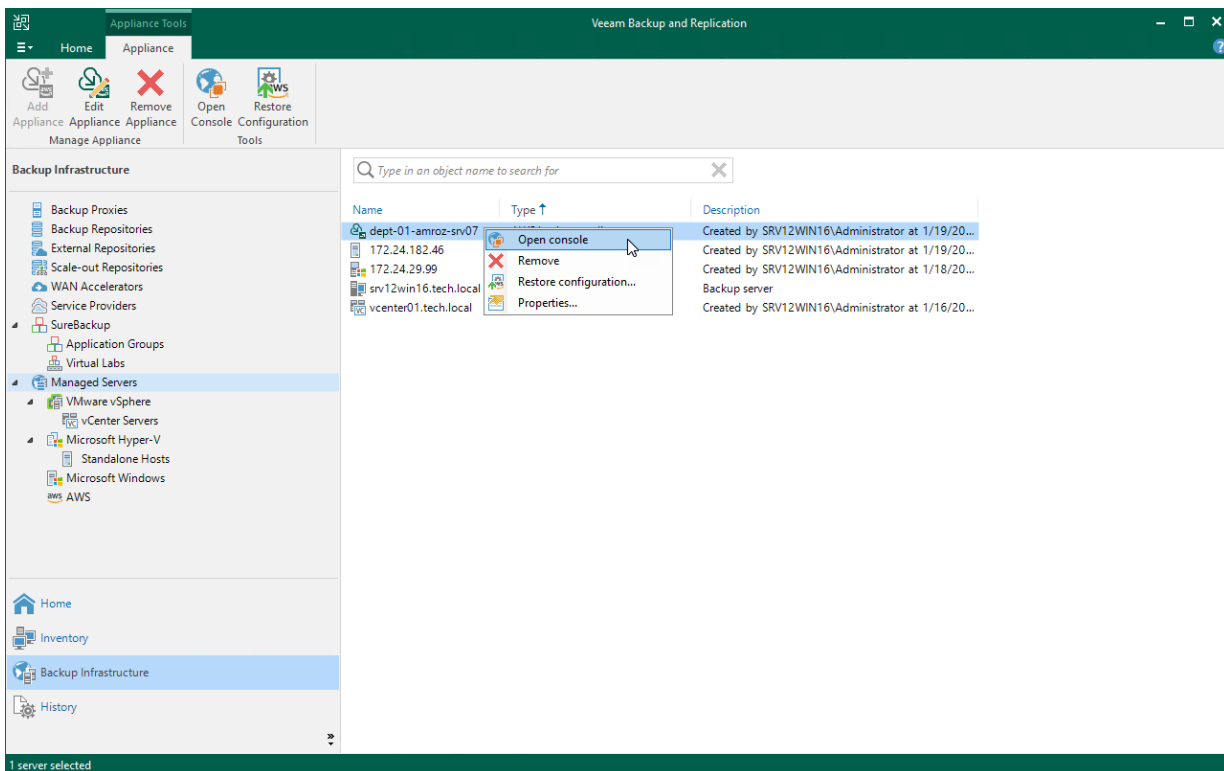
Accessing Web UI from Console

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 device 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**.



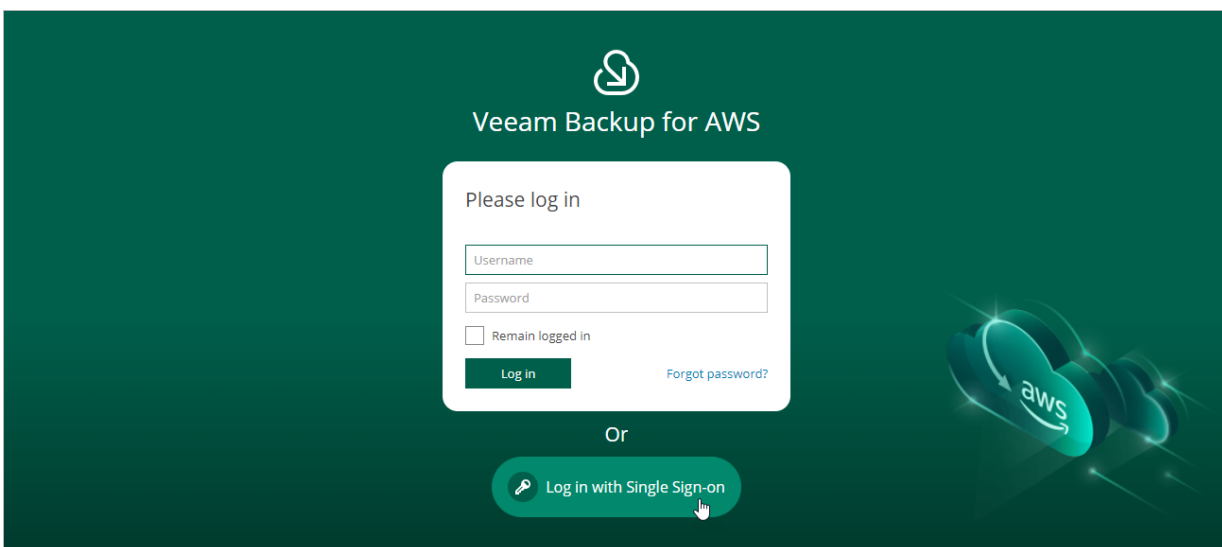
Logging In Using Identity Provider User Account

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. [Add users to control access to Veeam Backup for AWS.](#)
 - c. [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 launched.

- d. [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 tab](#) – 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 as described in section [Deploying Backup Appliance](#).

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

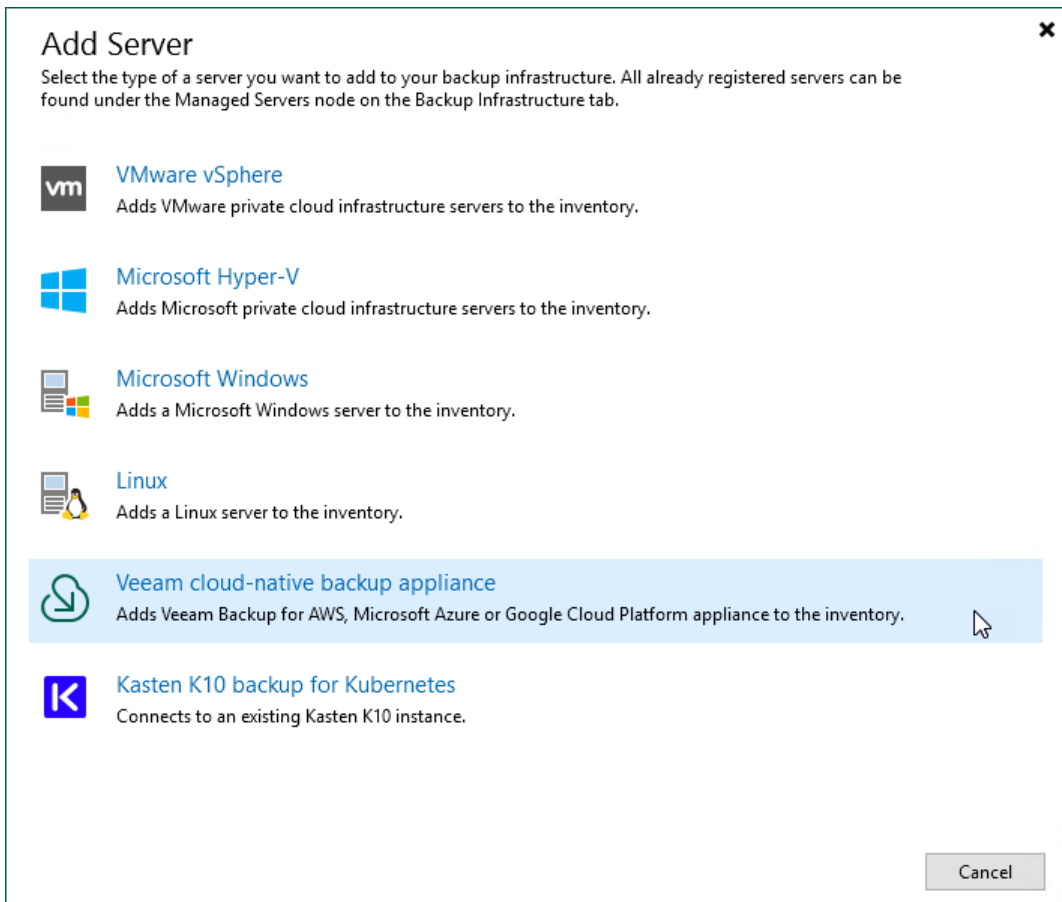
If you have already [deployed a backup appliance](#), you can add the appliance to the backup infrastructure:

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

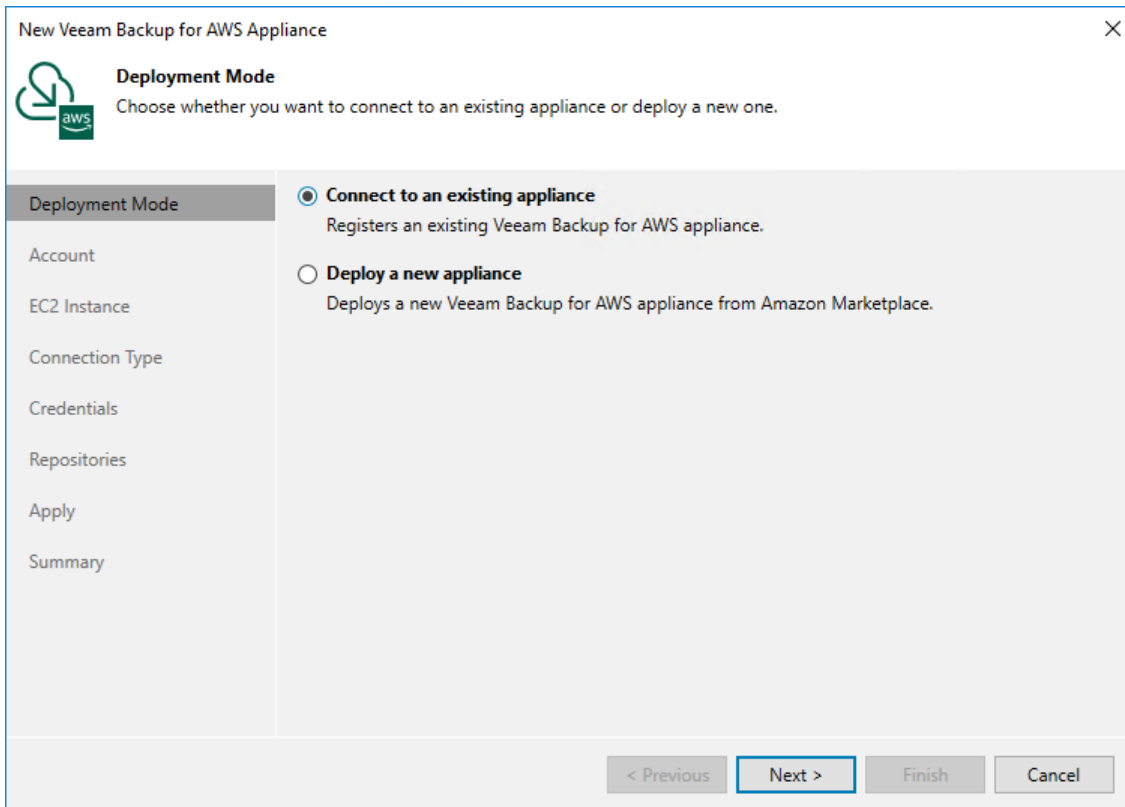
To launch the **New Veeam Backup for AWS Appliance** wizard, do one of the following:

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 (highlighted), 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, there 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. 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-down list, specify whether the backup appliance resides in the AWS Global or AWS GovCloud (US) region.

IMPORTANT

To check region availability, Veeam Backup & Replication establishes a temporary test connection to 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.

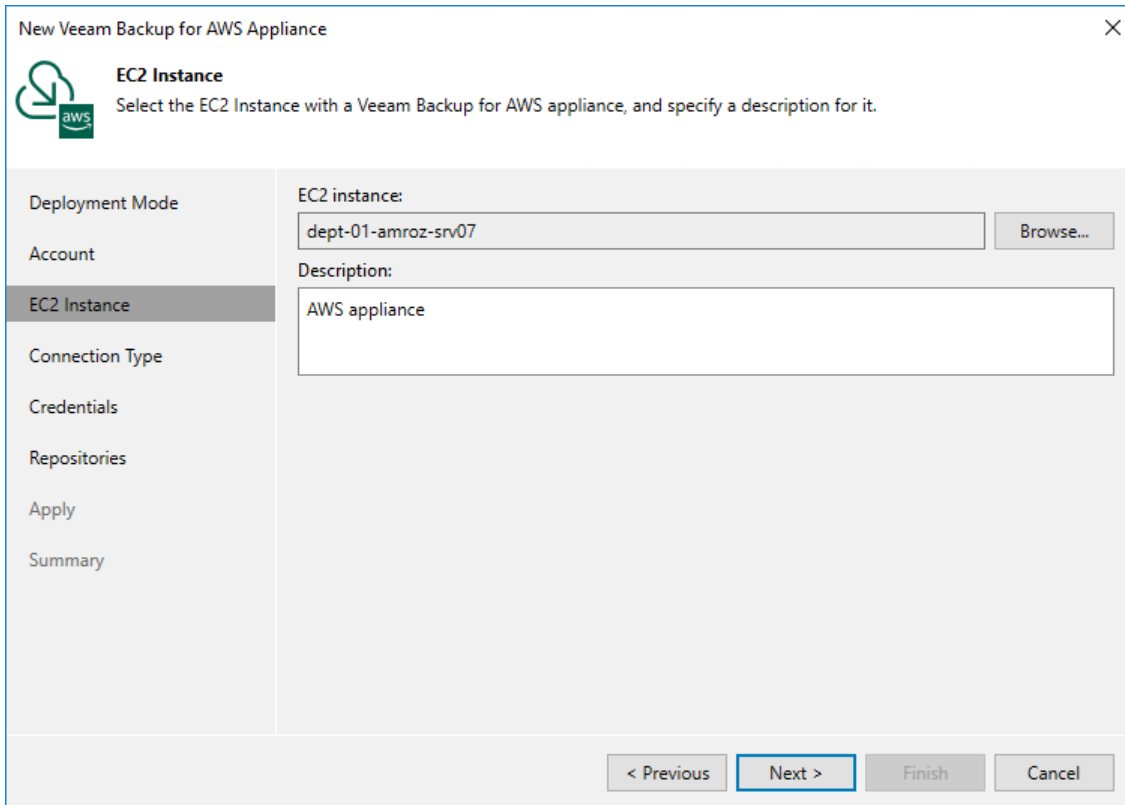
3. From the **Data center** drop-down list, select the AWS Region in which the backup appliance resides. For more information on regions and availability zones, see [AWS Documentation](#).

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' and it has a close button (X) in the top right corner. The main heading is 'Account' with the subtitle '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 'XXXXXXXXXXXXXXXXX (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 contains the following fields and instructions:

- EC2 Instance**: Select the EC2 Instance with a Veeam Backup for AWS appliance, and specify a description for it.
- EC2 instance:** A text input field containing "dept-01-amroz-srv07" and a "Browse..." button to its right.
- Description:** A text input field containing "AWS appliance".

At the bottom of the window, there are four navigation 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 contains 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 includes the instruction "Specify if the Veeam Backup for AWS appliance is connected to the Internet." Below this, 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, there are four buttons: "< Previous", "Next >" (highlighted with a blue border), "Finish", and "Cancel".

Step 6. Specify User Credentials

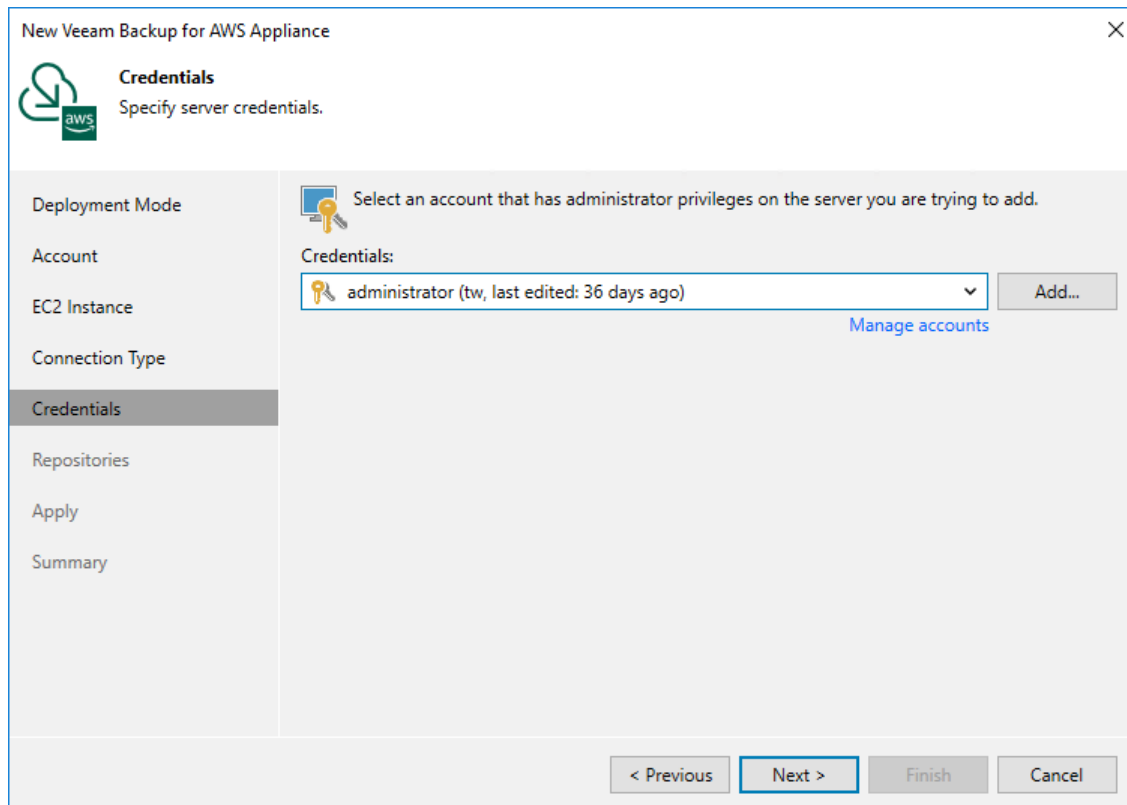
At the **Credentials** step of the wizard, specify a user whose credentials Veeam Backup & Replication will use to connect to the backup appliance.

For a user to be displayed in the **Credentials** list, it must be added to the Credentials Manager. If you have not added a user to the Credentials Manager beforehand, you can do it without closing the **New Veeam Backup for AWS Appliance** wizard. To add a new user, click either the **Manage cloud accounts** link or the **Add** button and specify a 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' with a close button (X) in the top right corner. The 'Credentials' step is highlighted in the left sidebar, which also includes 'Deployment Mode', 'Account', 'EC2 Instance', 'Connection Type', 'Repositories', 'Apply', and 'Summary'. The main content area has a heading 'Credentials' with the instruction 'Specify server credentials.' Below this, there is a key icon and the text 'Select an account that has administrator privileges on the server you are trying to add.' A dropdown menu labeled 'Credentials:' shows 'administrator (tw, last edited: 36 days ago)' selected. 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 >', 'Finish', and 'Cancel'. The 'Next >' button is highlighted with a blue border.

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 7.0](#), Veeam Backup & Replication will verify whether the IAM user whose one-time 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 permissions 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 details 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 one-time 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 one-time access keys of an IAM user to be displayed in the **Credentials** list, they must be added to the Cloud Credentials Manager. If you have not added the 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 one-time 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 backup infrastructure 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 add the password, click either the **Manage cloud accounts** link or the **Add** button and specify a hint and the password 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](#).

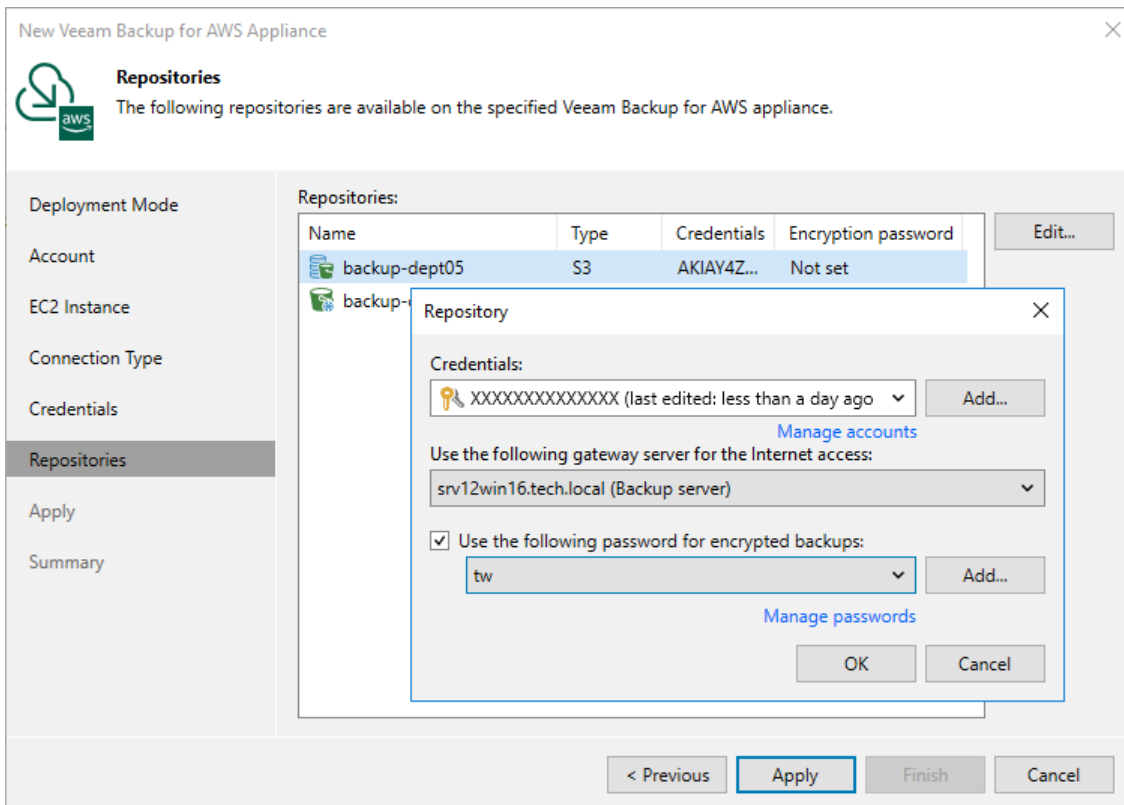
- If data in the repository is encrypted with a KMS key, Veeam Backup & Replication will show the used KMS key in the **Perform AWS encryption with the following KMS key** drop-down list but will not allow to change it.

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**.

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 [Configuration](#). If you want Veeam Backup & Replication to open the Web UI of the added backup appliance immediately, click the **backup appliance console** link.

The screenshot shows the 'New Veeam Backup for AWS Appliance' wizard at the 'Summary' step. The window title is 'New Veeam Backup for AWS Appliance' with a close button (X) in the top right corner. On the left, there is a navigation pane with the following items: Deployment Mode, Account, EC2 Instance, Connection Type, Credentials, Repositories, Apply, and Summary (which is highlighted). The main content area is titled 'Summary' and contains the text: 'You can copy the configuration information below for future reference.' Below this text is a scrollable box containing the following configuration details:

```
Summary:
New backup appliance has been registered successfully.
Account options:
  AWS account: XXXXXXXXXXXXXXXXXXXX
  AWS region: Global
  Data center: EU (Paris) (eu-west-3)
EC2 instance options:
  EC2 instance name: dept-01-amroz-srv07
  Connection type: Direct connection
  Guest OS credentials: administrator
Repositories:
  Name: backup-dept05
  Type: S3
  Credentials: XXXXXXXXXXXXXXXXXXXX
  Gateway server: srv12win16.tech.local (Backup server)

  Name: backup-dept06
  Type: S3 Glacier
  Credentials: N/A
```

Below the scrollable box, there is a link: 'Open [backup appliance console](#) to configure advanced settings'. At the bottom of the window, there are four buttons: '< Previous', 'Next >', 'Finish' (which is highlighted with a blue border), and 'Cancel'.

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:

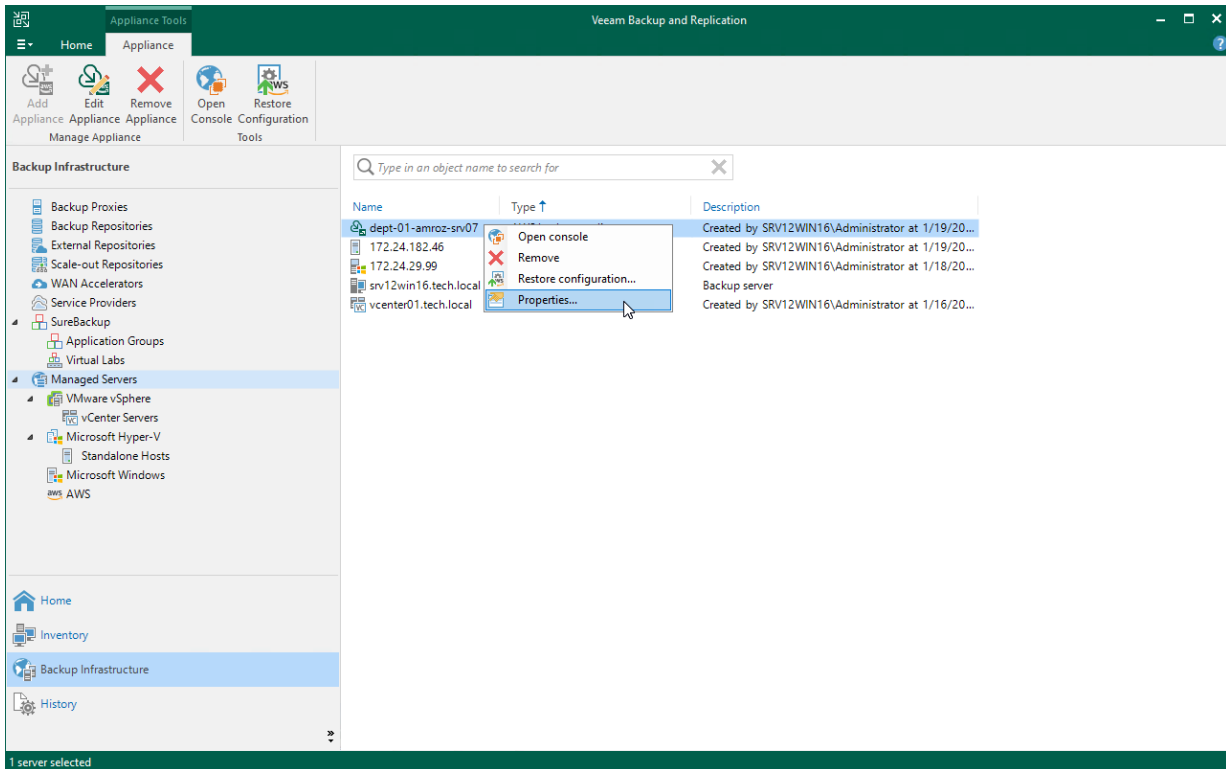
1. In the Veeam Backup & Replication console, open the **Backup Infrastructure** view.
2. Navigate to **Managed Servers**.
3. Select the necessary backup appliance and click **Edit Appliance** on the ribbon.
Alternatively, you can right-click the appliance and select **Properties**.
4. Complete the **Edit 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 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).
- f. At the **Apply** step of the wizard, wait for the changes to be applied and click **Next**.

g. At the **Summary** step of the wizard, review summary information and click **Finish**.



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.

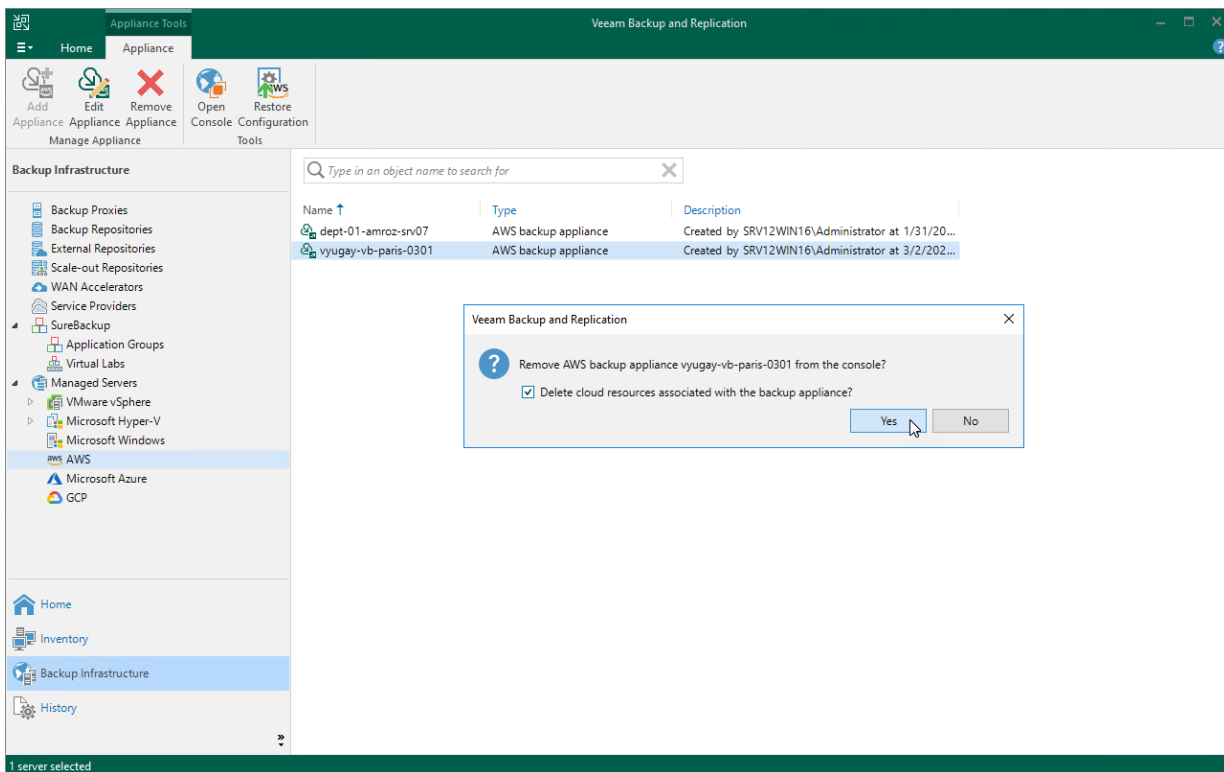
To remove a backup appliance, do the following:

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 [Uninstalling Veeam Backup for AWS](#).



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 launch 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, unless you deployed the product from an AMI and manually [assigned the role a minimum set of permissions](#).

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:

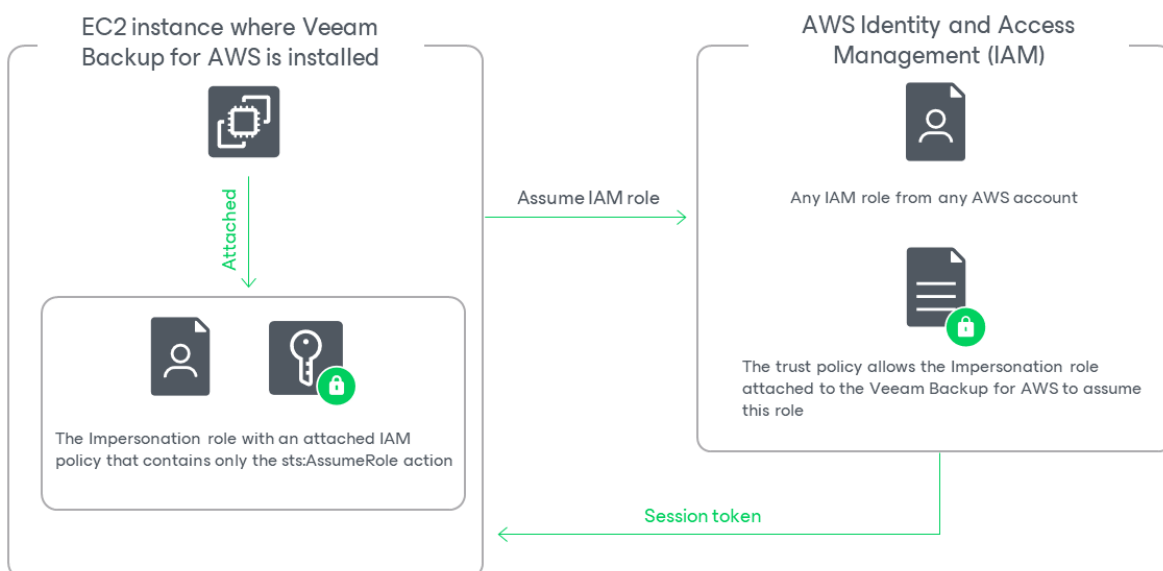
1. [Launch the Add IAM Role wizard.](#)
2. [Specify a name and description for the IAM role.](#)
3. [Specify IAM role settings.](#)
4. [Specify IAM role permissions.](#)
5. [Finish working with the wizard.](#)

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 all the permissions required to assume these roles.

IMPORTANT

The only exception to this behavior is the scenario where you [deploy the backup appliance from AMI](#) and perform the appliance configuration using the **Manual** configuration mode. In this case, Veeam Backup for AWS does not create the *Impersonation* IAM role automatically, and you must create it after the deployment manually as described in section [Required IAM Permissions](#).



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 **Roles** 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:

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>"
      }
    }
  ]
}
```

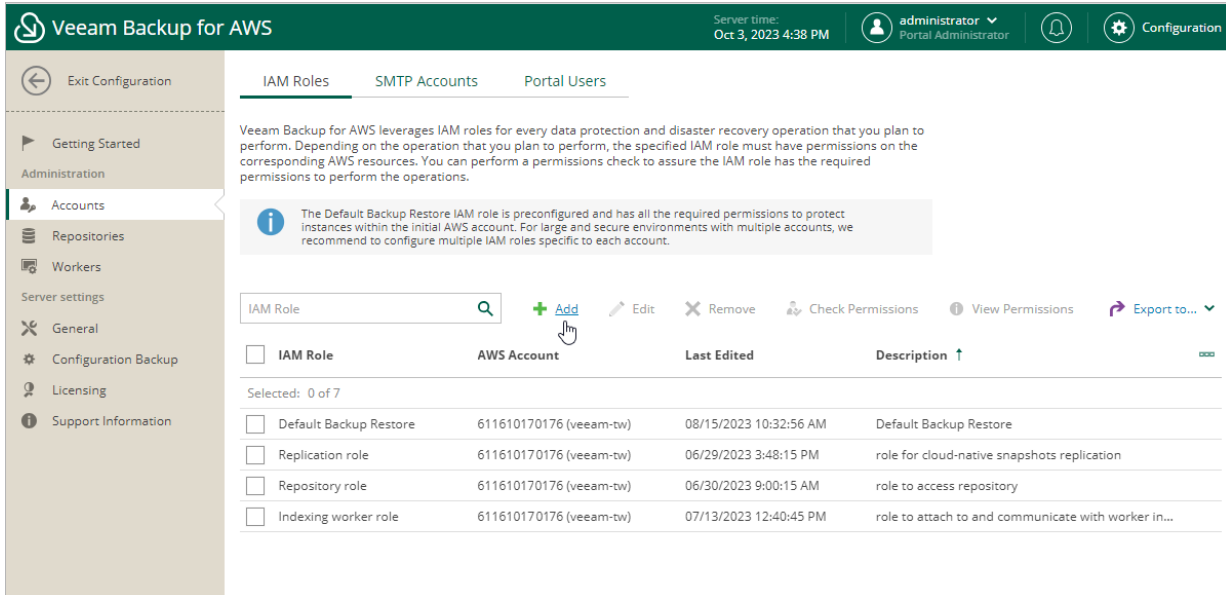
Where `<Role ARN>` is the ARN either of the *Impersonation* IAM role that you have copied at step 2a, or the ARN of the AWS account to which the backup appliance belongs.

- e. Click **Update policy**. Note that it may take up to 5 minutes for AWS to update the trust policy.

Step 1. Launch Add Account 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**.



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 titled 'Add IAM Role' and is currently on the 'Info' step. The main content area is titled 'Specify IAM role name and description' and contains the instruction 'Enter a name and description for the IAM role.' 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'. The left sidebar shows a navigation menu with 'Info' selected, and other options like 'Type', 'Permissions', and 'Summary'. The top navigation bar includes the Veeam logo, 'Veeam Backup for AWS', server time 'Oct 3, 2023 4:42 PM', user 'administrator Portal Administrator', and a 'Configuration' link. At the bottom right, 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.

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 selected in the left-hand navigation pane. The main content area is titled 'Specify IAM role type and settings' and includes a sub-header 'Select which type of IAM role to use and specify the settings for this role. For more information on IAM roles, see the [User Guide](#).' There are four radio button options: 'IAM role from current account' (which is selected), 'IAM role from another account', 'Create new IAM role', and 'Create template'. Below the first option, there is a text input field for 'AWS role name' containing the value 'production_worker_role'. At the bottom of the wizard, there are three buttons: 'Previous', 'Next', and 'Cancel'.

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.

- [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 selected in the left sidebar. The main content area is titled 'Specify IAM role type and settings' and includes the instruction: 'Select which type of IAM role to use and specify the settings for this role. For more information on IAM roles, see the User Guide.' There are four radio button options: 'IAM role from current account', 'IAM role from another account' (which is selected), 'Create new IAM role', and 'Create template'. The 'IAM role from another account' option has three input fields: 'Account ID' (3945681109), 'AWS role name' (production_worker_role), and 'External ID (optional)' (ad395958dept01). Each field has an information icon. At the bottom, 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:

- 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](#).

- 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 contains 'production_worker_role'. The 'Access key' field contains 'AKIAY4ZWOU4WMVRAGEVN'. The 'Secret key' field is masked with dots. A 'Previous' button is highlighted in yellow at the bottom.

Specifying Settings for Template

[This step applies only if you have selected the **Create template** option]

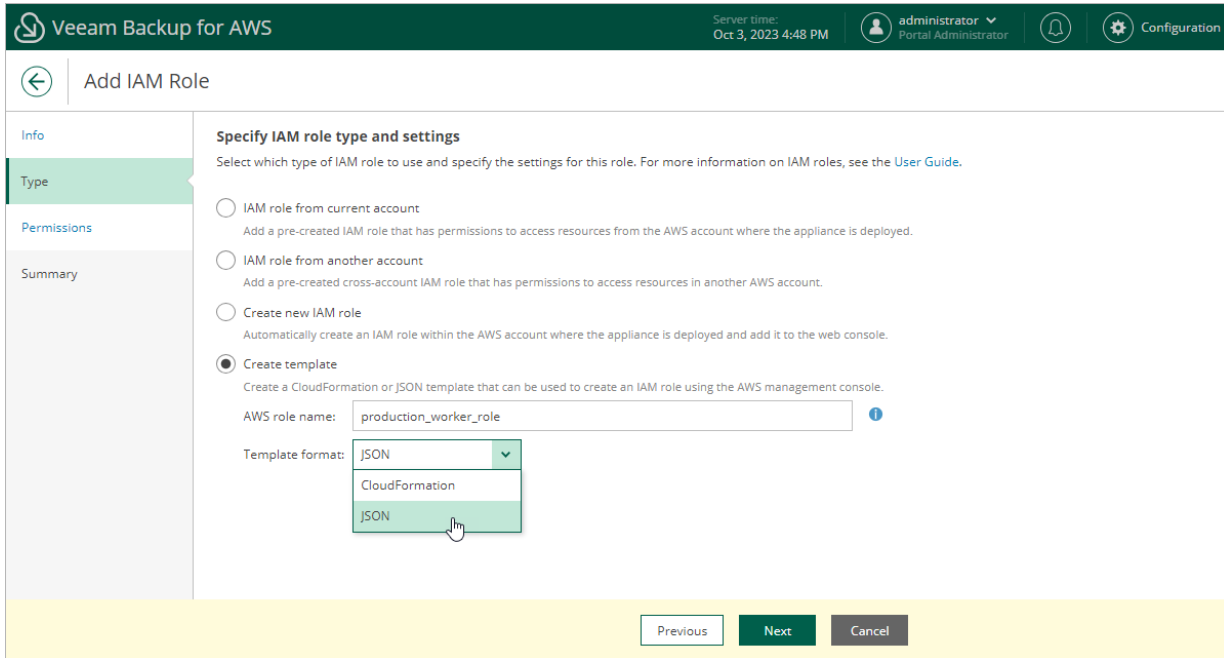
At the **Type** step of the wizard, specify the following settings:

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 launch 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 [Service IAM Role in Backup Account](#), [Service IAM Roles 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, DynamoDB, EFS and VPC resources. If you select this operation for an IAM role, you will be able to select it in the [EC2 backup](#), [DynamoDB backup](#), [EFS 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, DynamoDB, EFS 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](#), [DynamoDB restore](#), [EFS 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 shows the 'Add IAM Role' configuration interface in Veeam Backup for AWS. The interface includes a navigation sidebar on the left with options for Info, Type, Permissions (selected), and Summary. The main content area is divided into three panels:

- Specify IAM role permissions:** Contains an information icon and a note: "By default, the IAM role will be assigned to all workloads." Below this is a toggle for "Specify granular permissions:" which is currently turned on. A table lists permissions for various services:

Veeam management permissions	No permissions
Amazon EC2	No permissions
Amazon RDS	No permissions
Amazon EFS	No permissions
Amazon VPC	No permissions
Amazon DynamoDB	No permissions

 An "Edit Permissions" link is located below the table.
- Veeam management roles:** Includes a note about managing backup repositories in AWS S3 or worker instances. It features "Select All", "Clear All", and "Reset" buttons. The "Production worker role" is selected, while "Worker deployment role" and "Repository role" are not.
- Workload permissions:** A section titled "Select the workloads you are planning to protect and what actions this account should be able to perform." It lists several workload categories with expandable options:
 - Amazon EC2:** Backup (checked), Replication (unchecked), Restore (checked).
 - Amazon RDS:** Snapshot (unchecked), Replication (unchecked), Restore (unchecked).
 - Amazon EFS:** Backup (checked), Restore (checked).
 - Amazon VPC:** Backup (unchecked), Restore (unchecked).
 - Amazon DynamoDB:** Backup (unchecked), Restore (unchecked).

At the bottom of the configuration area, there are "Apply" and "Cancel" buttons. The "Apply" button is highlighted in green.

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 Account 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

An information message states: "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." Below this message, the checkbox "Perform permission check when I click Finish" is checked.

At the bottom of the wizard, there are three buttons: "Previous", "Finish", and "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 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](#).
 - c. 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](#).
 - d. At the **Summary** step of the wizard, review summary information and click **Finish**.

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 the permissions available. If you want to provide granular permissions to the roles, follow the instructions provided in section [Adding IAM Roles](#) (step 4).

Veeam Backup for AWS

Server time: Oct 3, 2023 4:52 PM

administrator Portal Administrator

Configuration

Exit Configuration

IAM Roles SMTP Accounts Portal Users

Getting Started

Administration

Accounts

Repositories

Workers

Server settings

General

Configuration Backup

Licensing

Support Information

Veeam Backup for AWS leverages IAM roles for every data protection and disaster recovery operation that you plan to perform. Depending on the operation that you plan to perform, the specified IAM role must have permissions on the corresponding AWS resources. You can perform a permissions check to assure the IAM role has the required permissions to perform the operations.

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.

IAM Role + Add Edit Remove Check Permissions View Permissions Export to...

<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...

Selected: 1 of 7

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 displays the Veeam Backup for AWS configuration interface. The top navigation bar includes 'Exit Configuration', 'IAM Roles', 'SMTP Accounts', and 'Portal Users'. The main content area shows a table of permissions for a repository role, with one entry marked as 'Failed' and others as 'Passed'. A 'Grant Permissions' dialog box is open, prompting for temporary credentials (Access key and Secret key) to grant permissions manually.

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

Grant Permissions

Provide temporary credentials

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: AKIAY4ZWOU4WMVRAGEVN

Secret key:

Apply Cancel

Removing IAM Roles

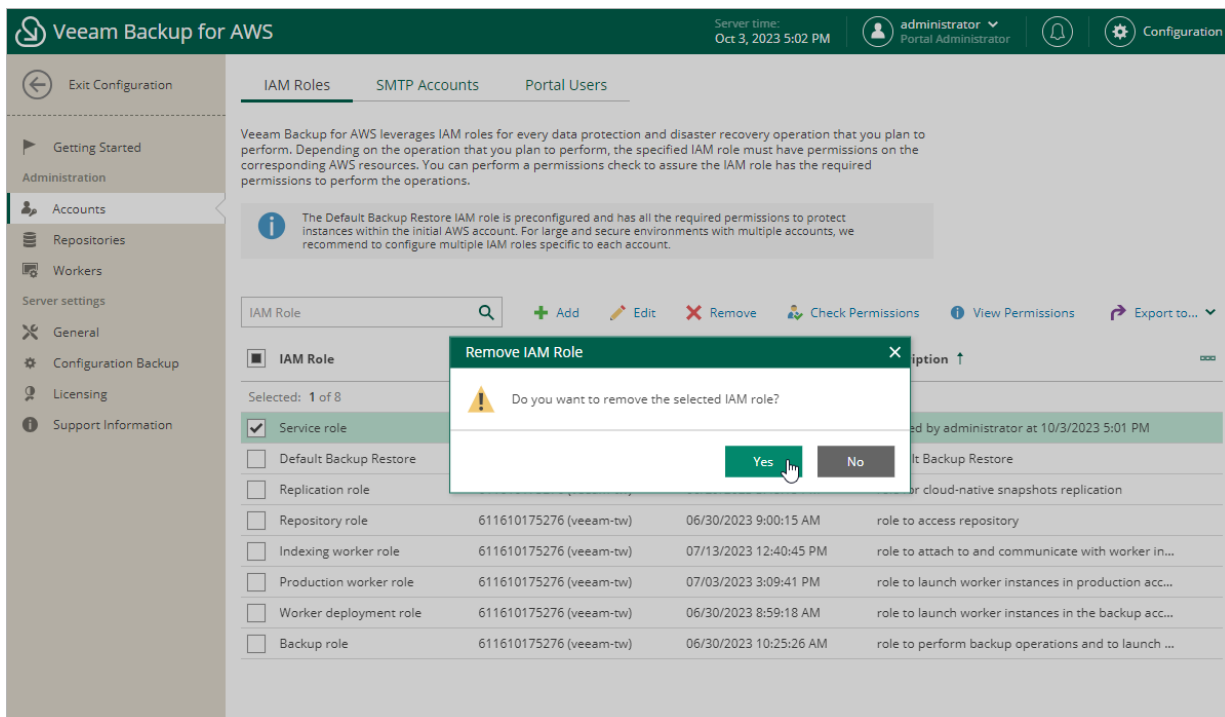
You can remove an IAM role from Veeam Backup for AWS if it is no longer used to perform data protection and disaster recovery operations.

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 be able to retrieve user identities from the identity provider, you must first [configure single sign-on settings](#).

To add a Veeam Backup for AWS user account, do the following:

1. Switch to the **Configuration** page.
2. Navigate to **Accounts > Portal Users**.
3. Click **Add**.
4. Complete the **Add Account** wizard:
 - a. 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.
 - b. At the **Account Info** step of the wizard, specify a name and description for the user account. An account name cannot be *admin*, can contain only lowercase Latin letters, numeric characters, underscores and dashes. You can use the dollar sign (\$) as the last character of the 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 maximum length of the description is 1024 characters.

IMPORTANT

If you have selected the **Identity Provider account** option at step 4a, 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](#).

- c. 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 4a, specify a password for the new Veeam Backup for AWS user account.

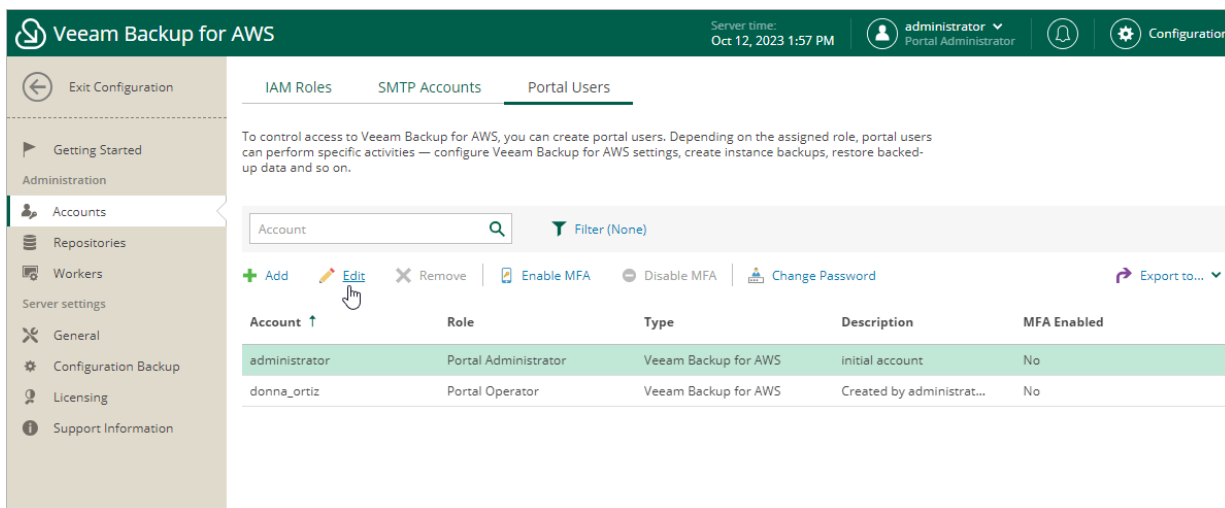
- d. At the **Summary** step of the wizard, review summary information and click **Finish**.

The screenshot shows the 'Add Account' wizard in the Veeam Backup for AWS interface. The top navigation bar includes the Veeam logo, 'Veeam Backup for AWS', server time 'Sep 10, 2021 12:38 PM', and user information 'administrator Portal Administrator'. The wizard title is 'Add Account'. On the left, a sidebar lists the steps: 'Account Type', 'Account Info', 'General Settings', and 'Summary' (which is highlighted). The main content area is titled 'Summary' and contains the text: 'Review the configured settings, and click Finish to exit the wizard.' Below this is a section titled 'Account' with the following details: Type: Identity Provider account, Name: sara_baker@company.com, Description: tw admin, and Role: Portal Administrator. 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. At the **AccountInfo** step of the wizard, edit a description of the user account.
 - b. At the **General Settings** step of the wizard, select a new role for the user account.
 - c. At the **Summary** step of the wizard, review summary information and click **Finish**.



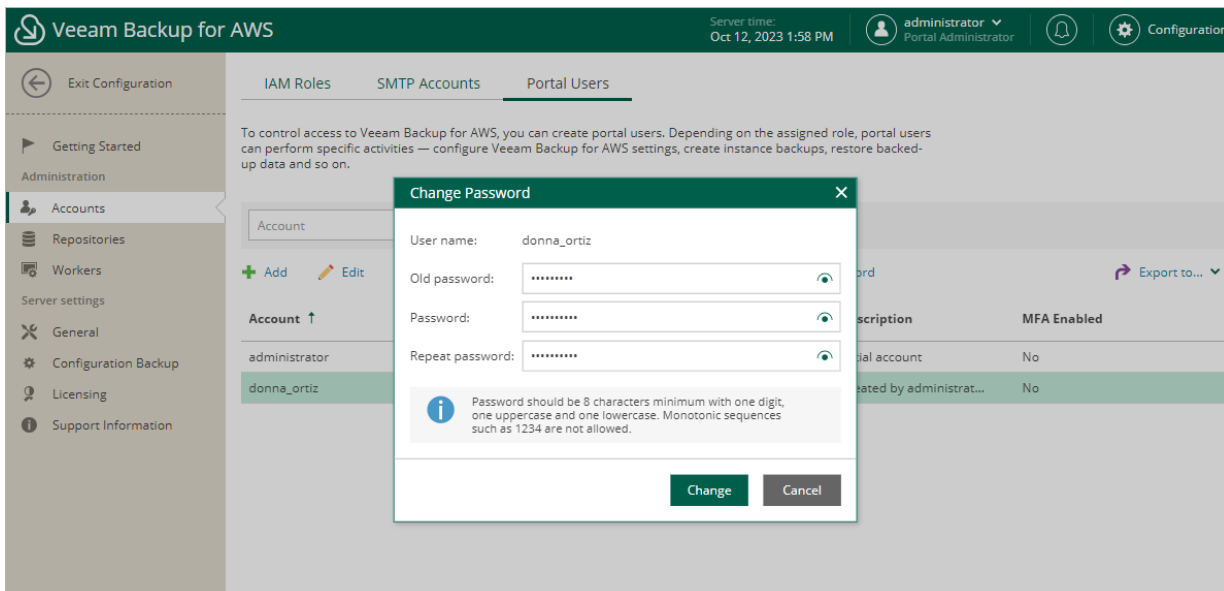
Changing User Passwords

For Veeam Backup for AWS user accounts, you can change the password specified while creating the account:

IMPORTANT

You cannot change the password for a user account whose user identity was obtained from an identity provider.

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**.



Configuring 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.

IMPORTANT

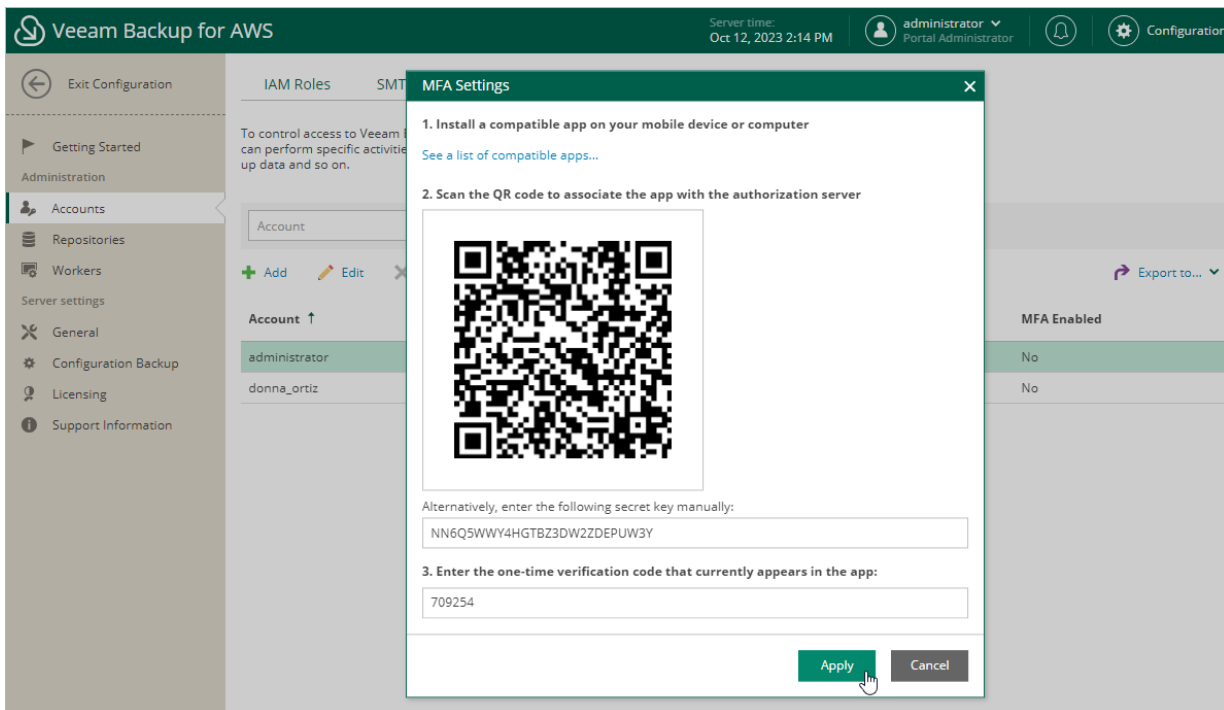
You cannot enable MFA for a user account whose user identity was obtained from an identity provider.

Enabling MFA

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**.



Disabling MFA

To disable MFA for a user account, select the account on the **Portal Users** tab and click **Disable MFA**.

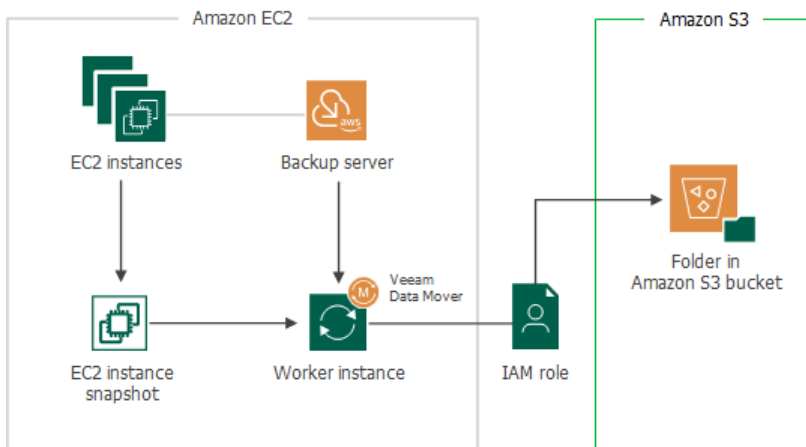
Managing Backup Repositories

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. That is why Veeam Backup for AWS verifies whether the backup repository is managed by any backup appliance – but only for the repository that was added to Veeam Backup for AWS version 7.0. If the backup repository is already managed by a backup appliance, you must import the repository to the current appliance in the repository settings to take ownership of this repository. For more information, see [Repository Ownership Alert](#).

To communicate with the backup repository, Veeam Backup for AWS uses the 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, the Veeam Data Mover establishes a connection with the repository enabling data transfer. To let the 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 Backups](#).

- **[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:

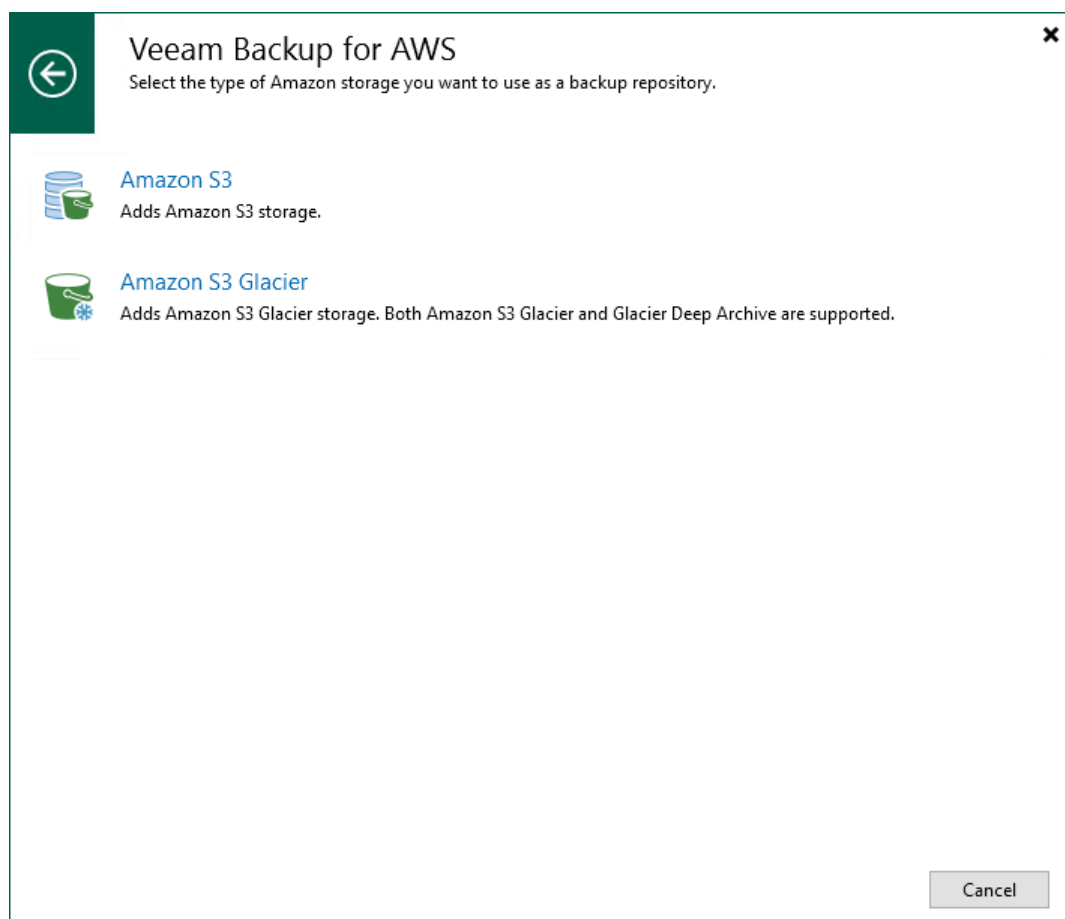
1. [Launch the Add External Repository wizard](#).

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. Wait for the repository to be added to the backup infrastructure.
8. 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 **Archive S3 Glacier** option if you want to create a repository with the S3 Glacier Flexible Retrieval or S3 Glacier Deep Archive storage class assigned.

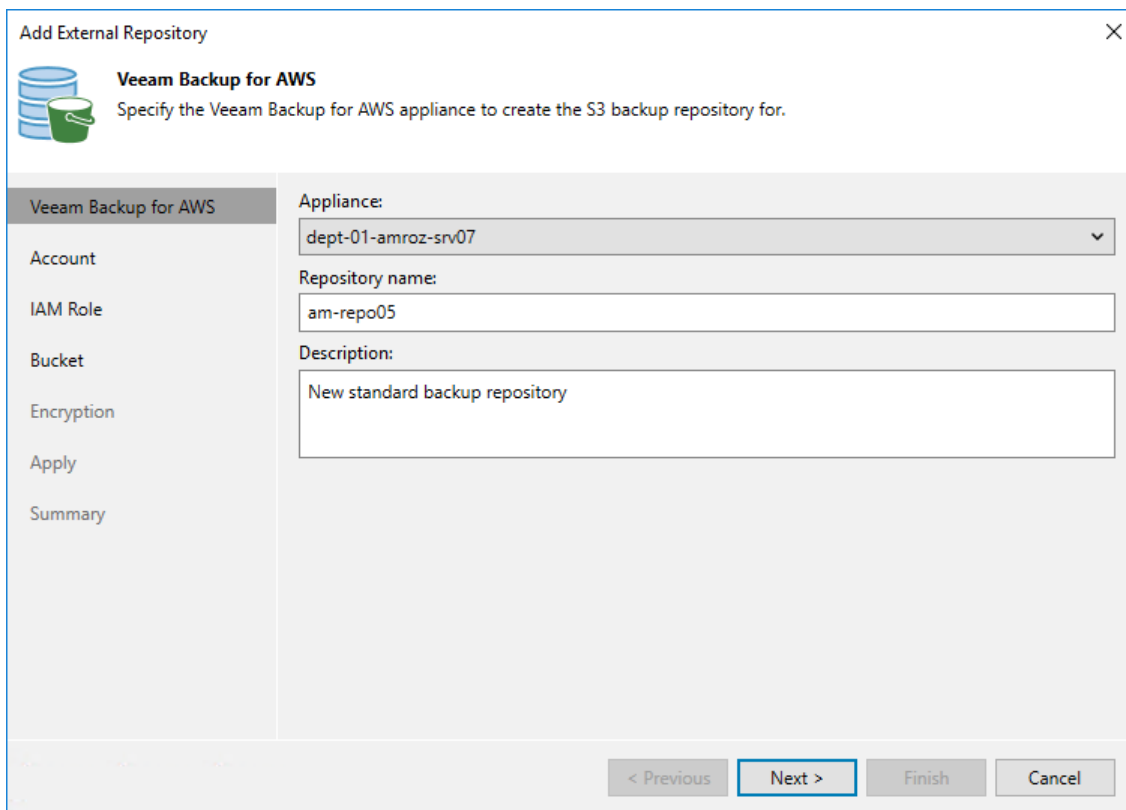


Step 2. Specify Repository Details

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 Appliance from Console](#) or [Adding 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 X

Veeam Backup for AWS
Specify the Veeam Backup for AWS appliance to create the S3 backup repository for.

Veeam Backup for AWS | Appliance: dept-01-amroz-srv07

Account | Repository name: am-repo05

IAM Role | Description: New standard backup repository

Bucket

Encryption

Apply

Summary

< 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. 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.

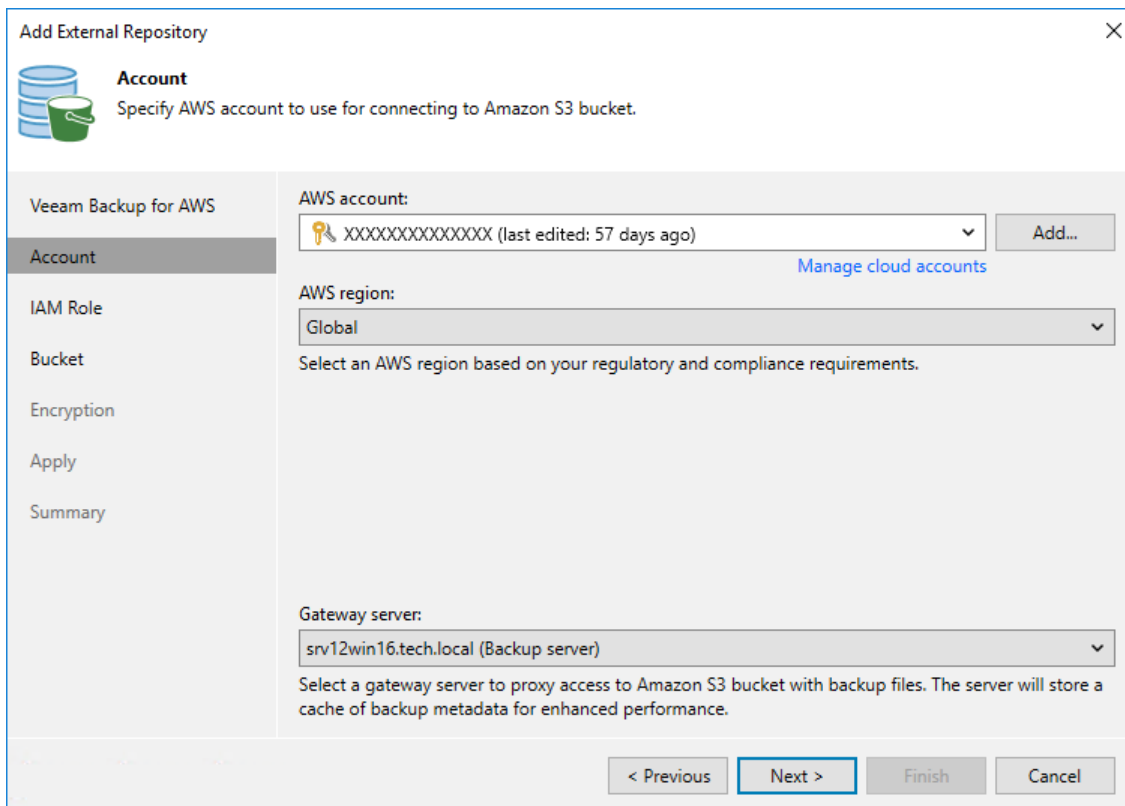
2. From the **AWS region** drop-down list, specify whether the repository will be located in an AWS Global or AWS GovCloud (US) region.

IMPORTANT

To check region availability, get infrastructure information and validate account permissions, Veeam Backup & Replication establishes a temporary test connection to 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.

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 (X) 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', 'Apply', and 'Summary'. The main area of the wizard is titled 'Account' and contains the following fields and controls:

- AWS account:** A dropdown menu showing 'XXXXXXXXXXXXXXXX (last edited: 57 days ago)' with a key icon on the left and an 'Add...' button on the right. Below the dropdown is a blue link labeled 'Manage cloud accounts'.
- AWS region:** A dropdown menu showing 'Global'. Below it is the text: 'Select an AWS region based on your regulatory and compliance requirements.'
- Gateway server:** A dropdown menu showing 'srv12win16.tech.local (Backup server)'. Below it is the text: '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.'

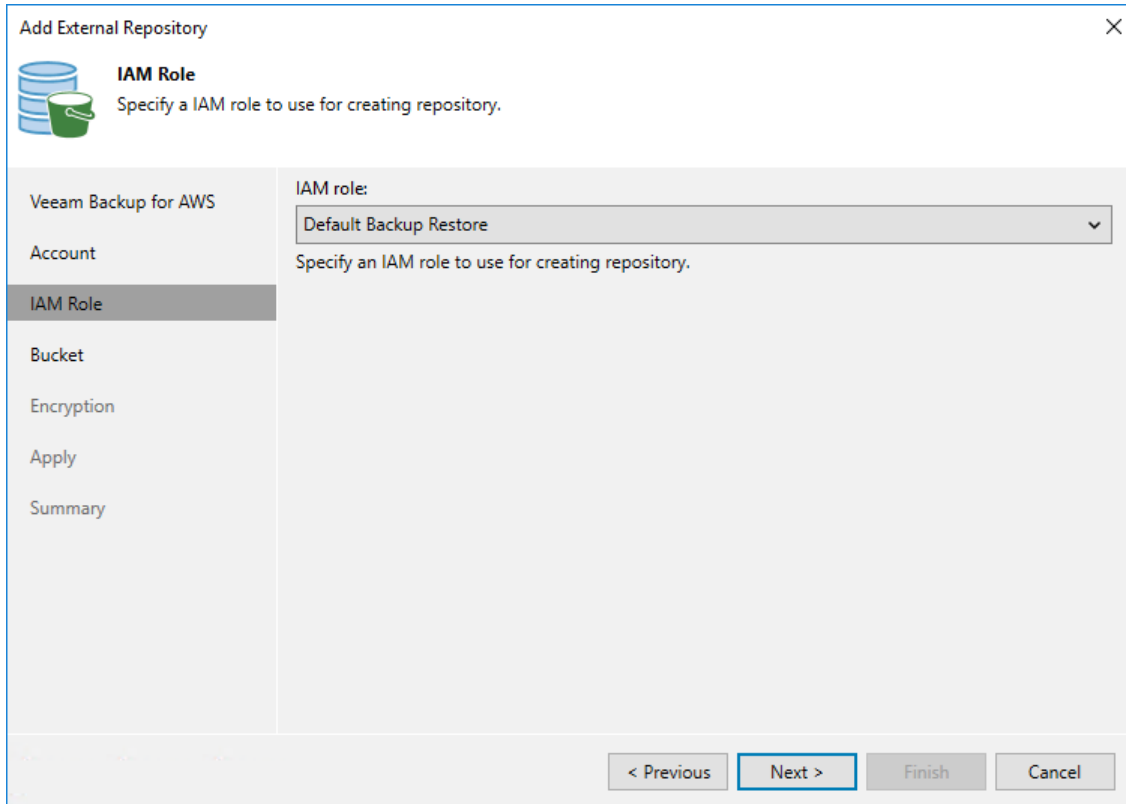
At the bottom of the wizard, there are four buttons: '< Previous', 'Next >' (which is highlighted with a blue border), 'Finish', and 'Cancel'.

Step 4. Specify IAM Role

[This step applies only if you have added multiple IAM roles to the backup appliance]

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. On the left side, there is a navigation pane with the following items: "Veeam Backup for AWS", "Account", "IAM Role" (which is highlighted), "Bucket", "Encryption", "Apply", and "Summary". Above the navigation pane, there 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." The main area of the wizard contains the label "IAM role:" above a dropdown menu. The dropdown menu currently displays "Default Backup Restore" with a downward arrow. Below the dropdown menu, the instruction "Specify an IAM role to use for creating repository." is repeated. At the bottom of the wizard, 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](#).

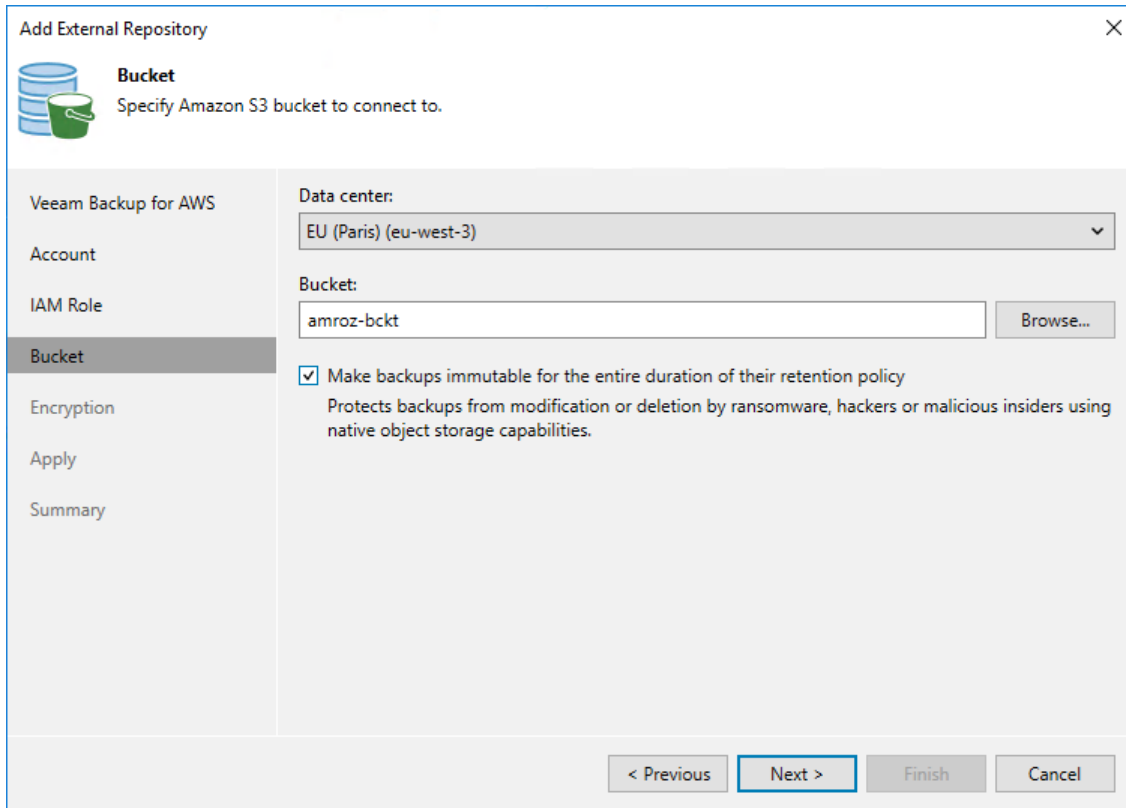
IMPORTANT

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.



The screenshot shows the 'Add External Repository' wizard in Veeam Backup for AWS, specifically the 'Bucket' step. The window title is 'Add External Repository' with a close button (X) in the top right corner. Below the title bar, there is a blue and green icon representing a bucket and the text 'Bucket Specify Amazon S3 bucket to connect to.' A left-hand navigation pane lists the steps: 'Veeam Backup for AWS', 'Account', 'IAM Role', 'Bucket' (which is highlighted), 'Encryption', 'Apply', and 'Summary'. The main area contains a 'Data center:' dropdown menu set to 'EU (Paris) (eu-west-3)'. Below that is a 'Bucket:' text input field containing 'amroz-bckt' and a 'Browse...' button. A checkbox is checked, labeled 'Make backups immutable for the entire duration of their retention policy', with a sub-note: 'Protects backups from modification or deletion by ransomware, hackers or malicious insiders using native object storage capabilities.' At the bottom of the window, there are four buttons: '< Previous', 'Next >' (which is highlighted with a blue border), 'Finish', and '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 backupfile 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 Role 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](#).

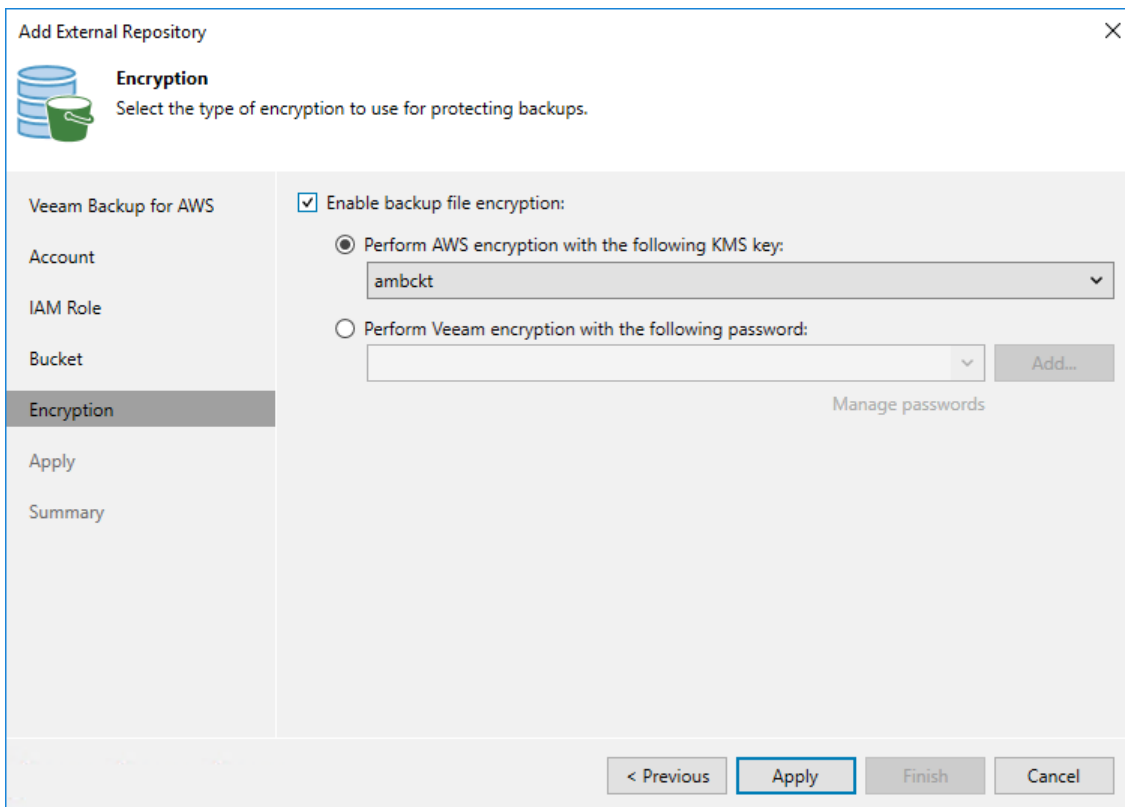
- To encrypt data using a password, select the **Perform Veeam encryption with the following password** option and choose the necessary password from the drop-down list.

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 password beforehand, you can do it without closing the wizard. To add the password, click either the **Manage passwords** link or the **Add** button, and specify a hint and the password in the **Password** window.

IMPORTANT

If you select the **Perform AWS encryption with the following KMS key** option, consider the following:

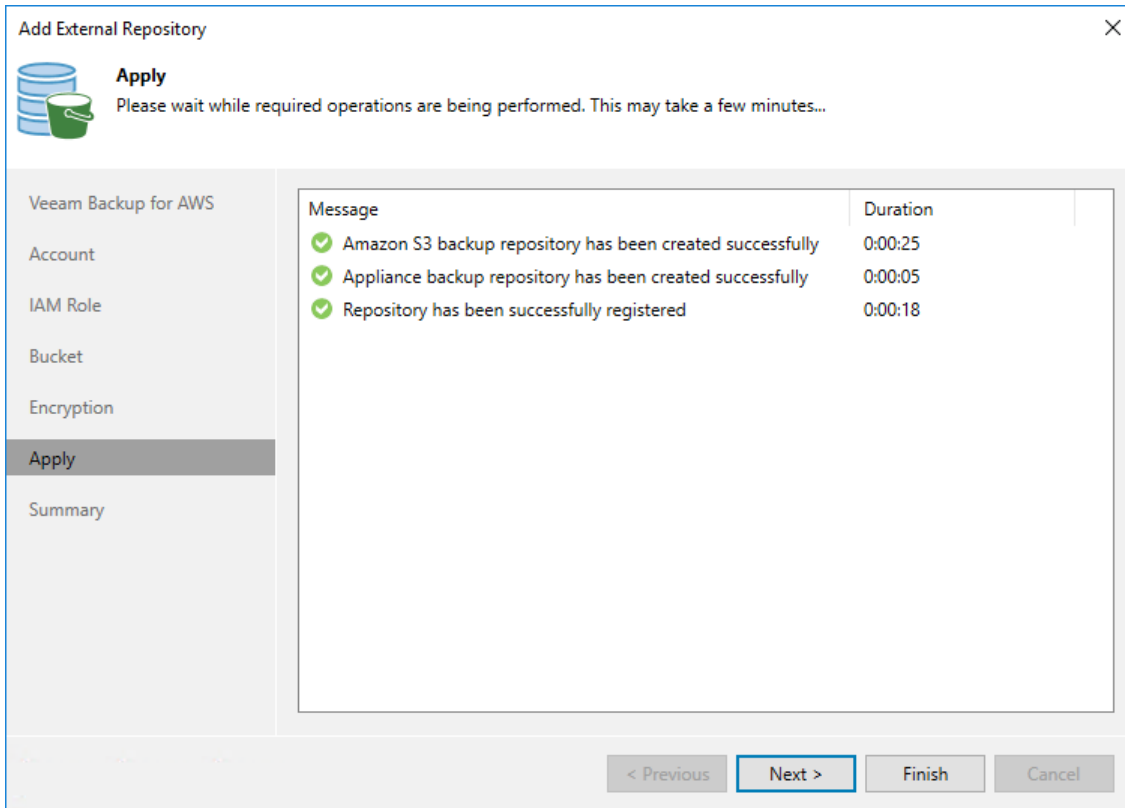
- Only symmetric KMS keys are supported.
- For Veeam Backup & Replication to be able to decrypt data stored in the created repository, the IAM user specified at [step 3](#) of the wizard must have [permissions to access the selected KMS key](#).
- Do not disable the KMS key specified in the repository 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' wizard in the 'Encryption' step. The window title is 'Add External Repository' with a close button (X) in the top right corner. Below the title bar, there is an icon of a database and a bucket, followed by the heading 'Encryption' and the instruction 'Select the type of encryption to use for protecting backups.' A left-hand navigation pane contains the following items: 'Veeam Backup for AWS', 'Account', 'IAM Role', 'Bucket', 'Encryption' (which is highlighted), 'Apply', and 'Summary'. The main content area features a checked checkbox labeled 'Enable backup file encryption:'. Below this, there are two radio button options. The first option, 'Perform AWS encryption with the following KMS key:', is selected and has a dropdown menu showing 'ambckt'. The second option, 'Perform Veeam encryption with the following password:', is unselected and has an empty text input field with a dropdown arrow and an 'Add...' button to its right. Below the password field is a 'Manage passwords' link. At the bottom of the window, there are four buttons: '< Previous', 'Apply' (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 creating the repository. At the **Apply** step of the wizard, wait for the process to complete and click **Next**.



Add External Repository [Close]

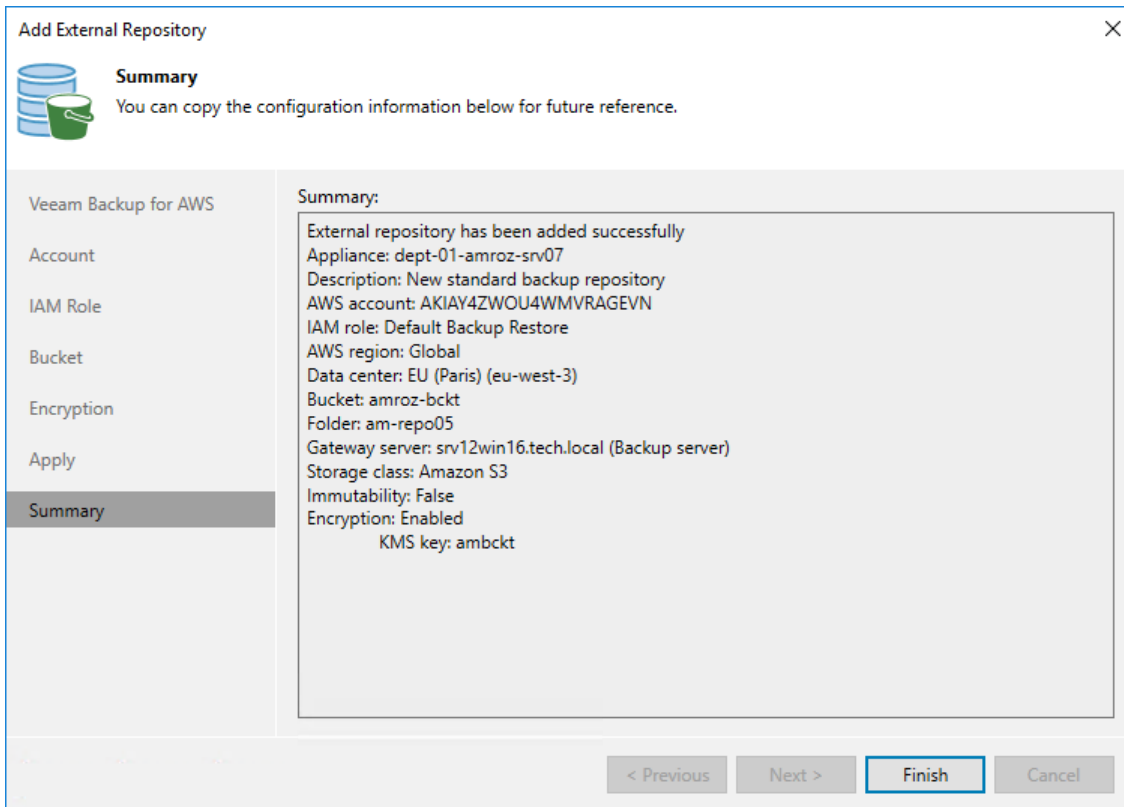
Apply
Please wait while required operations are being performed. This may take a few minutes...

Message	Duration
✓ Amazon S3 backup repository has been created successfully	0:00:25
✓ Appliance backup repository has been created successfully	0:00:05
✓ Repository has been successfully registered	0:00:18

< Previous **Next >** Finish Cancel

Step 8. 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

You can use only existing Amazon S3 buckets to create backup repositories. Before you add a backup repository, check [limitations for backup repositories](#).

To add a backup repository, do the following:

1. [Launch the Add Repository wizard](#).
2. [Specify a backup repository name and description](#).
3. [Configure backup repository settings](#).
4. [Enable data encryption for the backup repository](#).
5. [Specify an S3 interface endpoint](#).
6. [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.
- If the backup repository is already managed by any backup appliance, you must import the repository to the current appliance at [step 3](#) of the wizard to take ownership of this repository. Consider that as soon as you import the repository to the current appliance, the backup policies configured on the previous appliance will start failing.

- 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](#).
- Amazon S3 buckets with only S3 Object Lock enabled is not supported. It is recommended that S3 Object Lock and S3 Versioning are either both enabled or both disabled for a bucket.
- You cannot change immutability settings for the repository since these settings are based on the immutability settings of the selected Amazon S3 bucket, which are configured in the AWS Management Console upon bucket creation and cannot be modified afterward. For more information, see [AWS Documentation](#).
- 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 is divided into a top navigation bar, a left sidebar, and a main content area. The top bar includes the Veeam logo, the text 'Veeam Backup for AWS', the server time 'Nov 30, 2023 5:05 PM', the user 'administrator Portal Administrator', and a 'Configuration' link. The left sidebar has a back arrow and the title 'Add Repository', with tabs for 'Info', 'Bucket', 'Settings', and 'Summary'. The 'Info' tab is active, showing the title 'Specify repository name and description' and the instruction 'Enter a name and description for the repository.' Below this, there are two input fields: 'Name:' with the value 'Immutable Repository' and 'Description:' with the value 'Repository for storing immutable data'. At the bottom of the main area, there are 'Next' and 'Cancel' buttons.

Veeam Backup for AWS

Server time: Nov 30, 2023 5:05 PM

administrator Portal Administrator

Configuration

← Add Repository

Info **Specify repository name and description**

Enter a name and description for the repository.

Bucket

Settings

Summary

Name:
Immutable Repository

Description:
Repository for storing immutable data

Next Cancel

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 required permissions 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 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 **Select**. 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 a 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. The 'Bucket' tab is active, displaying the following configuration:

- 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.
- IAM role:** A dropdown menu is set to 'Repository role (role to access repository)'. There are '+ Add' and 'Check Permissions' buttons.
- Location:** The bucket is 'am-bckt-immutable'. There are two options for folder creation: 'Use existing folder' (unselected) and 'Create new folder' (selected) with the value 'dept_01_immutable'.
- Storage class:** A dropdown menu is set to 'S3 Standard'.
- Immutability settings:** A checkbox 'Backups stored in this repository will be immutable' is checked. A note states: 'Immutability settings cannot be changed. Immutability will be enabled for the entire duration of the retention policy.'

At the bottom, there are 'Previous', 'Next', and 'Cancel' buttons.

NOTE

As soon as you click **Next**, Veeam Backup for AWS verifies whether the backup repository is managed by a backup appliance. If the backup repository is already managed by any backup appliance, you will receive a warning. To learn how to eliminate this warning, see [Repository Ownership Alert](#).

Repository Ownership Alert

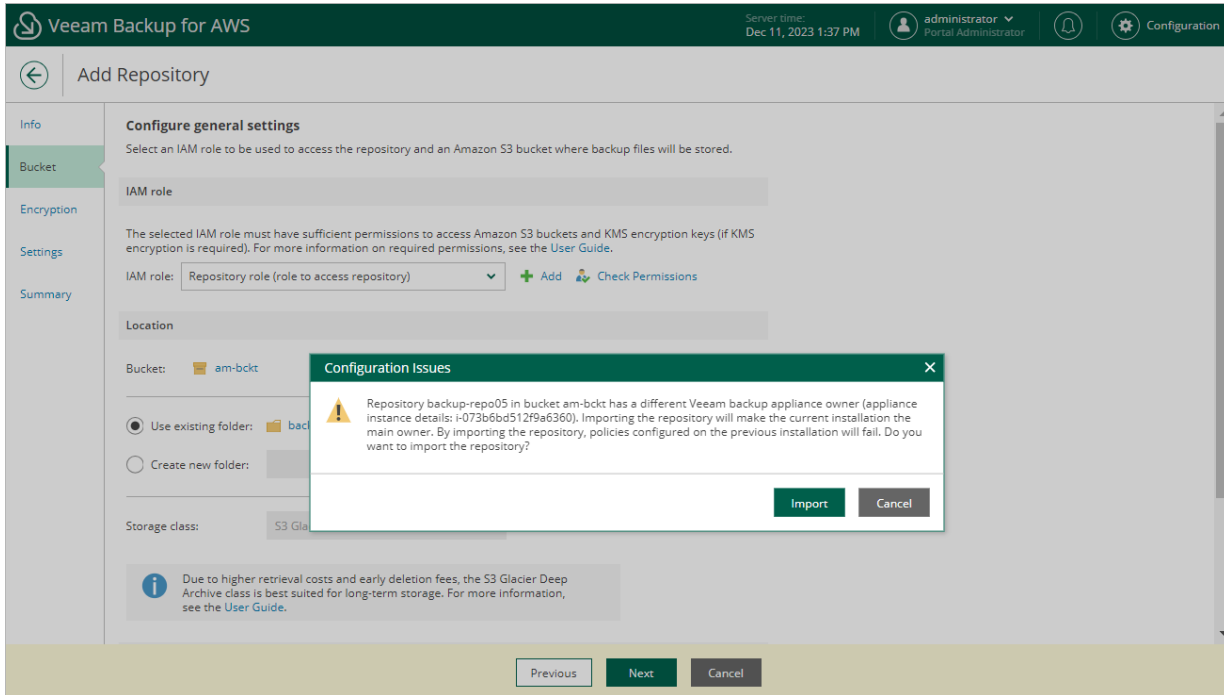
To prevent the same backup repository from being used simultaneously on different backup appliances, Veeam Backup for AWS verifies whether the backup repository is managed by any backup appliance when you add an existing folder as a target backup repository. Retention sessions running on different appliances may corrupt backup files stored in this repository, which may result in unpredictable data loss.

If the backup repository is already connected to any backup appliance, Veeam Backup for AWS will display a warning notifying that the backup repository has a different backup appliance owner. To allow Veeam Backup for AWS to take ownership of this repository, click **Import**. If you do not import the repository to the current backup appliance, you will not be able to proceed with the wizard.

IMPORTANT

Consider the following:

- Veeam Backup for AWS verifies the backup appliance owner only for those backup repositories that were added to Veeam Backup for AWS version 7.0.
- As soon as you import the backup repository to the current backup appliance, the backup policies configured on the previous backup 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.
- 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](#).

The screenshot shows the 'Add Repository' wizard in Veeam Backup for AWS. The 'Encryption' step is active, showing options to 'Configure encryption settings'. The 'Enable encryption' toggle is set to 'On'. Under 'Use password encryption', the 'Password' field is masked with dots, 'Repeat password' is also masked, and 'Password hint' is 'hint'. Under 'Use KMS encryption key', the 'Encryption key' dropdown is empty. A warning message states: 'The selected IAM role must have permissions to access the encryption key. For more information, see the User Guide.' The bottom navigation bar includes 'Previous', 'Next', and 'Cancel' buttons.

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.

3. After the required permissions are granted, close the **Permission check** window, review configuration information and click **Finish**.

As soon as you click **Finish**, Veeam Backup for AWS will start adding the backup repository to infrastructure. To track the progress, select the **Go to Sessions** check box to proceed to the [Sessions Log](#) tab.

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**.

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:

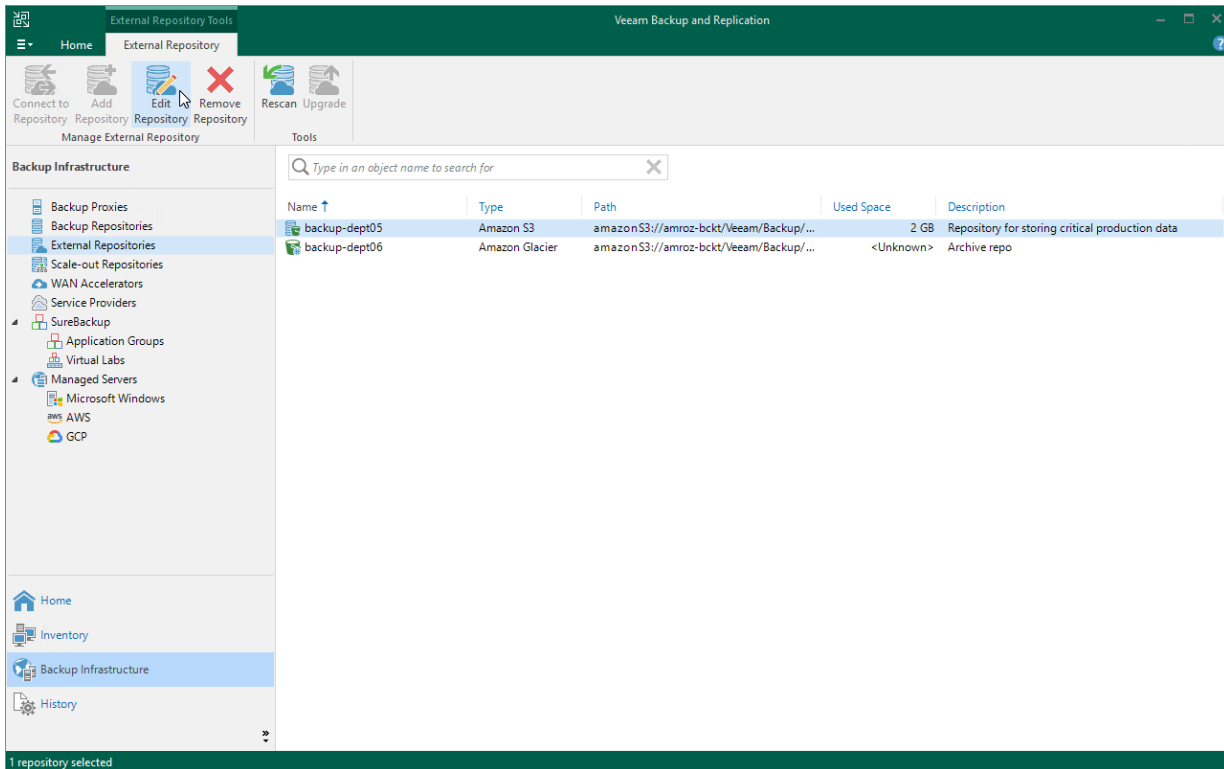
1. In the Veeam Backup & Replication console, open the **Backup Infrastructure** view.
2. Navigate to **External Repositories**.
3. Select the necessary repository and click **Edit Repository** on the ribbon.
Alternatively, you can right-click the repository and select **Properties**.
4. Complete the **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**.



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. To change the backup repository owner of a repository managed by another backup appliance, 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 and click **Finish** to confirm the changes.

As soon as you click **Finish**, Veeam Backup for AWS will start modifying the backup repository settings. To track the progress, select the **Go to Sessions** check box to proceed to the [Sessions Log tab](#).

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 "Nov 30, 2023 5:10 PM", and the user "administrator Portal Administrator". The left sidebar contains navigation options: Exit Configuration, Getting Started, Administration, Accounts, Repositories (selected), Workers, Server settings, General, Configuration Backup, Licensing, and Support Information. The main content area displays a table of backup repositories. A search bar and a filter dropdown are at the top. Below the search bar are buttons for Add, Edit, View Info, and Remove. The table has columns for Repository, Description, Bucket, Folder, Storage Class, and Immutability. Three repositories are listed: backup-repo-02 (archive backup repository), backup-repo-03 (primary backup repository), and backup-repo-04 (Created by administrator at...). The "backup-repo-03" row is selected, indicated by a checkmark in the first column and a green background. A "Selected: 1 of 3" indicator is shown above the table. An "Export to..." button is located at the top right of the table area.

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

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.

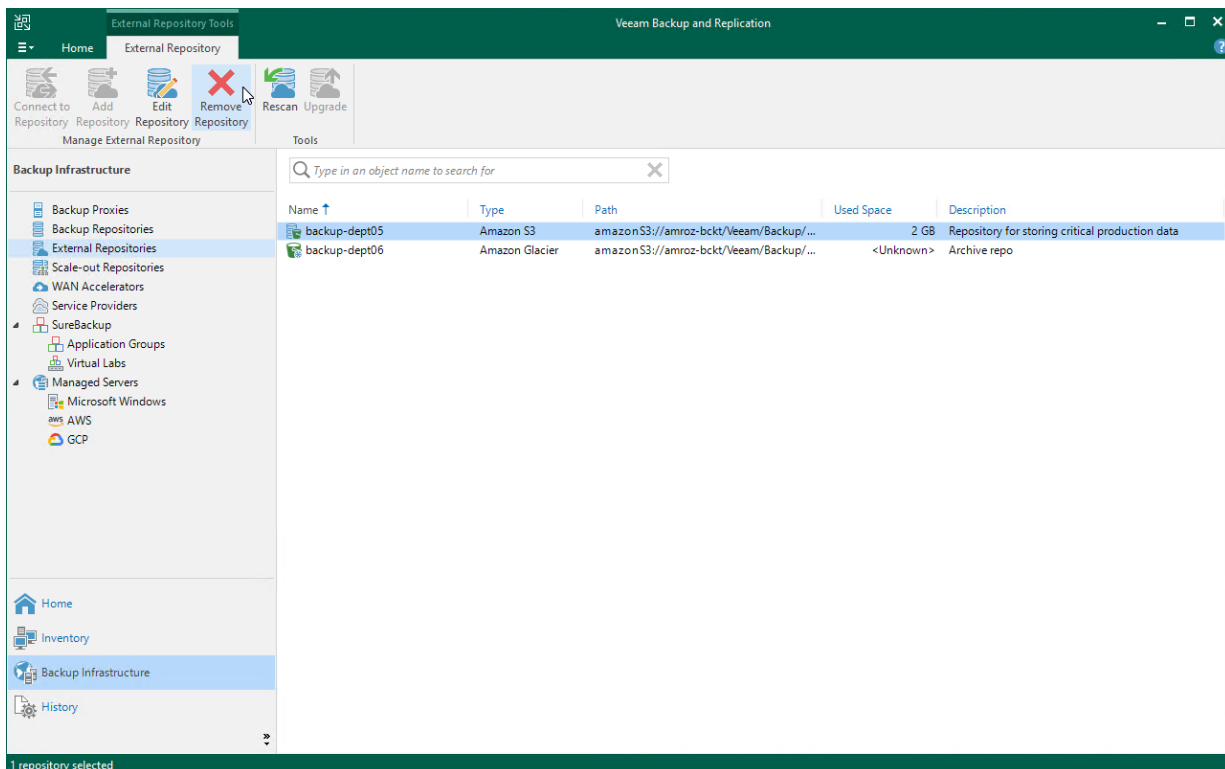
Removing Backup Repository Using Veeam Backup & Replication Console

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](#).

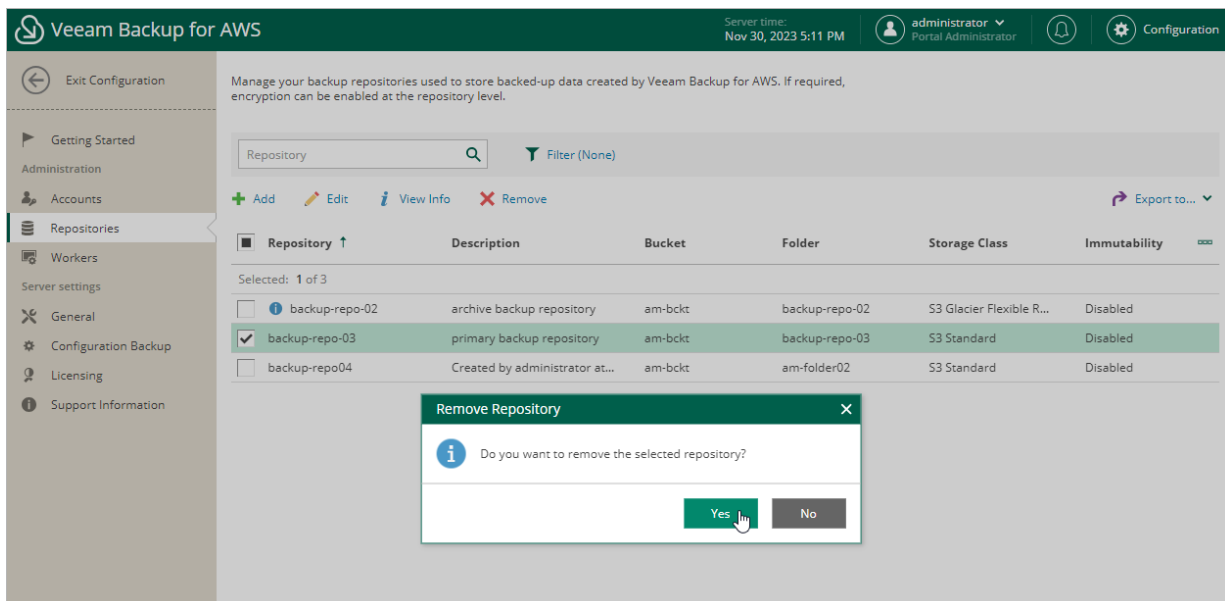
If you no longer need the backed-up data, either delete it as described in section [Managing Backed-Up Data](#) before you remove the repository from Veeam Backup for AWS, or [use the AWS Management Console](#) to delete the data if the repository has already been removed.

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 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 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 launched in a specific AWS Region for the duration of the backup, restore and retention process. AWS Regions in which Veeam Backup for AWS launches worker instances to perform operations are predefined and described in section [Worker Instances](#). However you can choose whether you want Veeam Backup for AWS to launch worker instances in the backup account or in production AWS accounts, specify network settings and instance types that will be used to launch worker instances.

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.

Veeam Backup for AWS can launch worker instances in the following AWS accounts:

- The backup account – an AWS account to which the service IAM role specified to launch worker instances belongs. Veeam Backup for AWS uses this account to launch worker instances for backup, restore and backup retention operations – unless instructed to launch the worker instances in production accounts.
- Production accounts – the same AWS accounts to which the processed resources belong. Veeam Backup for AWS uses these accounts to launch worker instances for EFS indexing, RDS backup and restore operations by default.

To allow Veeam Backup for AWS to use a production account when deploying worker instances for EC2 backup or restore operations, this functionality must be enabled in the backup policy or restore settings.

Adding Configurations for Backup Account

By default, Veeam Backup for AWS launches 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 launch worker instances – the role is preconfigured and has all the required permissions. Therefore, the default backup account is an AWS account where 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:

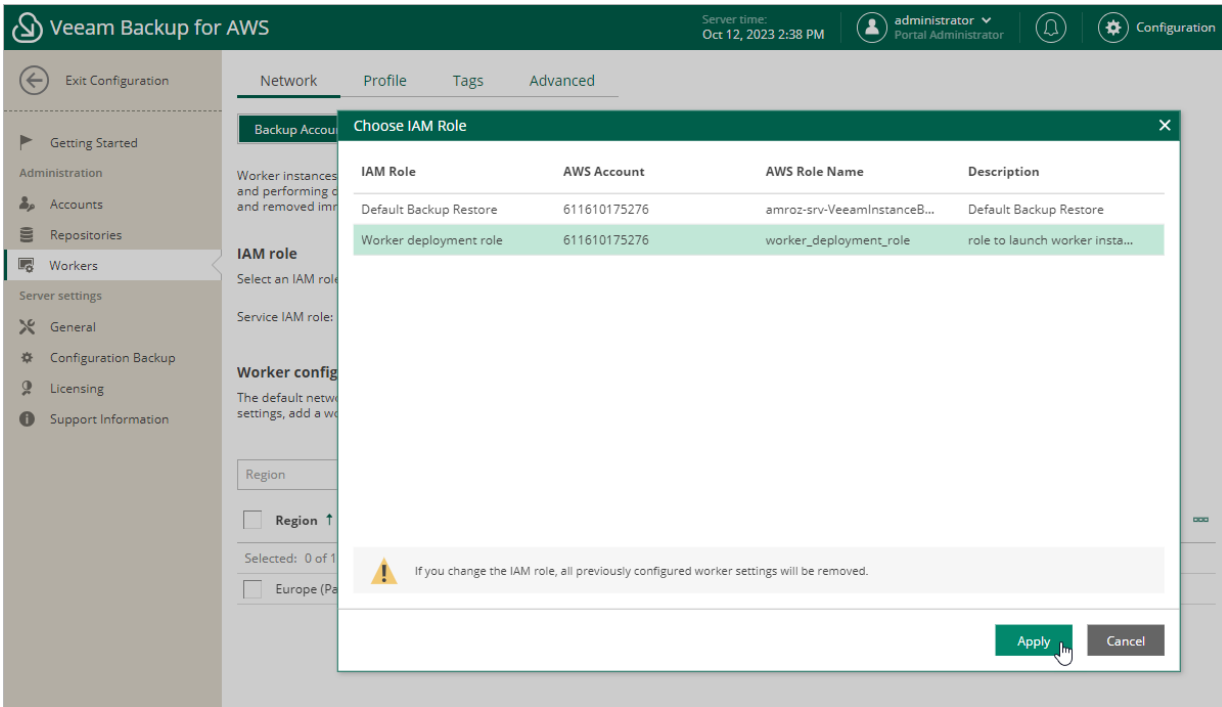
1. Switch to the **Configuration** page.
2. Navigate to **Workers > Network**.
3. At the **Backup Accounts** tab, click the link in the **Service IAM role** field.
4. In the **Choose IAM Role** window, select the necessary IAM role, and then click **Apply**.

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 launch 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 launch worker instances, see [Service IAM Role in Backup Account](#).



Adding Worker Configurations

To launch 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. 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, 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 launched 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](#).

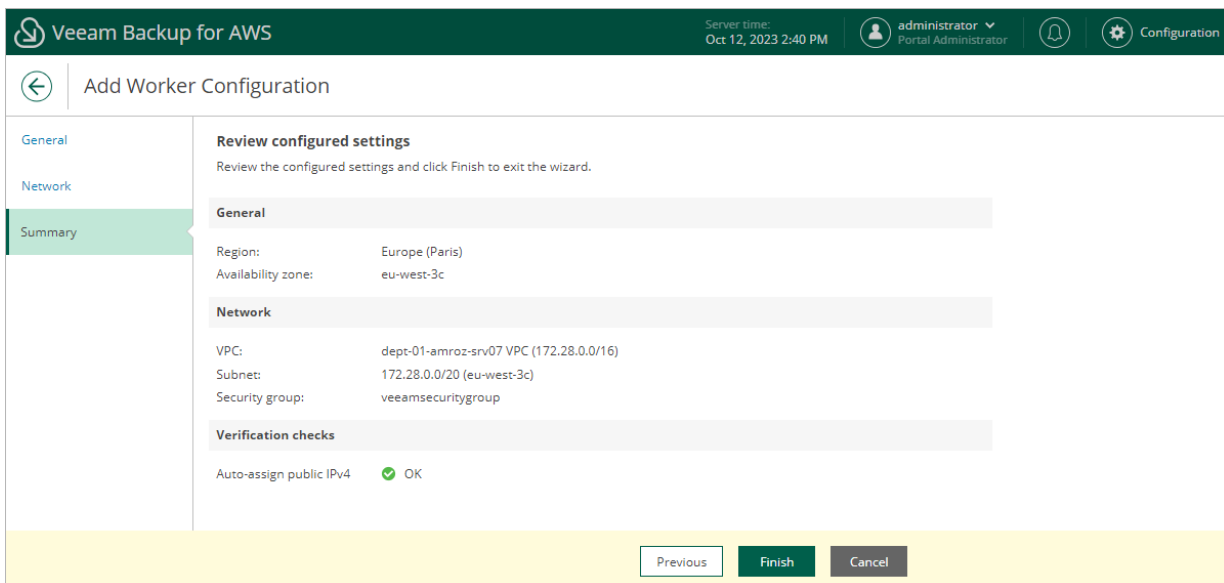
By default, Veeam Backup for AWS uses public IPv4 addresses to communicate with worker instances. If the public IPv4 addressing attribute is disabled for the selected subnet, Veeam Backup for AWS will display a warning at the **Summary** step of the wizard. In this case, do either of the following:

- 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 [Enabling 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**.



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 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 launch 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.

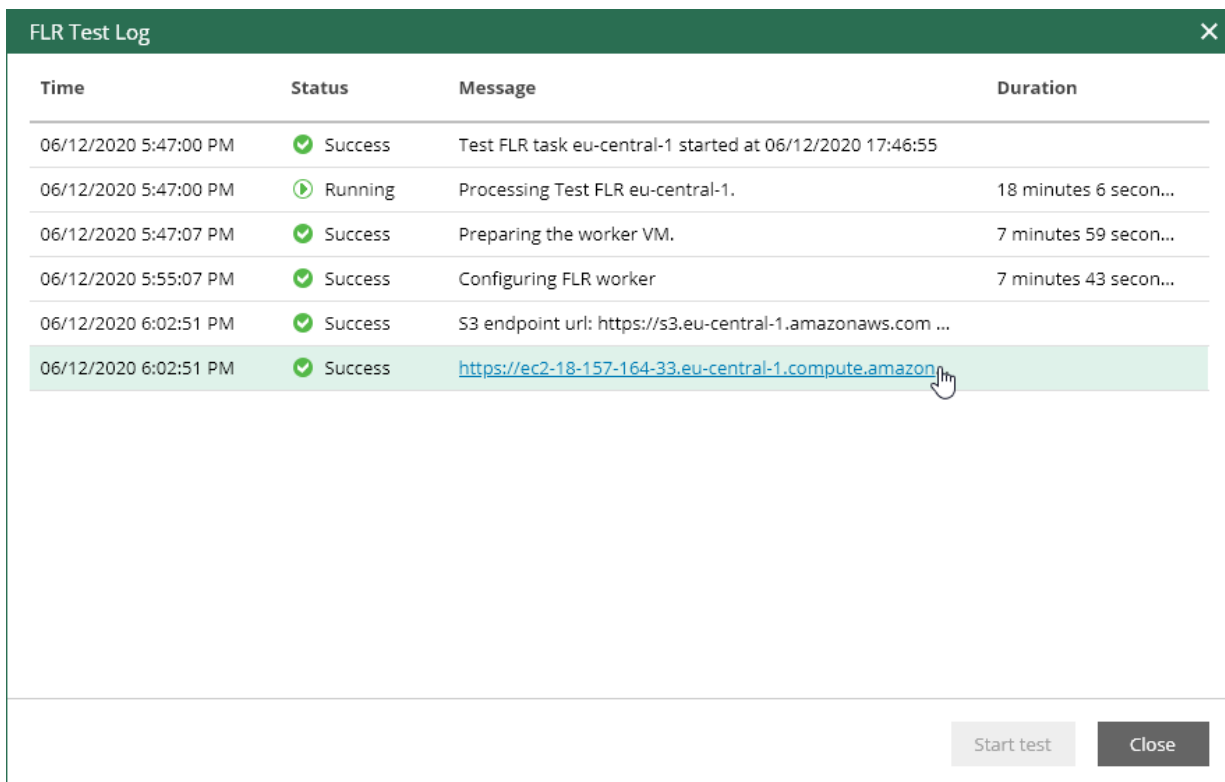
Note that the security group associated with worker instances must allow inbound internet access from the machine from which you plan to open the file-level recovery browser.

- 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, in the **FLR Status** column, click the status of the file-level recovery test, and then, click **Start test** in the **FLR Test Log** window.



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 launched in production accounts by default. However, if you also want Veeam Backup for AWS to launch 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 launch 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 launch 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 launched worker instances and further used by Veeam Backup for AWS to communicate with the instances. The role must be assigned permissions listed in section [Indexing Worker IAM Role Permissions](#), [Worker IAM Role Permissions](#) 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 launch worker instances.

Adding Worker Configurations

To launch 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 where 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 add an IAM role, 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 launching 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.

TIP

If the newly created worker configuration will be used to perform only EC2 backup operations, there is no need to select the availability zone where the processed EC2 instances reside – you can select any zone in the specified region.

- 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 launched in the specified location. For EFS indexing, Veeam Backup for AWS will also apply these settings to worker instances launched 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 [Enabling 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 the Veeam Backup for AWS console. The 'Summary' step is selected in the left-hand navigation pane. The main content area displays the following information:

- Review configured settings**
Review the configured settings and click Finish to exit the wizard.
- General**
 - Account: 359000203834 (veeam-backup)
 - Region: US East (Virginia)
 - Availability zone: us-east-1a
- Network**
 - VPC: veeamvpc (172.31.0.0/16)
 - Subnet: 172.31.0.0/16 (us-east-1a)
 - Security group: veeamsecuritygroup
- Verification checks**
 - Auto-assign public IPv4: ⚠ Warning: Auto IP-assignment is off. To work in a private network, configure endpoints for Amazon services. For more information, see the [User Guide](#).

At the bottom of the wizard, there are three buttons: 'Previous', 'Finish' (highlighted with a mouse cursor), and 'Cancel'.

Editing Configurations

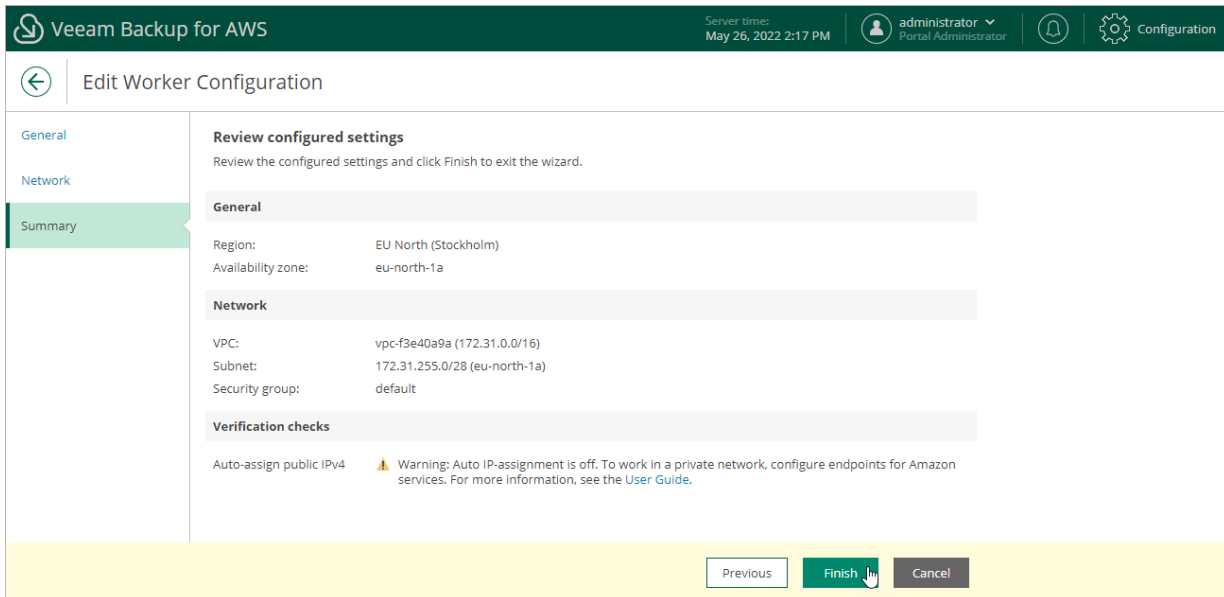
You can edit worker configurations added for AWS Regions:

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 launched 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 console interface. The top navigation bar includes the Veeam logo, the text 'Veeam Backup for AWS', the server time 'Oct 12, 2023 2:42 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 divided into tabs: 'Network', 'Profile', 'Tags', and 'Advanced'. Under the 'Network' tab, there are sub-tabs for 'Backup Accounts' and 'Production Accounts'. A text block explains that worker instances are temporary Linux-based EC2 instances. Below this, the 'IAM role' section is partially visible. 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 'Europe (Paris)'. A dialog box titled 'Remove Worker Configuration' is overlaid on the table, containing an information icon and the text: '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.

Remove Worker Configuration

If you remove the worker configuration, the default network settings will be used to launch worker instances in the selected region.

Remove No

Region	Availability Zone	Virtual Private Cl...	Subnet	Security Group	FLR Status
Europe (Paris)	eu-west-3b	amroz-srv VPC (17...	172.28.0.0/20 (eu-...	amroz-srv-VcbSecu...	Never Executed
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 launches 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 backups if the total size of all EBS volumes of the processed instance is more than 6 TB.

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 launch 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 Instances](#). 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 launched, 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.



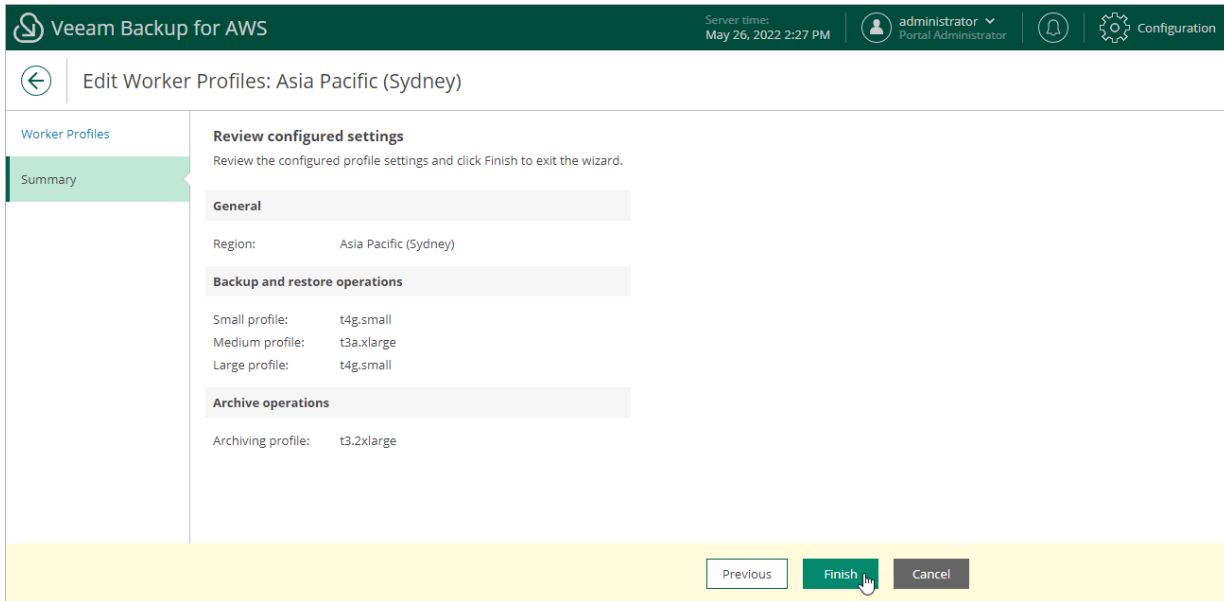
Editing Profiles

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.

Remove Worker Profile

If you remove the worker profile, the default profile settings will be used to launch worker instances in the selected region.

Remove No

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 launched 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 displays the Veeam Backup for AWS configuration interface. The top navigation bar shows the user is logged in as 'administrator' (Portal Administrator) and is in the 'Configuration' section. The left sidebar lists various configuration areas, with 'Workers' selected. The main content area is titled 'Tags' and includes a 'Save' button and a warning: 'Your changes are not saved yet.' Below this, a message states: 'You can assign custom tags to workers and use this for billing, security, monitoring and reporting services.' The interface features two input fields: 'Key:' with the value 'policyname' and 'Value:' with the value '%policyName%'. An '+ Add' button is positioned to the right of the Value field. Below these fields, a tag 'applianceid: %applianceid%' is shown with a close icon. A note indicates 'A maximum of 25 custom tags is allowed.' An information box at the bottom provides details on runtime variables:

Parameter	Description
<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.

- [Enable 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.](#)

Enabling 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 launch 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 enable the private network deployment functionality, 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 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 navigation options: 'Exit Configuration', 'Getting Started', 'Administration', 'Accounts', 'Repositories', 'Workers', 'Server settings', 'General' (selected), '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 cloud-native snapshots and 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 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. The number must be between 15 and 36135.

If you select this option, Veeam Backup for AWS will first wait for the specified period of time after an instance stops being processed by a backup policy, and then will remove its obsolete snapshots from the configuration database as soon as the 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.

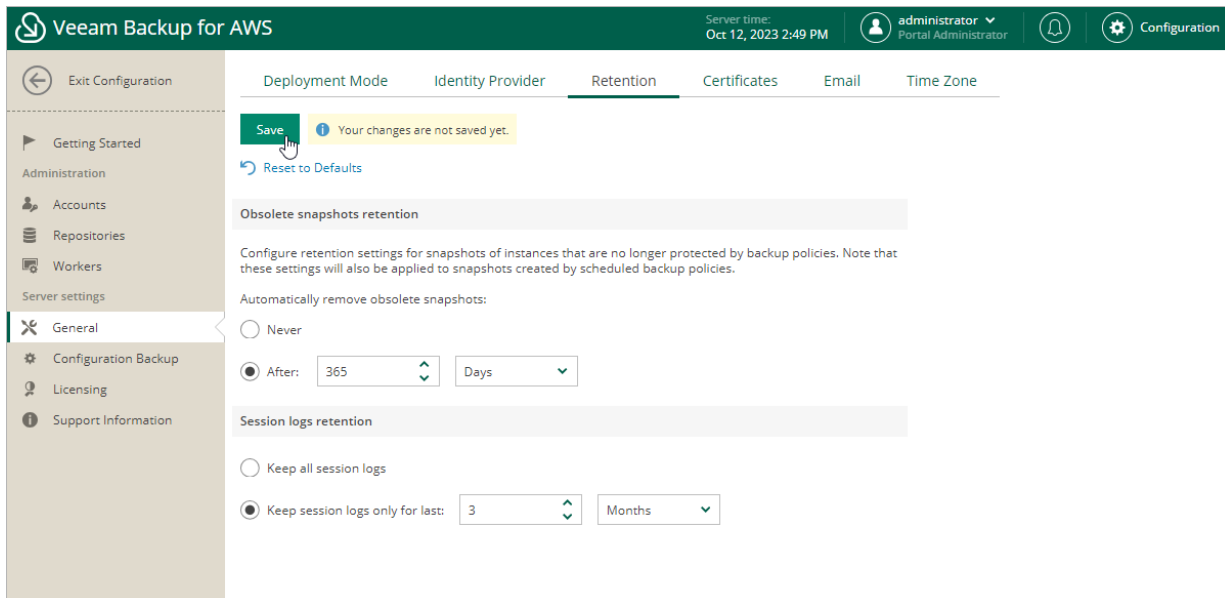
To configure retention settings for session records, do the following:

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 > E-mail**.
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 E-mail**. 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.

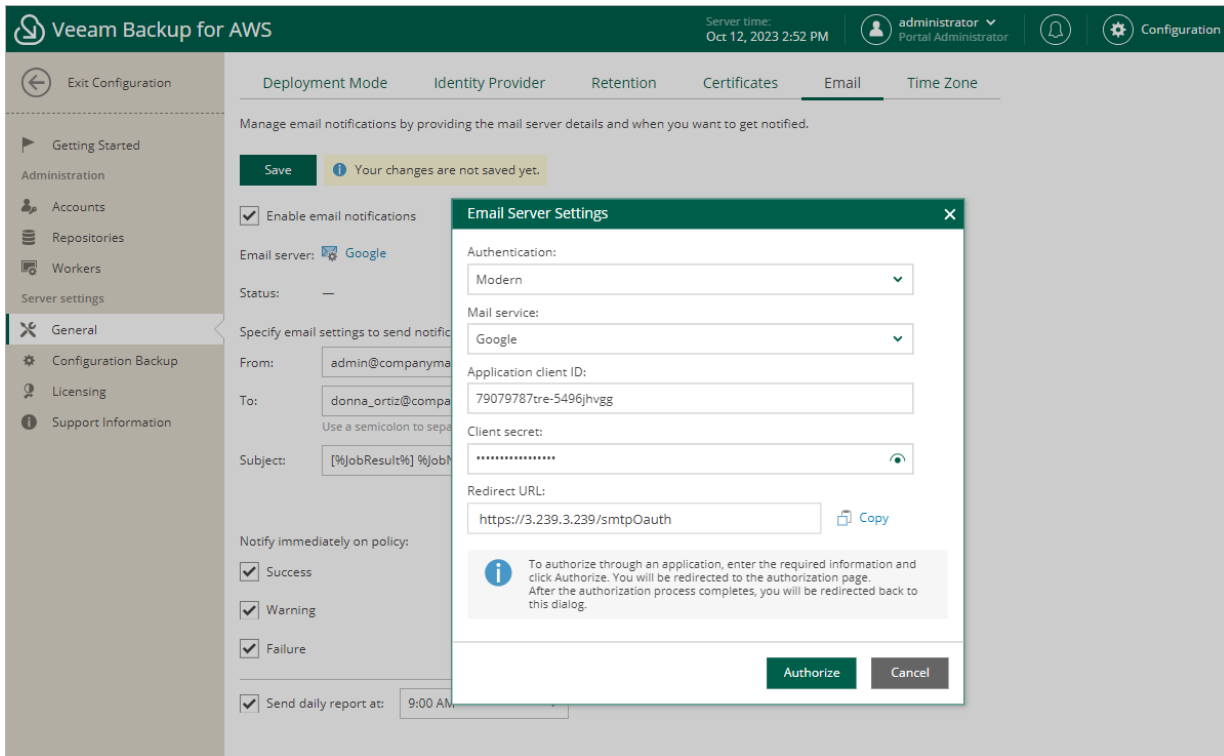
7. Click **OK**.

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 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 URI specified for the application matches the URL copied from the Veeam Backup for AWS Web UI.
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.
- e. Click **OK**.



Adding SMTP Accounts

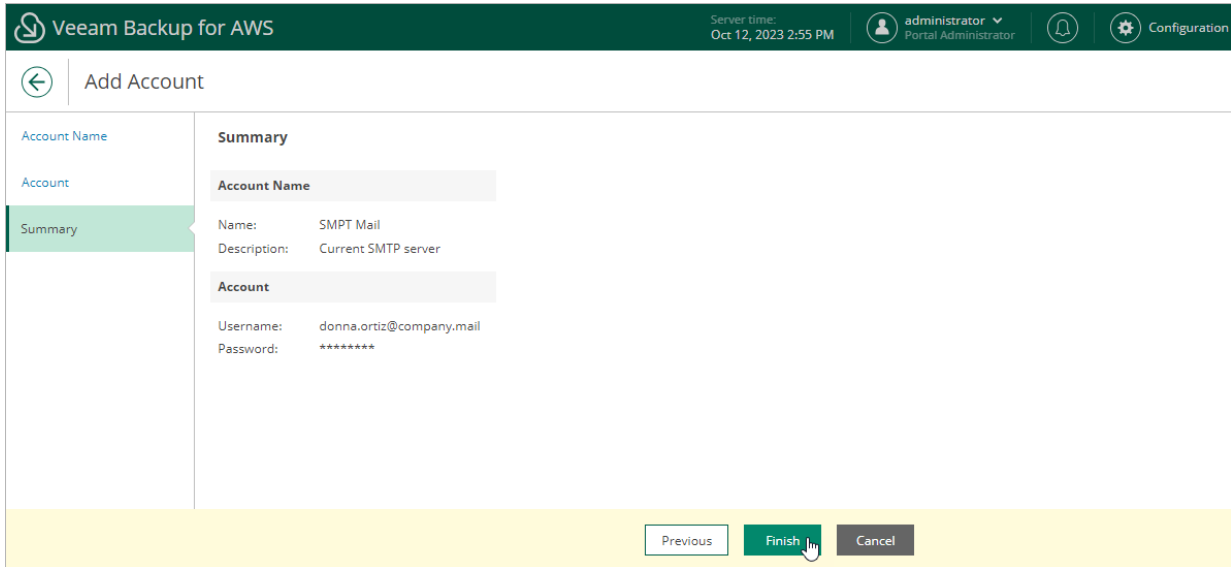
To add an account that will be used to connect to an SMTP server, do the following:

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**.



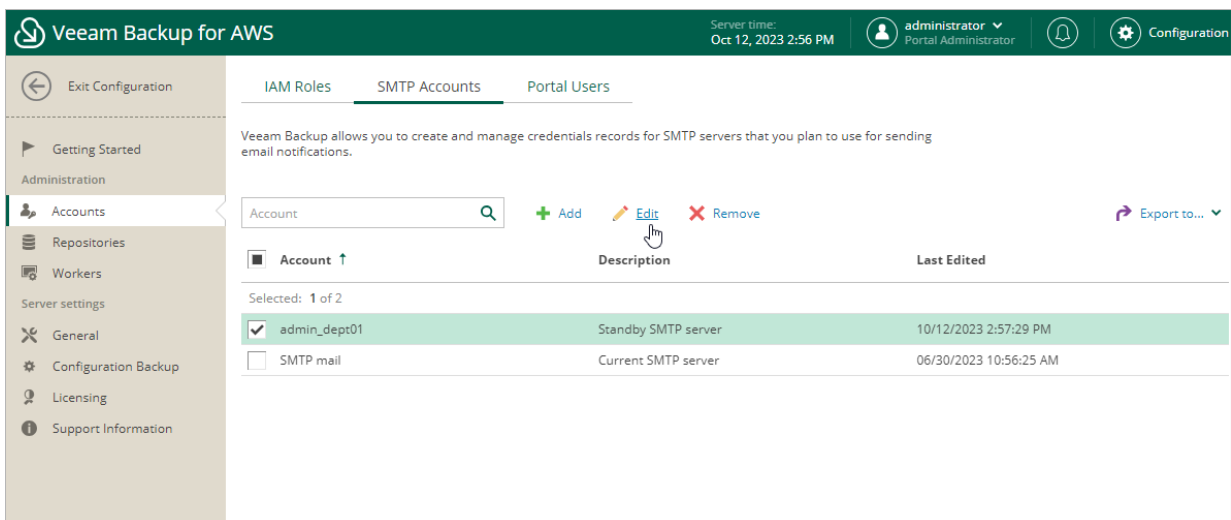
Editing SMTP Accounts

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

When updating to Veeam Backup for AWS version 5.0, note that only the TLS v1.3 certificates are now supported. Veeam Backup for AWS will automatically recreate all previously generated self-signed certificates.

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 CertificateWizard**.

- 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.

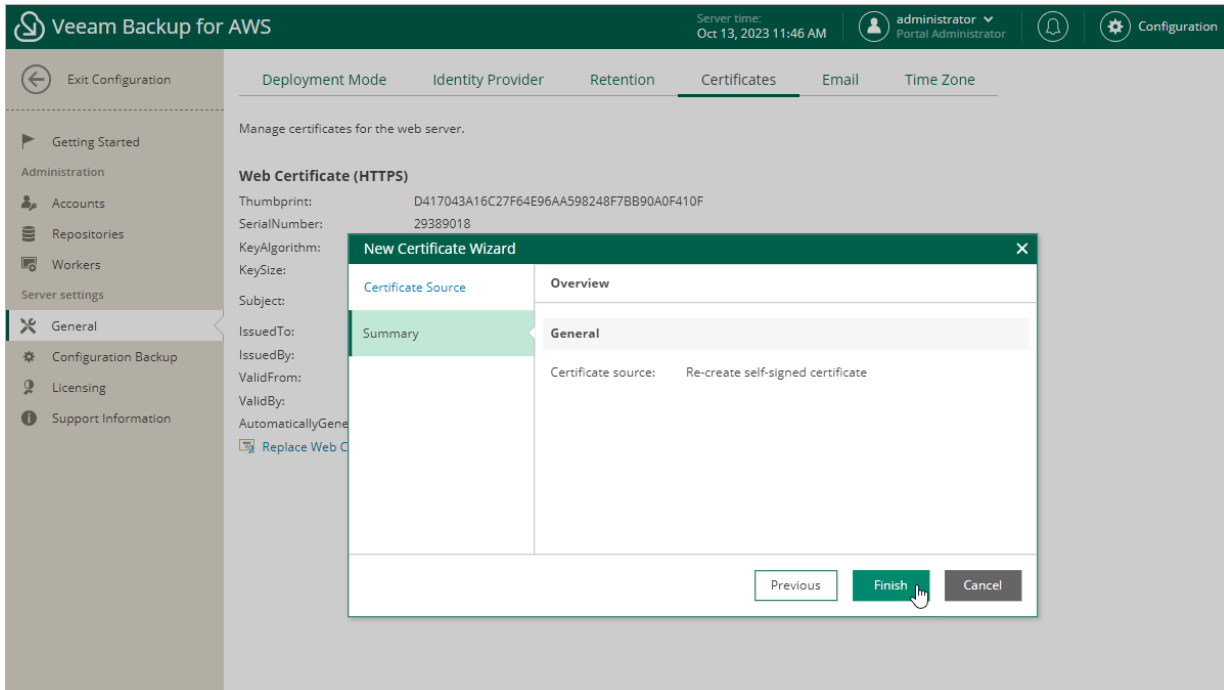
NOTE

Only .PFX and .P12 certificate files are supported.

- c. At the **Summary** step of the wizard, review summary information and click **Finish**. To allow Veeam Backup for AWS to discover the newly installed certificate, restart the backup appliance.

NOTE

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.



Changing Time Zone

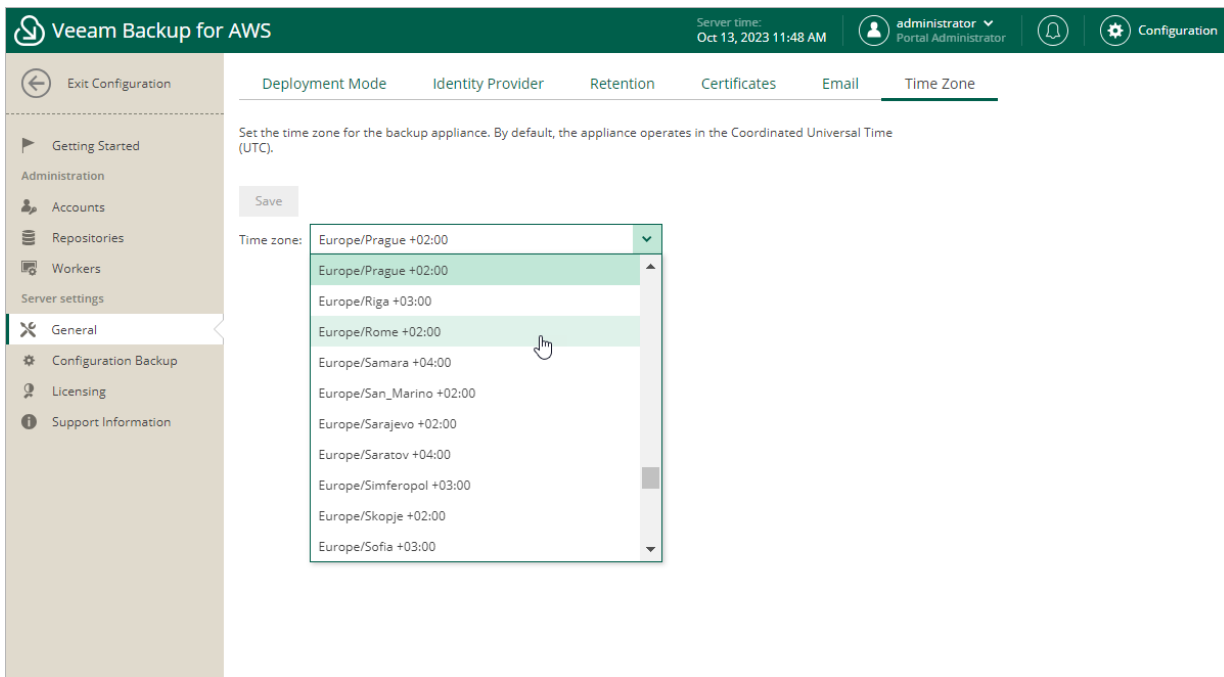
Veeam Backup for AWS runs daily reports and performs all data protection and disaster recovery operations according to the time zone set on the backup appliance. Since the backup appliance is deployed on an EC2 instance in Amazon EC2, the time zone is set to Coordinated Universal Time (UTC) by default. 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

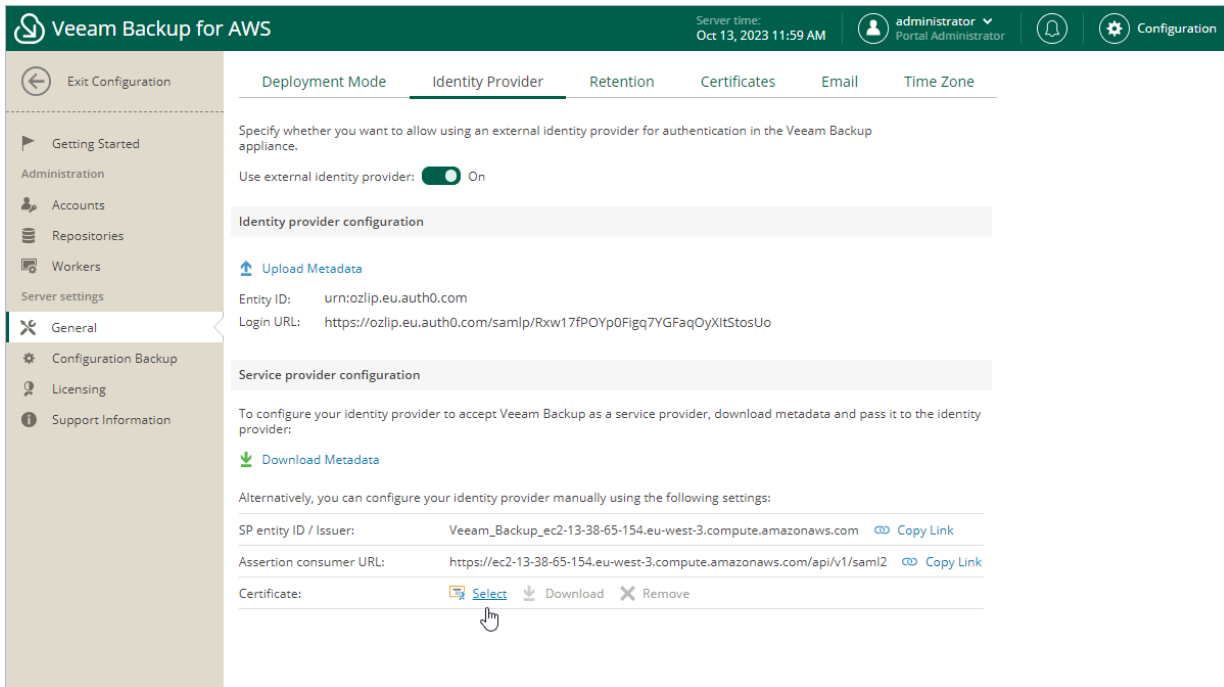
Only .PFX and .P12 certificate files are supported.

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 console. 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:** Shows 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 EC2 instances, RDS resources, DynamoDB tables, EFS file systems and VPC configurations, 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 EC2 instances, RDS resources, DynamoDB tables, EFS file systems and VPC configurations, 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.

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](#).

IMPORTANT

For Veeam Backup & Replication to be able to back up configurations of managed backup appliances, you must enable backup file encryption in the configuration backup settings.

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](#). 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.

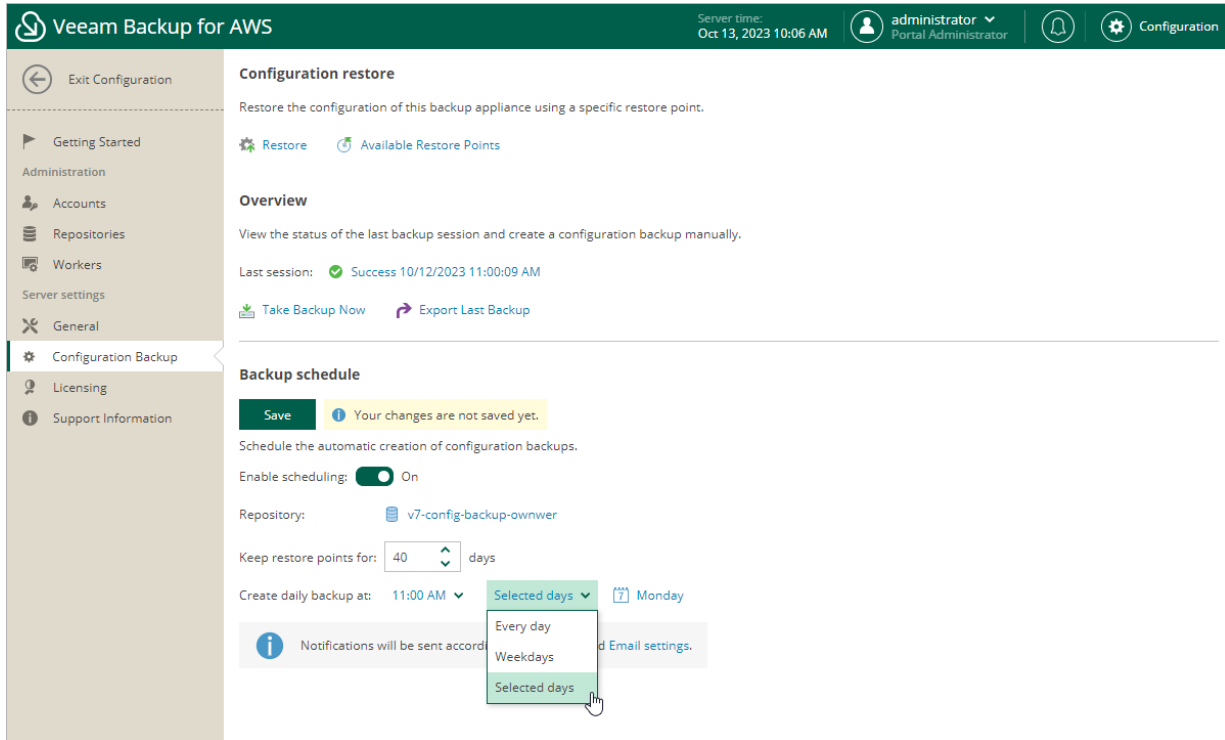
For a backup repository to be displayed in the list of available repositories, it must be added to Veeam Backup for AWS as described in section [Adding Backup Repositories](#). The list shows only backup repositories of the *S3 Standard* storage class that have encryption enabled and immutability disabled.

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**.



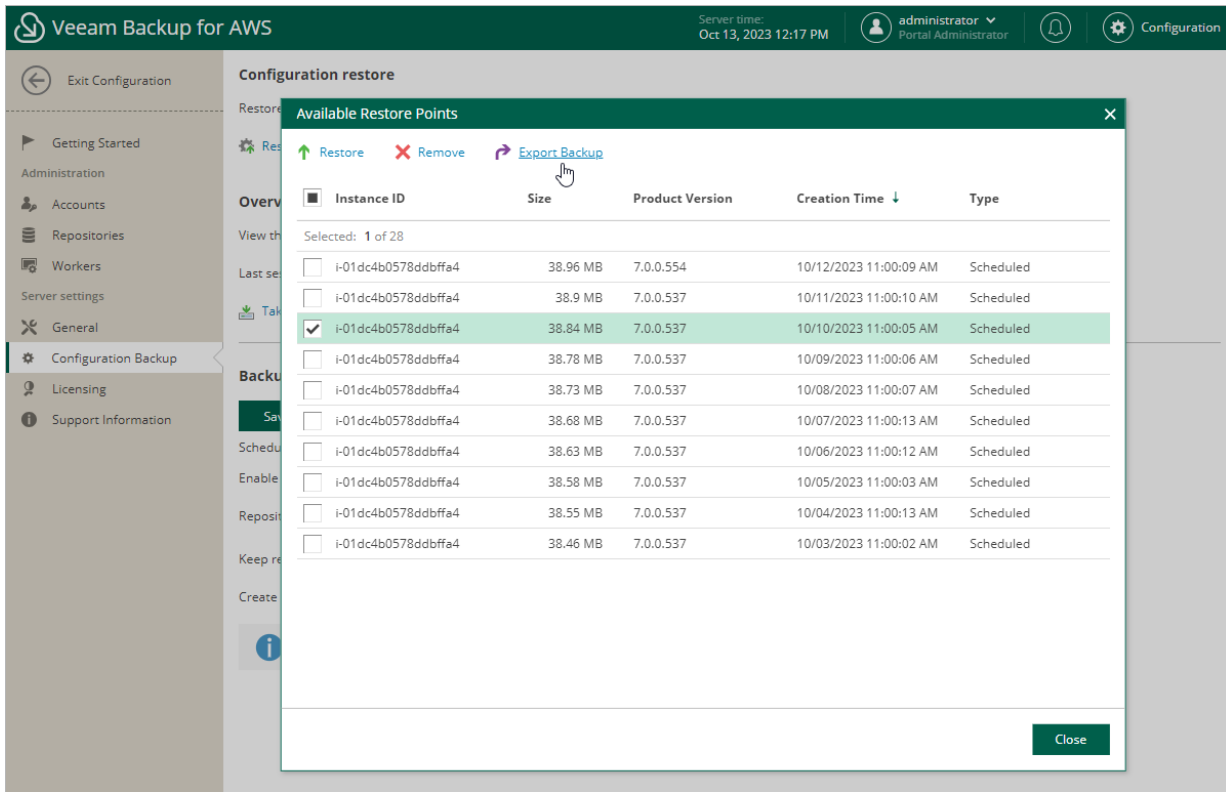
Exporting Configuration Backup Data

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](#).

Before You Begin

Before you restore configuration 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.

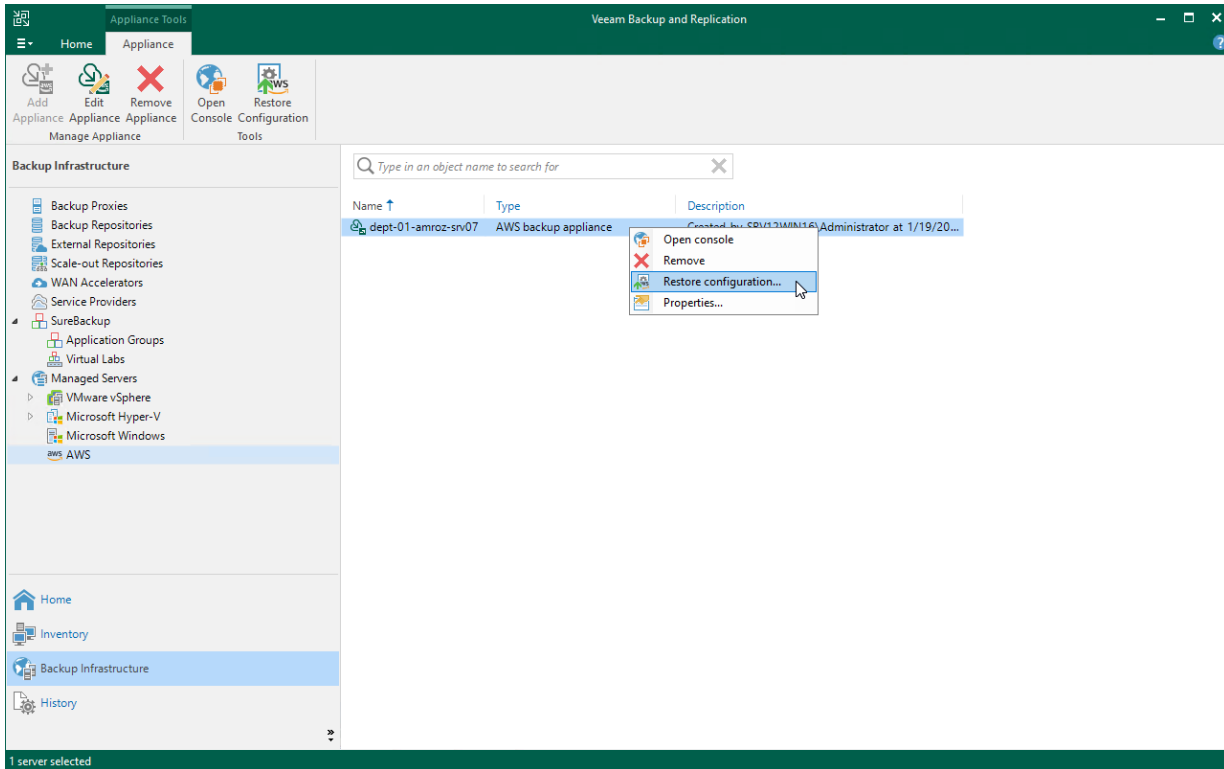
- If the backup appliance requires an upgrade, perform it before you start configuration restore. Otherwise, Veeam Backup & Replication will not be able to perform the restore operation. To learn how to upgrade appliances, see [Upgrading Appliances Using Console](#).
- If you remove the backup appliance 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 of the backup appliance to another one, 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 of the backup appliance to another one] During restore, Veeam Backup & Replication removes the 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:

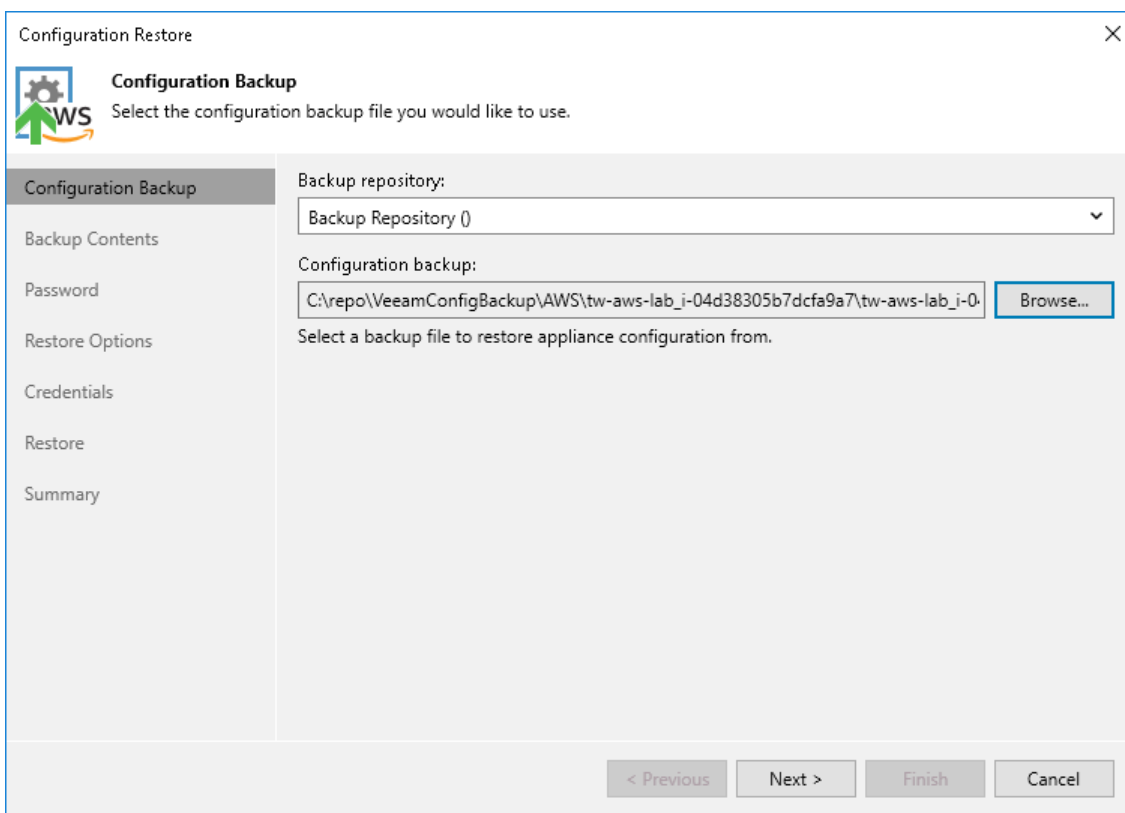
1. From the **Backup repository** list, select a repository where the configuration backup file is stored.

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' (highlighted), 'Backup Contents', 'Password', 'Restore Options', 'Credentials', 'Restore', and 'Summary'. The main area of the wizard is titled 'Configuration Backup' and contains the instruction 'Select the configuration backup file you would like to use.' Below this instruction, there are two input fields: 'Backup repository:' with a dropdown menu showing 'Backup Repository ()', and 'Configuration backup:' with a text box containing the path 'C:\repo\VeeamConfigBackup\AWS\tw-aws-lab_j-04d38305b7dcfa9a7\tw-aws-lab_j-0-'. To the right of the text box is a 'Browse...' button. Below the input fields, there is a note: '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

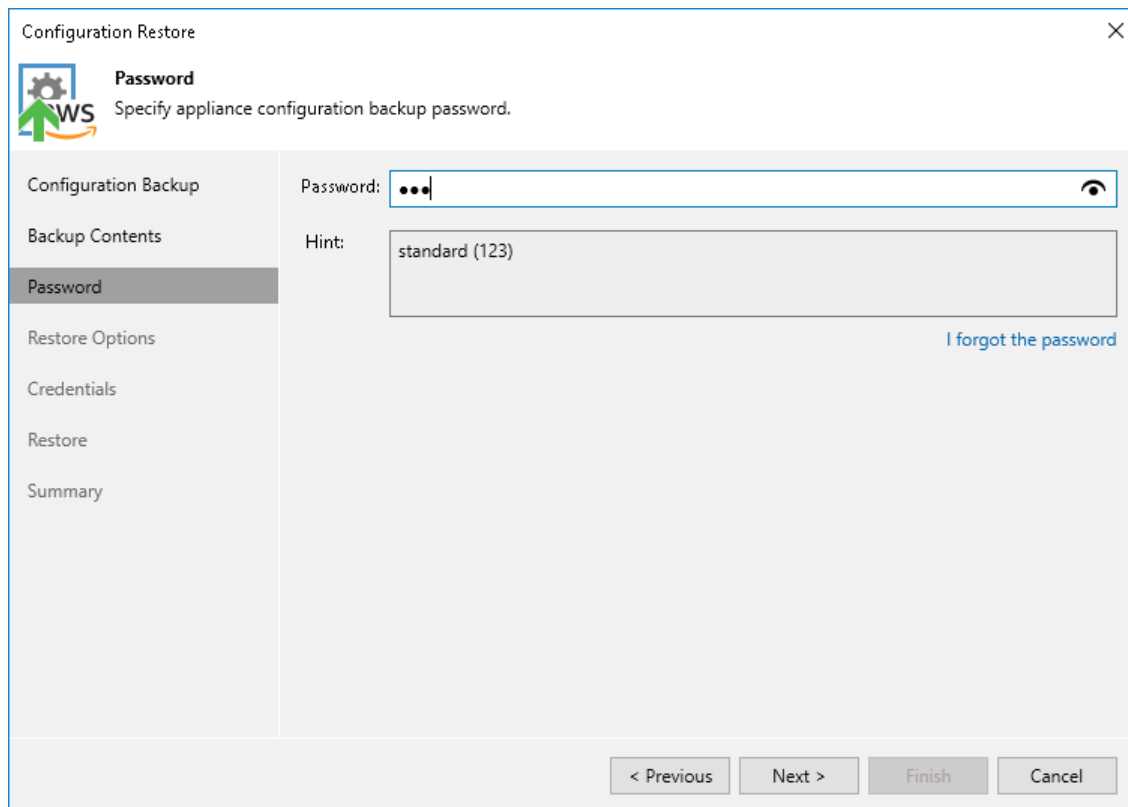
At the **Password** step of the wizard, specify a password used to encrypt the configuration backup.

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 wizard is titled 'Password' and contains the instruction 'Specify appliance configuration backup password.' There are two input fields: 'Password:' with a masked password '...' and a toggle icon to show/hide it, and 'Hint:' with the text 'standard (123)'. A blue link 'I forgot the password' is located below the hint field. At the bottom of the window are four buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'.

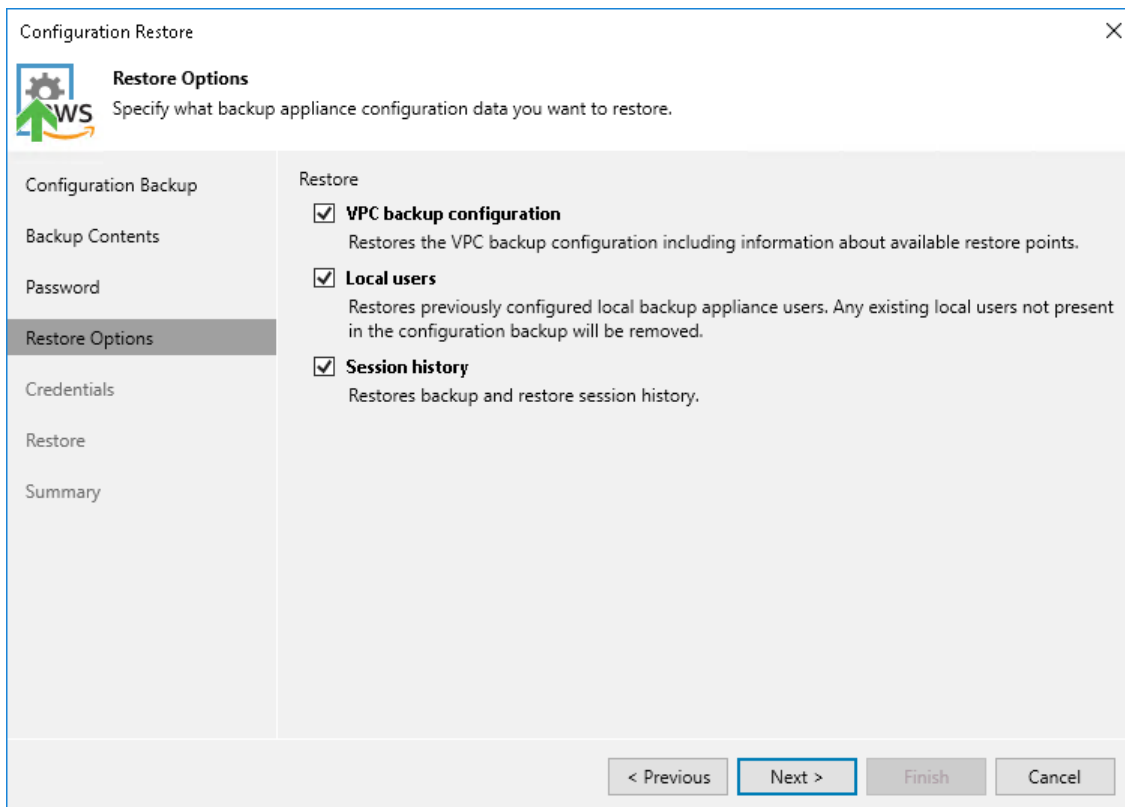
Step 5. Choose Restore Options

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]

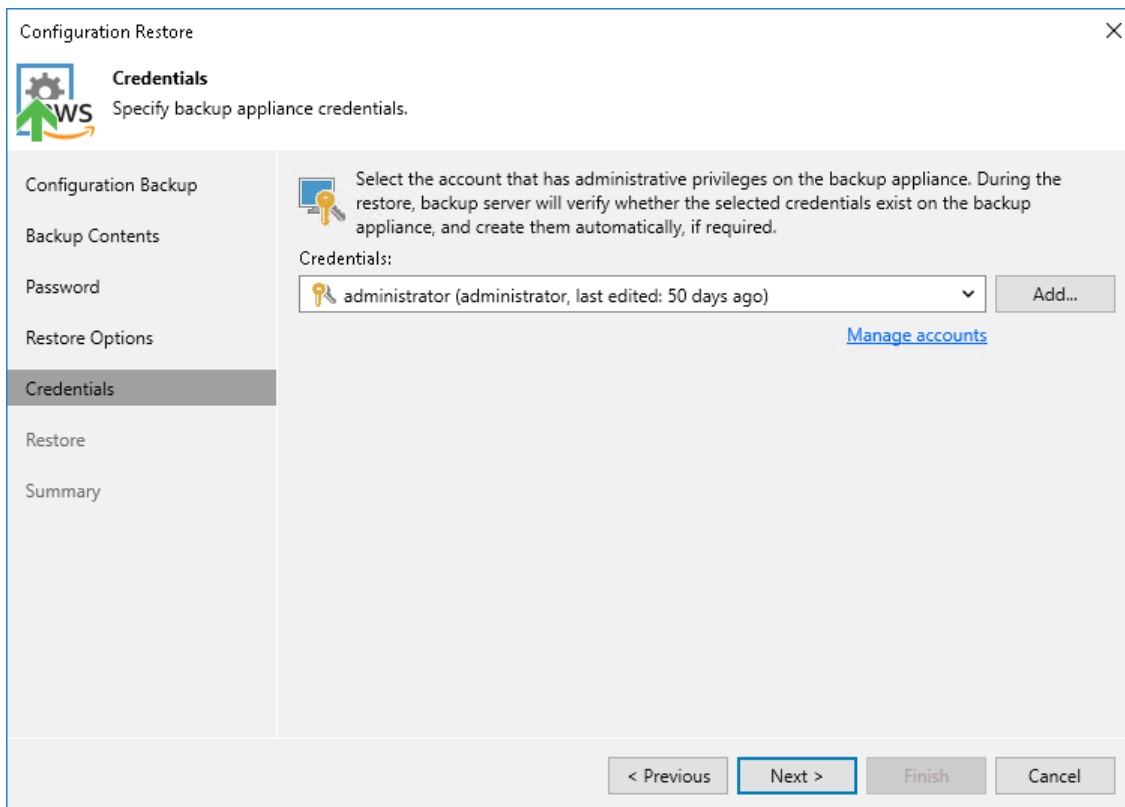
After the configuration restore process completes, Veeam Backup & Replication will try to connect to the backup appliance using credentials of the user specified [when adding the appliance](#) to the backup infrastructure. However, since you have chosen to restore all users saved to the configuration 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. If you have not added a user to the Credentials Manager beforehand, you can do it without closing the **Configuration Restore** wizard. To add a new user, click either the **Manage accounts** link or the **Add** button and specify a user name, password and description in the **Credentials** window.

IMPORTANT

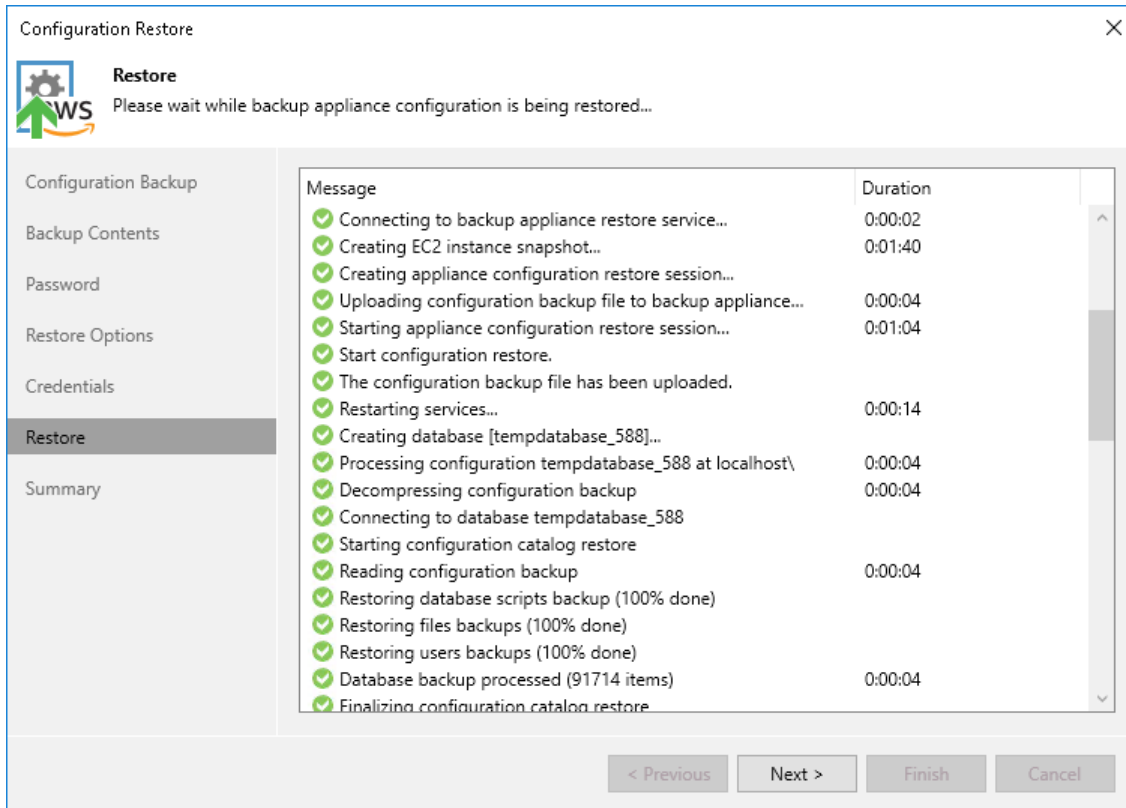
After you click **Next**, the restore process will start. You will not be able to halt the process or edit the restore settings.



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 sub-instruction: '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:

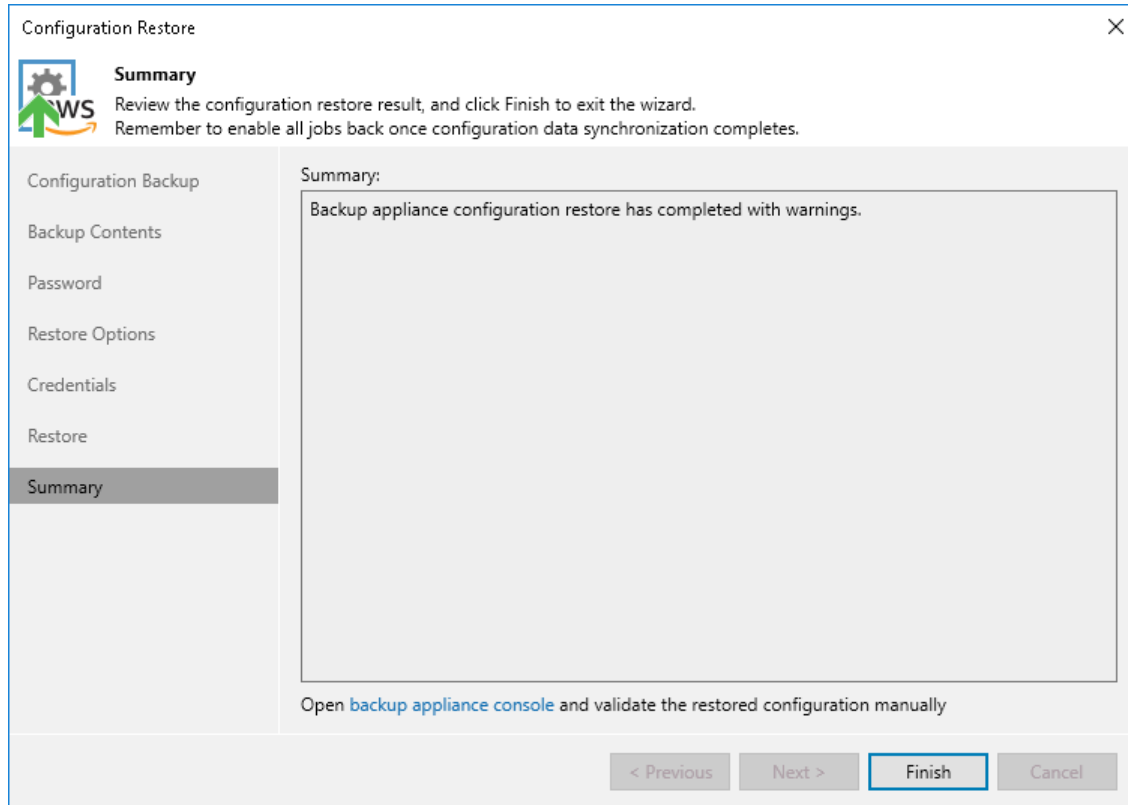
Message	Duration
✓ Connecting to backup appliance restore service...	0:00:02
✓ Creating EC2 instance snapshot...	0:01:40
✓ Creating appliance configuration restore session...	
✓ Uploading configuration backup file to backup appliance...	0:00:04
✓ Starting appliance configuration restore session...	0:01:04
✓ 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 that the restore process is complete and the user can proceed to the next step.

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](#). 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

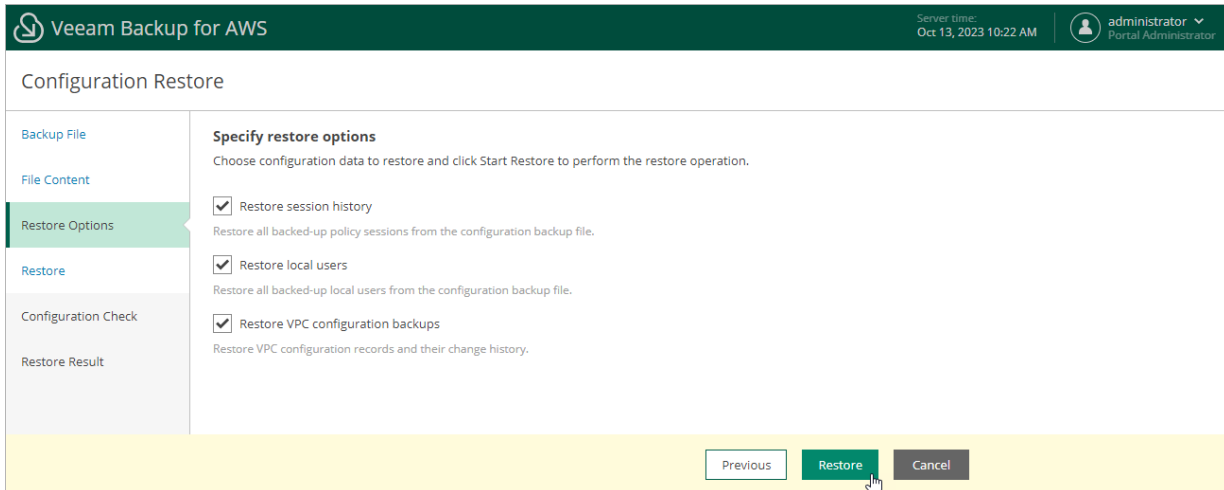
At the bottom of the wizard, there are three buttons: 'Previous', 'Next' (highlighted with a mouse cursor), and 'Cancel'.

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.

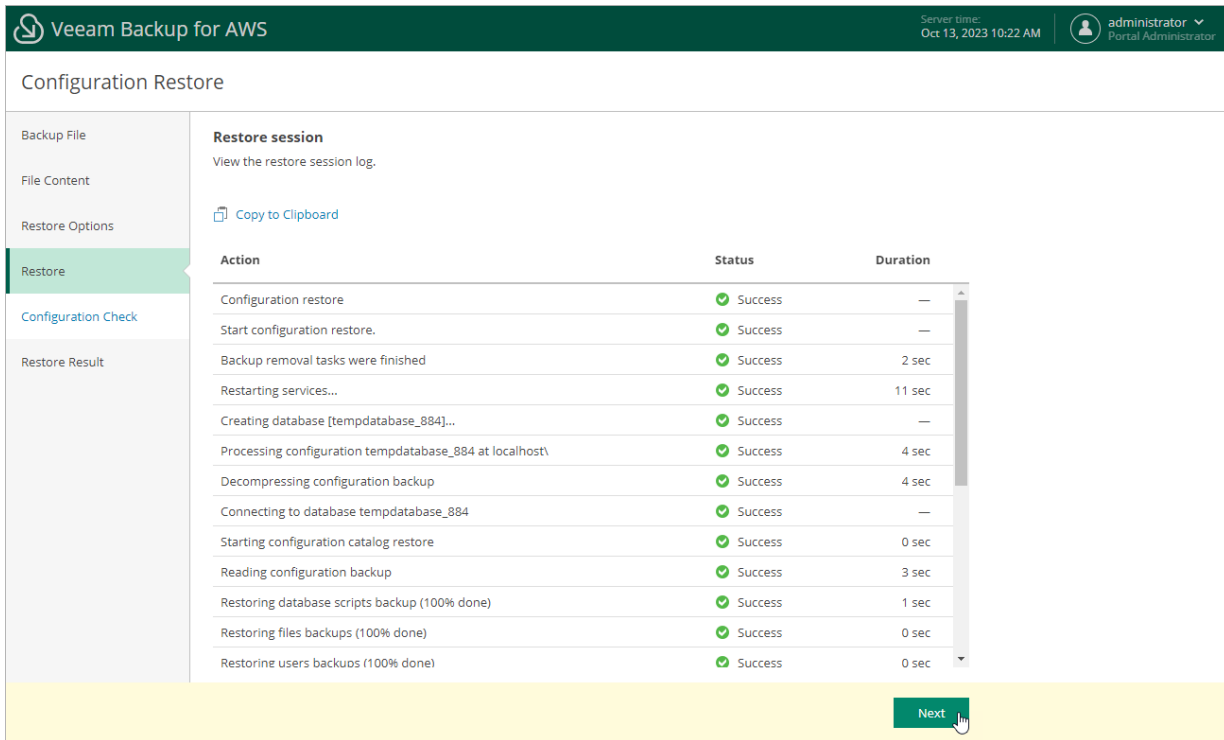
IMPORTANT

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 and a table of restore actions.

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

A "Next" button is visible at the bottom right of the main content area.

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. 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.

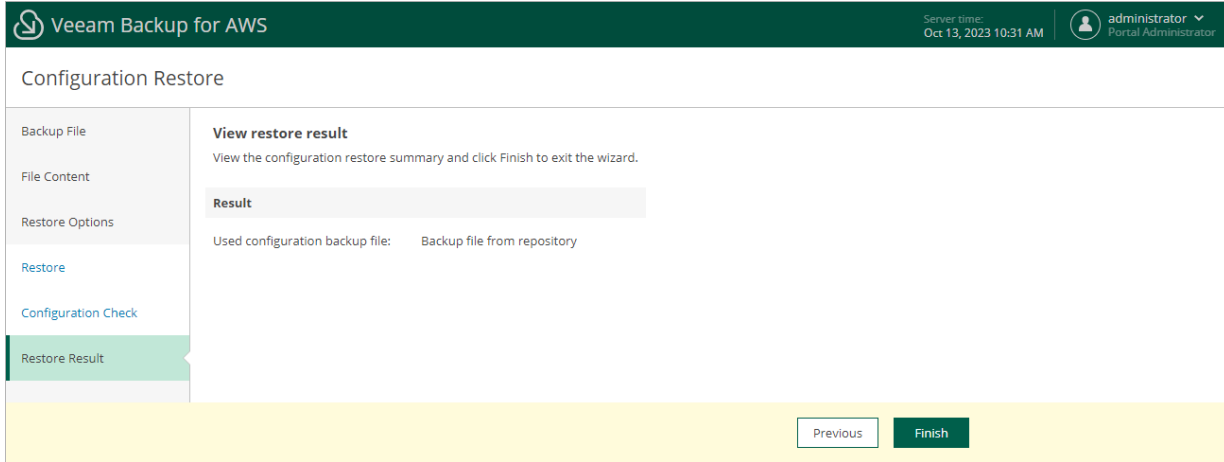
The screenshot displays the Veeam Backup for AWS Configuration Restore interface. The main window shows a list of verification steps under the 'Configuration Restore' section. The 'IAM roles' step is highlighted, indicating a '1 error found'. A table lists the IAM roles with their status and missing permissions.

IAM Role	AWS Account	Used As	Result
<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...

A 'Grant Permissions' dialog box is open, titled 'Provide temporary credentials'. It contains an information message and two input fields: 'Access key' (with the value AKIAX2DLKQIWIUFFYT4) and 'Secret key' (with masked characters). The 'Grant' button is highlighted with a mouse cursor.

Step 7. Finish Working with Wizard

At the Summary step of the wizard, click **Finish** to finalize the process of configuration data restore.



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 "Oct 13, 2023 10:31 AM", and the user "administrator Portal Administrator". The main content area is titled "Configuration Restore" and features a left-hand sidebar with the following menu items: "Backup File", "File Content", "Restore Options", "Restore", "Configuration Check", and "Restore Result" (which is highlighted in green). The main panel displays the "View restore result" section, which includes the instruction "View the configuration restore summary and click Finish to exit the wizard." Below this, a "Result" box shows the text "Used configuration backup file: Backup file from repository". At the bottom of the page, there are two buttons: "Previous" and "Finish".

Viewing Available Resources

After you create a backup policy to protect a specific type of AWS resources (EC2 instances, RDS resources, DynamoDB tables or EFS file systems), Veeam Backup for AWS rescans AWS Regions specified in the policy settings and populates the resource list on the **Resources** tab 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** tab displays AWS resources that can be protected by Veeam Backup for AWS. Each resource is represented with a set of properties, such as:

- **Instance** or **Name** – the name of the resource.
- **Instance ID** or **File System ID** – the unique identification number of the resource.
- **Instance Size**, **Source 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 where 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** tab you can also perform the following actions:

- Manually create cloud-native snapshots of RDS and EC2 instances, as well as backups of DynamoDB tables and EFS file systems. For more information, see sections [Creating EC2 Snapshots Manually](#), [Creating RDS Snapshots Manually](#), [Creating DynamoDB Backups Manually](#) and [Creating EFS 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 entire Amazon 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](#).

- Remove all cloud-native snapshots created for EC2 instances, DB instances or Aurora DB clusters manually, as well as remove all backups created for EFS file systems and DynamoDB tables 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.

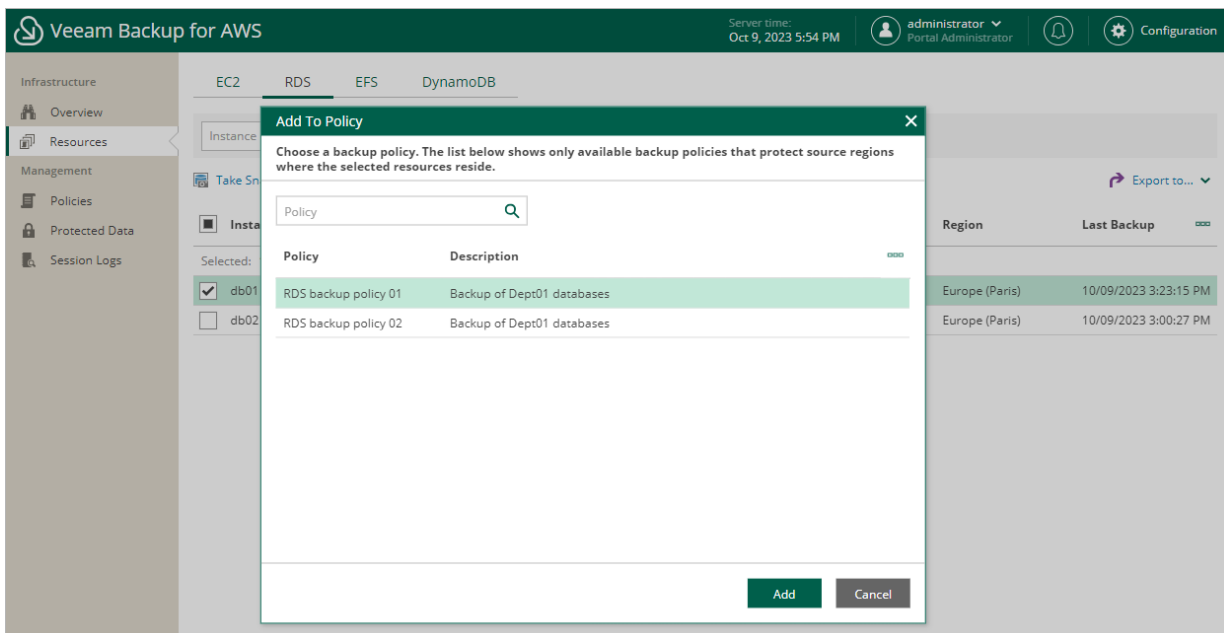
3. Click **Add to Policy**.

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 captures the whole image of the processed EC2 instance (including instance configuration, OS data, application data and so on) and DB instance at a specific point in time. The backup is saved to a backup repository in the native Veeam format.

- **Create backups and backup copies of your DynamoDB tables and EFS file systems**

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](#).

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](#).

By default, DynamoDB and 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. 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 and EFS file systems. To learn how to perform backup manually, see sections [Creating EC2 Snapshots Manually](#), [Creating RDS Snapshots Manually](#), [Creating DynamoDB Backups Manually](#), [Creating EFS Backups Manually](#).

TIP

You can perform advanced data protection operations with image-level backups from the Veeam Backup & Replication console. For details, see the [External Repository](#) section in the Veeam Backup & Replication User Guide.

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

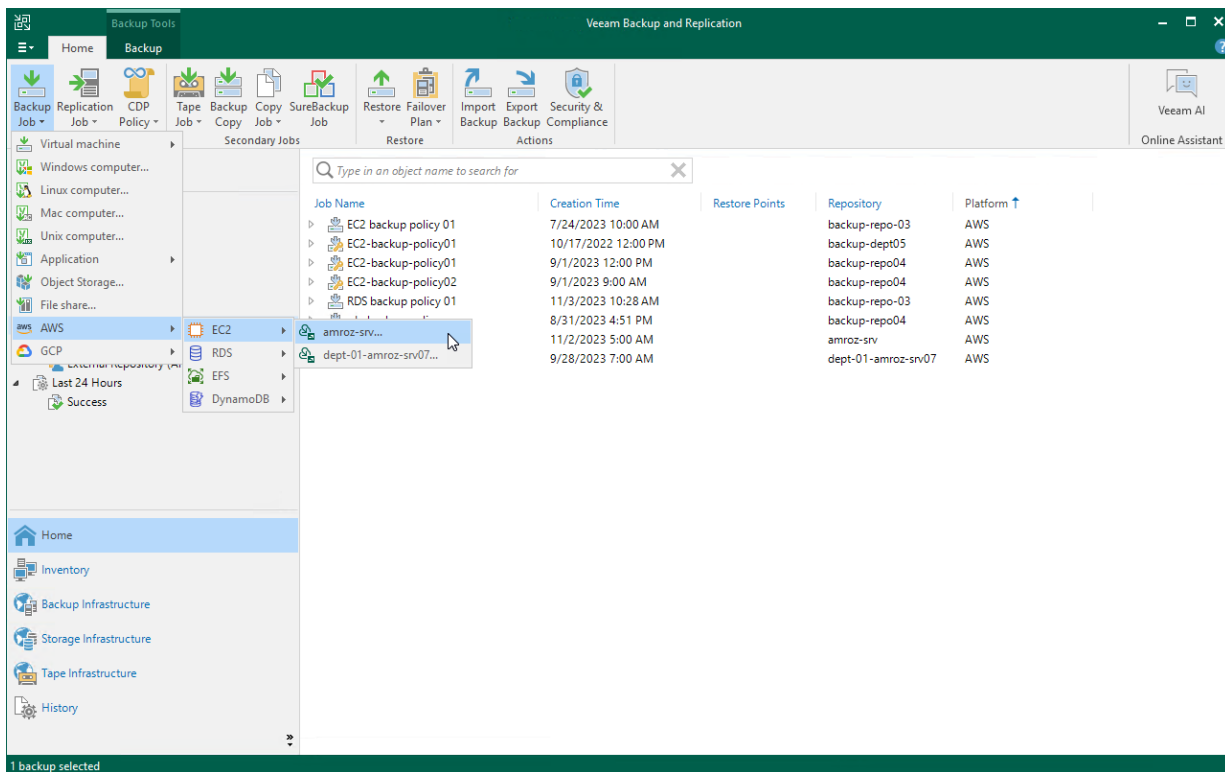
You can create backup policies in the Veeam Backup for AWS Web UI only. However, you can launch the add policy wizard directly from the Veeam Backup & Replication console – to do that, use either of the following options:

- Switch to the **Home** tab, click **Backup Job** on the ribbon, navigate to **AWS > EC2, RDS, EFS or DynamoDB**, 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 or DynamoDB**, 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** or **Add DynamoDB 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](#) or [Creating DynamoDB 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](#).



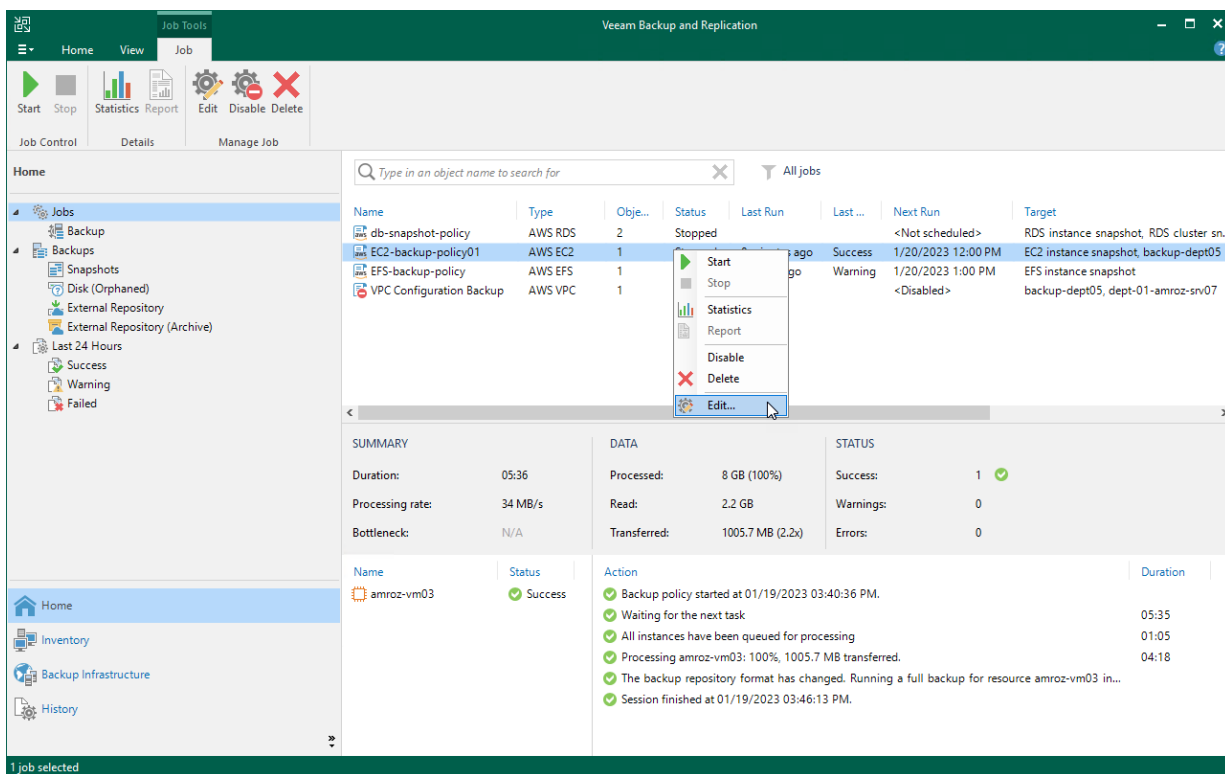
Editing Backup Policy Settings

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 EFS Backup Policies](#), [Creating DynamoDB Backup Policies](#) or [Editing VPC Configuration Backup Policy](#).



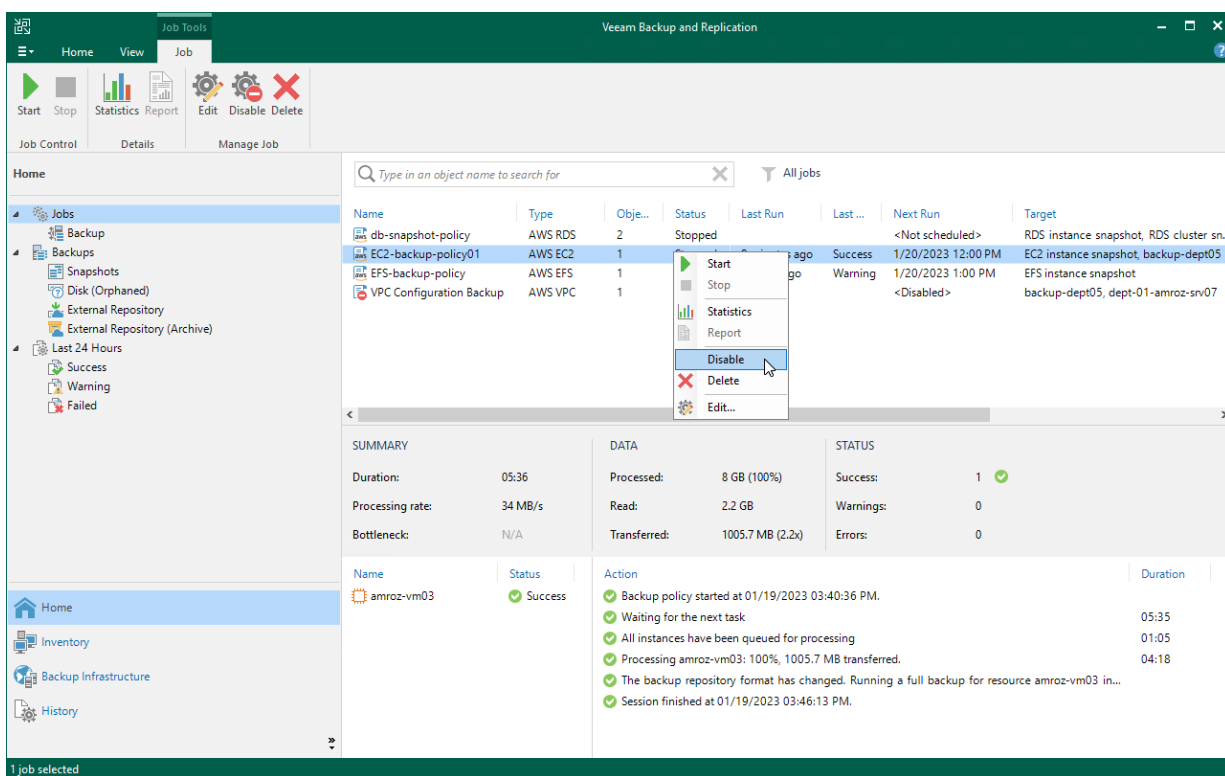
Enabling and Disabling Backup Policies

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**.



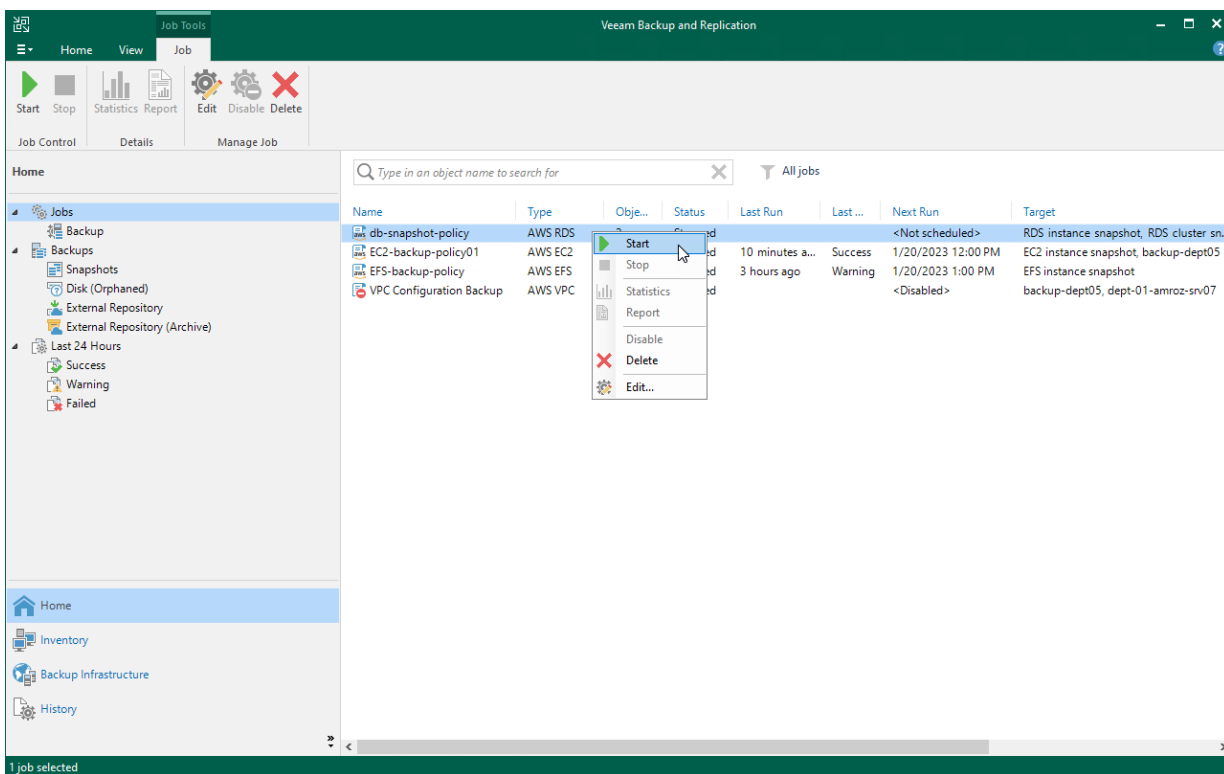
Starting and Stopping Backup Policies

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.

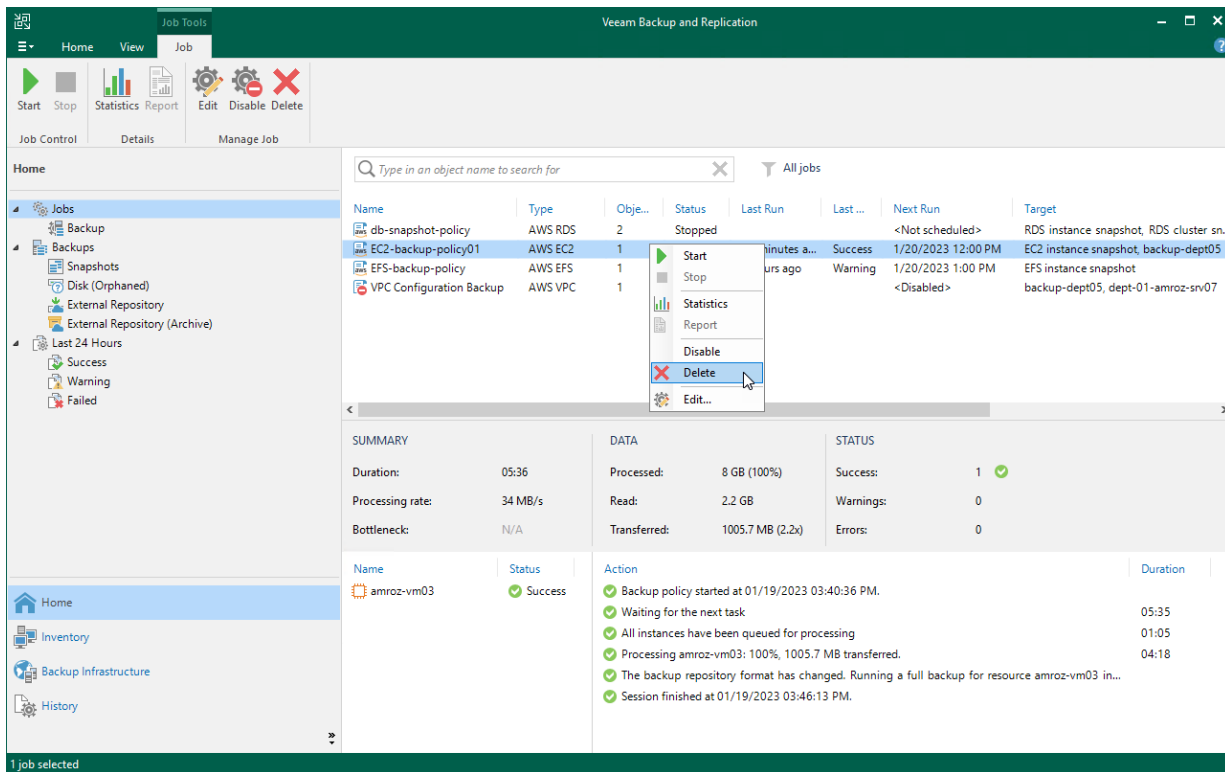
To delete a backup policy, do the following:

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

If you delete a backup policy from Veeam Backup & Replication, the policy is automatically deleted from the backup appliance as well.



The screenshot displays the Veeam Backup and Replication console interface. The 'Home' view is active, showing a list of backup jobs. A context menu is open over the 'EC2-backup-policy01' job, with the 'Delete' option highlighted. The console also shows a summary of the selected job's performance and a detailed action log.

Name	Type	Obj...	Status	Last Run	Last ...	Next Run	Target
db-snapshot-policy	AWS RDS	2	Stopped			<Not scheduled>	RDS instance snapshot, RDS cluster sn...
EC2-backup-policy01	AWS EC2	1	Running	minutes a...	Success	1/20/2023 12:00 PM	EC2 instance snapshot, backup-dept05
EFS-backup-policy	AWS EFS	1	Warning	urs ago	Warning	1/20/2023 1:00 PM	EFS instance snapshot
VPC Configuration Backup	AWS VPC	1	Disabled			<Disabled>	backup-dept05, dept-01-amroz-srv07

Summary	Data	Status
Duration: 05:36	Processed: 8 GB (100%)	Success: 1 ✓
Processing rate: 34 MB/s	Read: 2.2 GB	Warnings: 0
Bottleneck: N/A	Transferred: 1005.7 MB (2.2x)	Errors: 0

Name	Status	Action	Duration
amroz-vm03	Success ✓	<ul style="list-style-type: none">Backup policy started at 01/19/2023 03:40:36 PM.Waiting for the next taskAll instances have been queued for processingProcessing amroz-vm03: 100%, 1005.7 MB transferred.The backup repository format has changed. Running a full backup for resource amroz-vm03 in...Session finished at 01/19/2023 03:46:13 PM.	05:35 01:05 04:18

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.

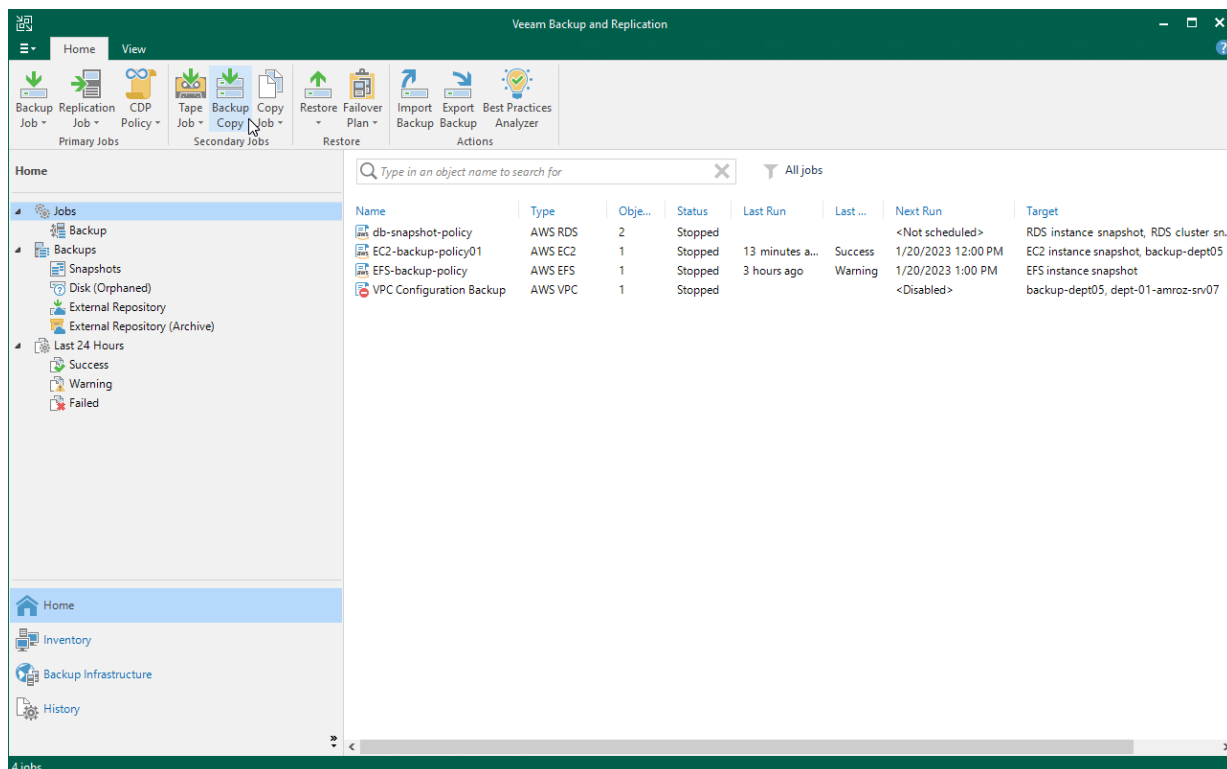
Backup-copy is a job-driven process. Veeam Backup & Replication fully automates the backup copy process and lets you specify retention settings to maintain the desired number of restore points, as well as full backups for archival purposes. For more information on the backup copy functionality, see the 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](#).



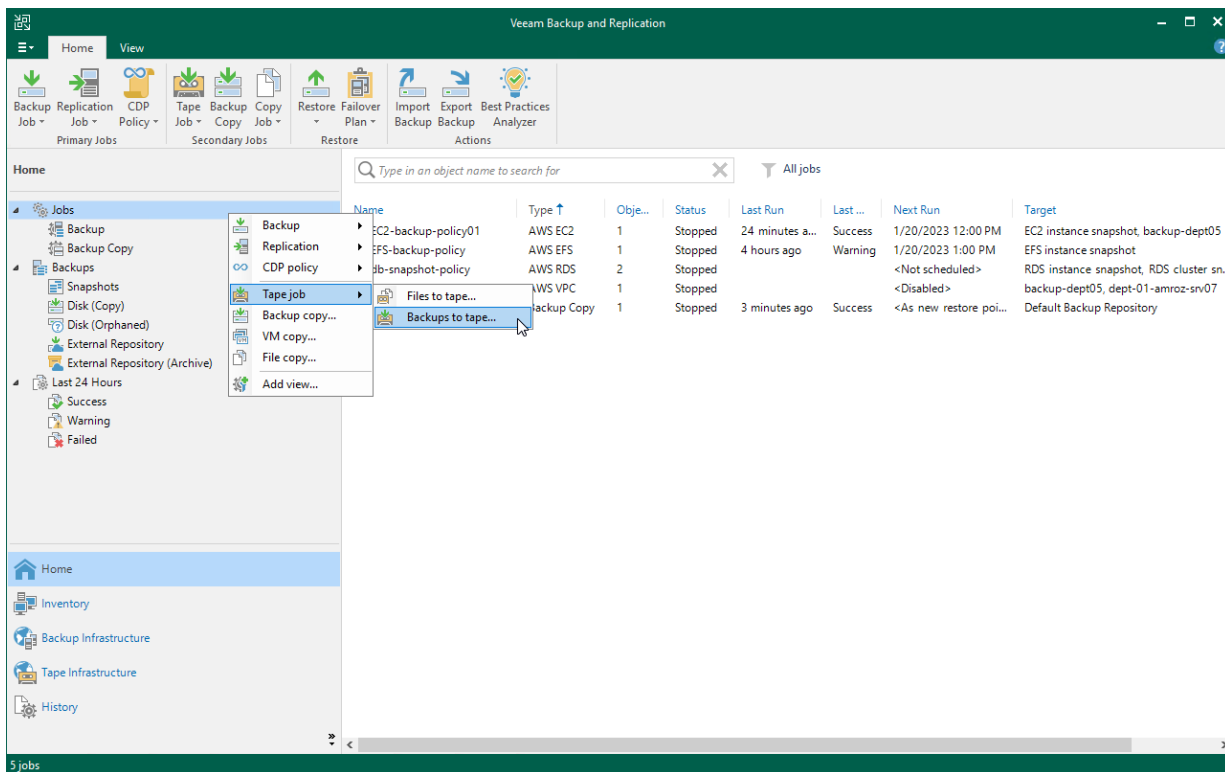
Copying Backups to Tapes

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 [global notification settings](#) first.
- If you plan to create transactionally consistent backups of EC2 instances, check the requirements for [application-aware processing](#) and [guest scripting](#).

For EC2 instances residing in any of the regions added to the backup policies, you can also [take cloud-native snapshots manually](#) when needed.

IMPORTANT

In Veeam Backup for AWS, you can protect only EC2 instances that run in VPCs. EC2-Classic instances are not supported. For details, see [this Veeam KB article](#).

Creating EC2 Backup Policies

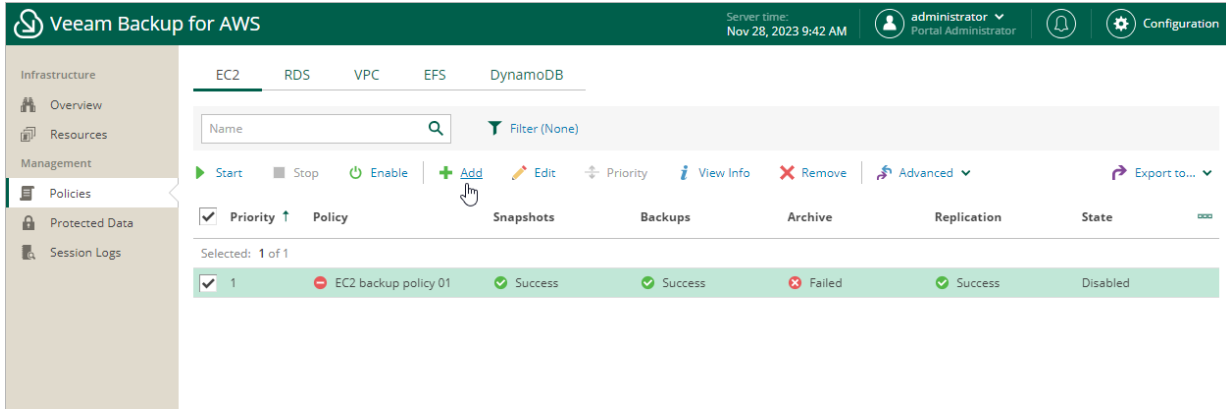
To create an EC2 backup policy, do the following:

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 for protecting 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 includes a top navigation bar with the Veeam logo, server time (Nov 28, 2023 9:44 AM), user information (administrator, Portal Administrator), and a Configuration icon. The main header shows a back arrow, the title 'Add EC2 Policy', and the cost 'Cost: N/A'. A left sidebar contains a navigation menu with 'Info' selected, and other options: Sources, Guest Processing, 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.' Below this are two input fields: 'Name:' with the value 'EC2 backup policy 02' and 'Description:' with the value 'EC2 backup policy for dept-01'. At the bottom of the form 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 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 add an IAM role, click **Add** and complete the [Add IAM Role](#) wizard.

IMPORTANT

It is recommended that you check whether the selected IAM role has all the required permissions 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 box is open. The dialog box displays a green checkmark and the message 'Your account meets the required permissions.' Below this, there are buttons for 'Grant', 'Recheck', and 'Export Missing Permissions'. A table lists various permissions and their status:

Type	Status	Missing Permissions
IAM permissions	Passed	—
SQS permissions	Passed	—
EVENTS permissions	Passed	—
SNS permissions	Passed	—
EBS permissions	Passed	—
EC2 permissions	Passed	—
SSM permissions	Passed	—
SERVICEQUOTAS permissions	Passed	—
KMS permissions	Passed	—

The dialog box also includes a 'Close' button at the bottom.

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 'Add EC2 Policy' wizard in Veeam Backup for AWS. The 'Sources' step is active, and the 'Specify source settings' section is expanded to 'Regions'. The 'Choose regions' dialog is open, showing a list of 17 available regions. The 'Europe (Milan)' region is selected in the 'Available Regions' list. The 'Add' button is highlighted, and the 'Europe (Paris)' region is listed in the 'Selected Regions' list. The 'Apply' button is visible at the bottom of the dialog.

Server time: Nov 28, 2023 9:45 AM

administrator Portal Administrator

Configuration

Cost: N/A

Info

Sources

Guest Processing

Targets

Schedule

Tags

General Settings

Cost Estimation

Summary

Specify source settings

Select an IAM role to use, regions to cover and resources to protect. This selection automatically changes the policy scope when tags are used.

IAM role

The selected IAM role must have sufficient permissions to create and manage resources. For more information on required permissions, see the User Guide.

IAM role: Backup role (role to perform backup operations an...

Regions

Specify one or more regions.

Choose regions...

Resources

Specify resources to protect or exclude.

Choose resources to protect...

Choose resources to exclude...

Volumes

Specify volumes that will be excluded from the backup.

Exclude volumes: Off

Choose regions

Available Regions (17)

- Asia Pacific (Mumbai)
- Asia Pacific (Osaka)
- Asia Pacific (Seoul)
- Asia Pacific (Singapore)
- Asia Pacific (Sydney)
- Asia Pacific (Tokyo)
- Canada (Central)
- Europe (Frankfurt)
- Europe (Ireland)
- Europe (London)
- Europe (Milan)
- Europe (Stockholm)
- South America (Sao Paulo)
- US East (N. Virginia)
- US East (Ohio)
- US West (N. California)
- US West (Oregon)

Selected Regions (1)

- Europe (Paris)

Apply Cancel

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 Veeam Backup for AWS interface. The main window is titled "Add EC2 Policy" and has a "Cost: N/A" indicator. The left sidebar contains navigation options: Info, Sources, Guest Processing, Targets, Schedule, Tags, General Settings, Cost Estimation, and Summary. The "Specify source settings" panel is active, showing fields for IAM role, Regions (2 regions selected), Resources (with "Choose resources to protect..." and "Choose resources to exclude..." buttons), and Volumes (with an "Exclude volumes" toggle set to "Off").

The "Choose resources to protect" dialog is open, showing two radio button options: "All resources" and "Protect only following resources" (which is selected). Below these are dropdowns for "Type" (set to "Instance") and "Name or ID" (set to "amroz-vm05 (i-0ecc7291291e1f49a)"). A "Protect" button is visible next to the "Name or ID" field. A search bar and a "Browse to select specific resources from the global list..." link are also present. Below the search bar, it says "Protected resources (2)".

Item ↑	ID	Value	Region
Selected: 0 of 2			
amroz-vm03	i-0a4daf7f697f8321e	—	Europe (Paris)
amroz-vm04	i-05e46f817f6bca045	—	Europe (Paris)

At the bottom of the dialog are "Apply" and "Cancel" buttons.

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 'Specify source settings' panel is active, showing options for IAM role, regions, and resources. The 'Volumes' section is expanded, and the 'Exclude volumes' toggle is turned on. A modal dialog titled 'Choose volumes to exclude' is open, allowing the user to select specific volumes to exclude from the backup. The dialog includes a search bar, a list of excluded resources, and buttons for 'Apply' and 'Cancel'.

Specify source settings

Select an IAM role to use, regions to cover and resources to protect. The selection that automatically changes the policy scope when tags are selected.

IAM role

The selected IAM role must have sufficient permissions to create snapshots. For more information on required permissions, see the User Guide.

IAM role: Backup role (role to perform backup operations and...)

Regions

Specify one or more regions.

2 regions selected

Resources

Specify resources to protect or exclude.

2 resources will be protected

Choose resources to exclude...

Volumes

Specify volumes that will be excluded from the backup.

Exclude volumes: On

Choose volumes to exclude

Exclude system volumes: Off

Exclude specific volumes

Type: Volume Volume ID: vol-01b2b116a1d749fd9 [Exclude](#)

Browse to select specific resources from the global list...

Excluded resources (1)

Item	ID	Value	Region
—	vol-05c6318e5eabe4c8e	—	Europe (Paris)

Selected: 0 of 1

Apply Cancel

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*.

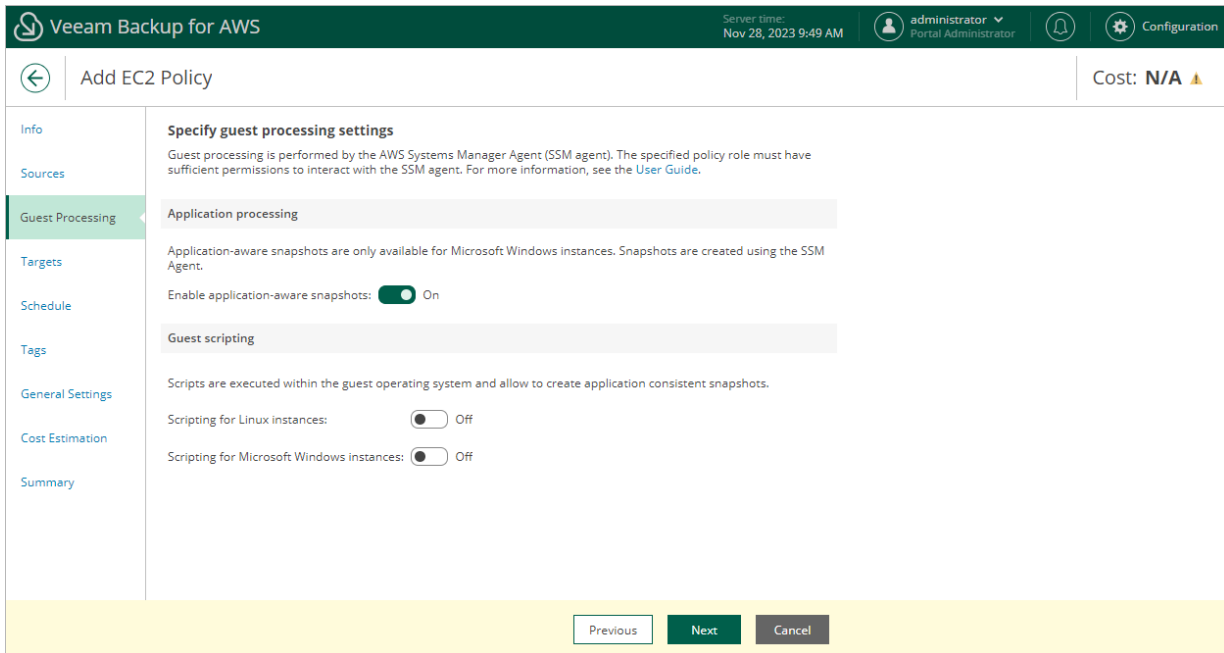
Limitations and Requirements for Application-Aware Processing

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 shows the Veeam Backup for AWS interface. At the top, the server time is Nov 28, 2023 9:50 AM, and the user is administrator (Portal Administrator). The main heading is 'Add EC2 Policy' with a cost of N/A. The left sidebar shows navigation options: Info, Sources, Guest Processing (selected), Targets, Schedule, Tags, General Settings, Cost Estimation, and Summary. The 'Guest Processing' section is expanded to show 'Specify guest processing settings' and 'Guest scripting'. The 'Guest scripting' section has a toggle for 'Scripting for Linux instances' set to 'On' and a warning 'Script settings are not configured...'. The 'Specify scripting settings for Linux instances' dialog box is open, showing:

- Pre-snapshot script:** Path in guest: /scripts/pre-snapshot.cmd; Arguments: %policy% %instance%
- Post-snapshot script:** Path in guest: /scripts/post-snapshot.cmd; Arguments: %instance%
- Additional options:**
 - 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

 The dialog has 'Apply' and 'Cancel' buttons at the bottom.

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

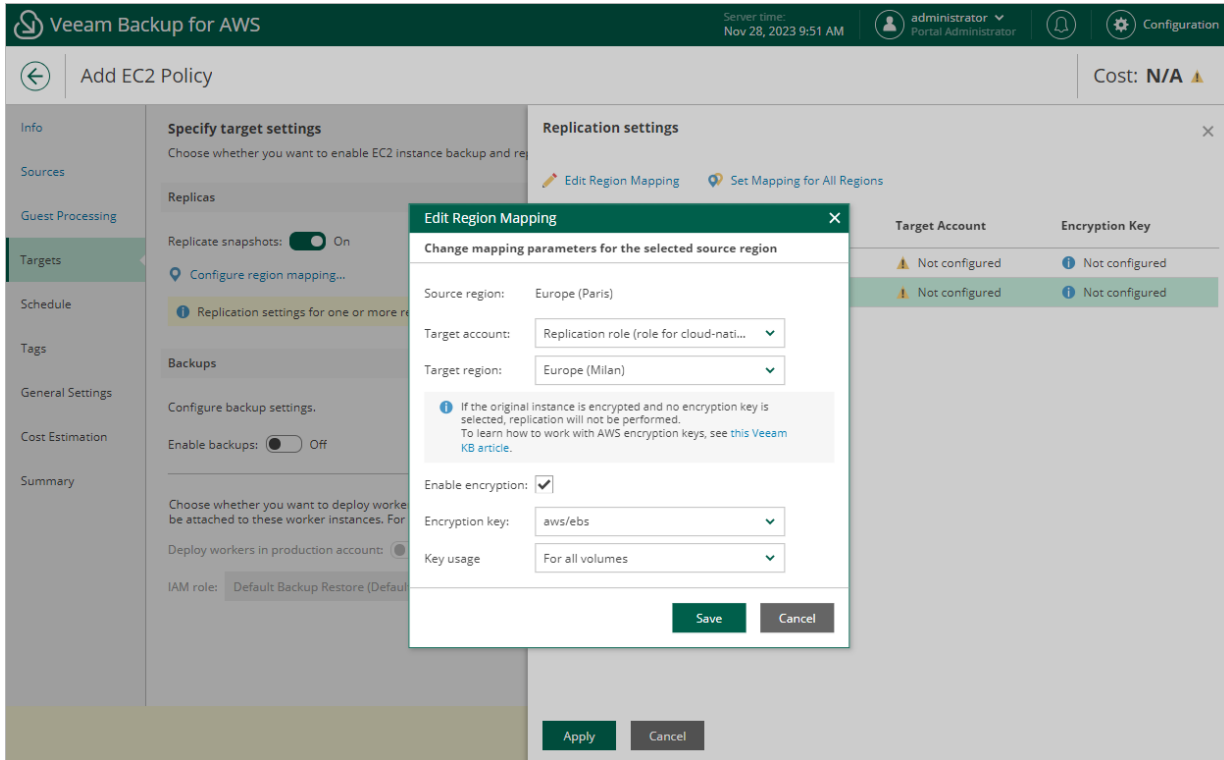
If the original EBS volume is encrypted, you must enable encryption for replicated snapshots, otherwise, the replication process will fail.

- 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](#). 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](#). 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 launches worker instances used to perform backup operations in the backup account. However, you can instruct Veeam Backup for AWS to launch worker instances in a production account – that is, an account to which the processed instances 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 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 *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 add an IAM role, 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](#).

The screenshot shows the 'Add EC2 Policy' wizard in Veeam Backup for AWS, specifically the 'Targets' step. The interface includes a top navigation bar with the Veeam logo, server time (Nov 28, 2023 9:53 AM), and user information (administrator, Portal Administrator). The left sidebar contains navigation options: Info, Sources, Guest Processing, Targets (selected), Schedule, Tags, General Settings, Cost Estimation, and Summary. 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 status indicator 'Mapping for 1 region is configured'.
- Backups:** A section for configuring backup settings, including a toggle for 'Enable backups' set to 'On', and storage locations: 'Backups will be stored in: backup-repo-03' and 'Archives will be stored in: backup-repo-02'. An information icon indicates it is recommended to use 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' and an 'IAM role' dropdown menu currently showing 'Production worker role (role to launch worker insta...'. There are '+ Add' and 'Check Permissions' buttons next to the dropdown.

At the bottom of the wizard, there are three buttons: 'Previous', 'Next', and 'Cancel'.

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 and RDS 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' configuration interface. On the left, a navigation pane includes 'Info', 'Sources', 'Guest Processing', 'Targets', 'Schedule' (selected), 'Tags', 'General Settings', 'Cost Estimation', and 'Summary'. The main area is divided into 'Specify scheduling options' and 'Create daily schedule'. Under 'Specify scheduling options', 'Daily schedule' is toggled 'On', while 'Weekly', 'Monthly', and 'Yearly' schedules are 'Off'. The 'Create daily schedule' window shows a 24-hour grid with 'Snapshots' every hour (Total: 24), 'Replicas' every 12 hours (Total: 2), and 'Backups' every 24 hours (Total: 1). Below the grid, 'Run at' is set to 'Every day'. The 'Daily retention' section has 'Snapshots to keep' at 24, 'Replicas to keep' at 2, and 'Keep backups for' at 14 days. 'Apply' and 'Cancel' buttons are at the bottom.

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](#).

- 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 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 and RDS 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](#).

- To save changes made to the backup policy settings, click **Apply**.

The screenshot displays the Veeam Backup for AWS configuration interface for adding an EC2 policy. The top navigation bar includes the Veeam logo, server time (Nov 28, 2023 10:01 AM), and user information (administrator, Portal Administrator). The main content area is titled 'Add EC2 Policy' and shows a cost of \$2.09. The interface is divided into three main sections:

- Specify scheduling options:** This section allows users to configure the frequency of backups. The 'Daily schedule' and 'Weekly schedule' are both turned on. The 'Monthly schedule' and 'Yearly schedule' are turned off. The 'Create restore points at' is set to 12:00 AM.
- Create weekly schedule:** This section allows users to specify how often the policy must produce snapshots, replicas, and backups. It includes a grid for selecting days of the week (Sun, Mon, Tue, Wed, Thu, Fri, Sat) for snapshots, replicas, and backups. The 'Creation' toggle is set to 'On'.
- Weekly retention:** This section allows users to specify the number of restore points to keep. The 'Snapshots to keep' is set to 6, 'Replicas to keep' is set to 4, and 'Keep backups for' is set to 1 month.

The 'Apply' button is highlighted, indicating that the user is ready to save the changes.

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](#).

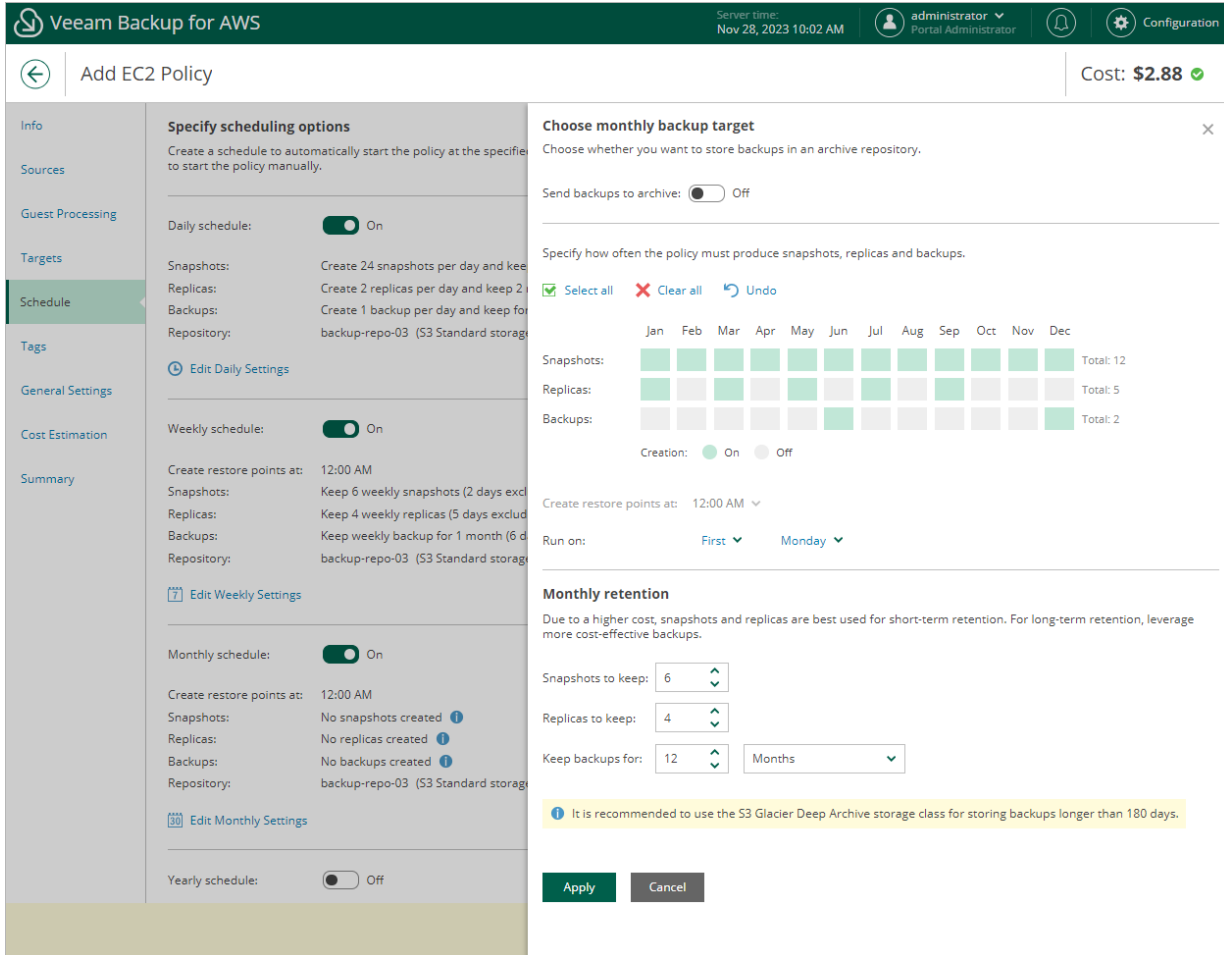
IMPORTANT

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 displays the Veeam Backup for AWS configuration interface. The main window is titled 'Add EC2 Policy' and shows a sidebar with navigation options like Info, Sources, Guest Processing, Targets, Schedule (selected), Tags, General Settings, Cost Estimation, and Summary. The 'Schedule' section is expanded, showing three schedule types: Daily, Weekly, and Monthly, each with a toggle switch set to 'On'. The 'Create yearly schedule' dialog is open, with the following settings: 'Send backups to archive' is turned on; 'Yearly schedule is applied only to image-level backups. Specify for how many years the policy must keep backup files.'; 'Create restore points on' is set to 'First' of 'Monday' of 'June' at '12:00 AM'; and 'Keep archives for' is set to '2' years. There are 'Apply' and 'Cancel' buttons at the bottom of the dialog. The top right of the console shows the server time as 'Nov 28, 2023 10:03 AM' and the user as 'administrator Portal Administrator'. A cost estimation of '\$9.01' is shown in the top right corner.

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.

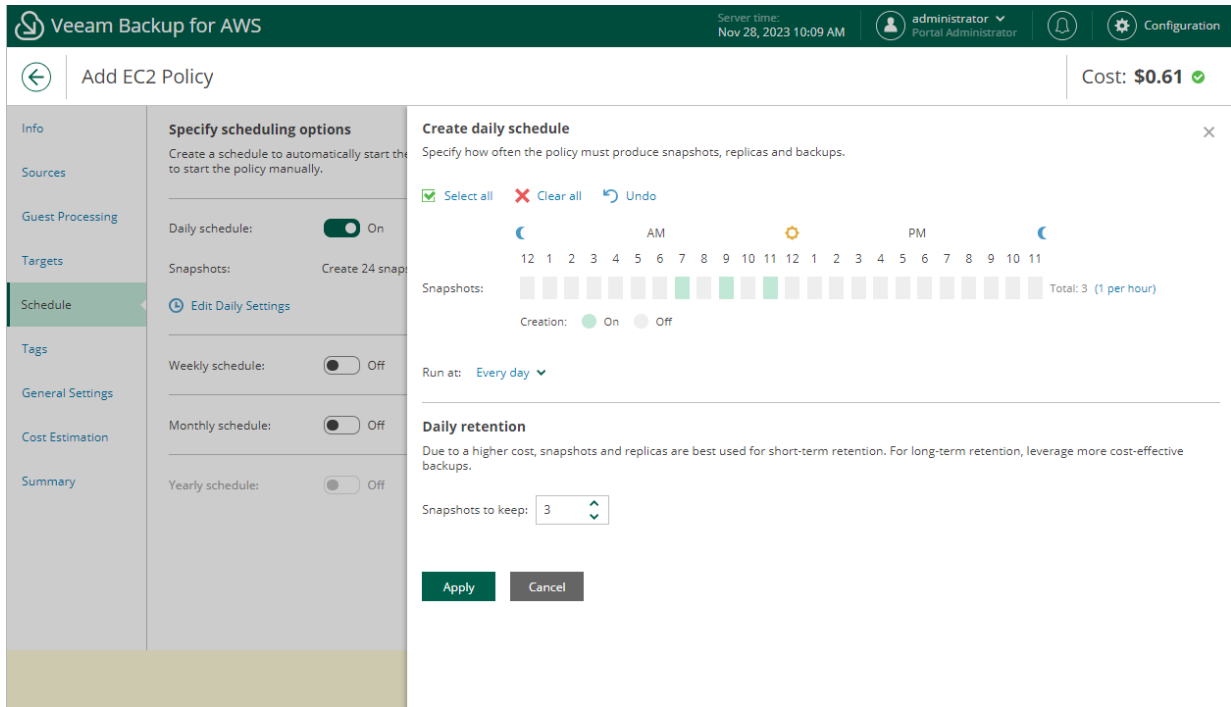
NOTE

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:

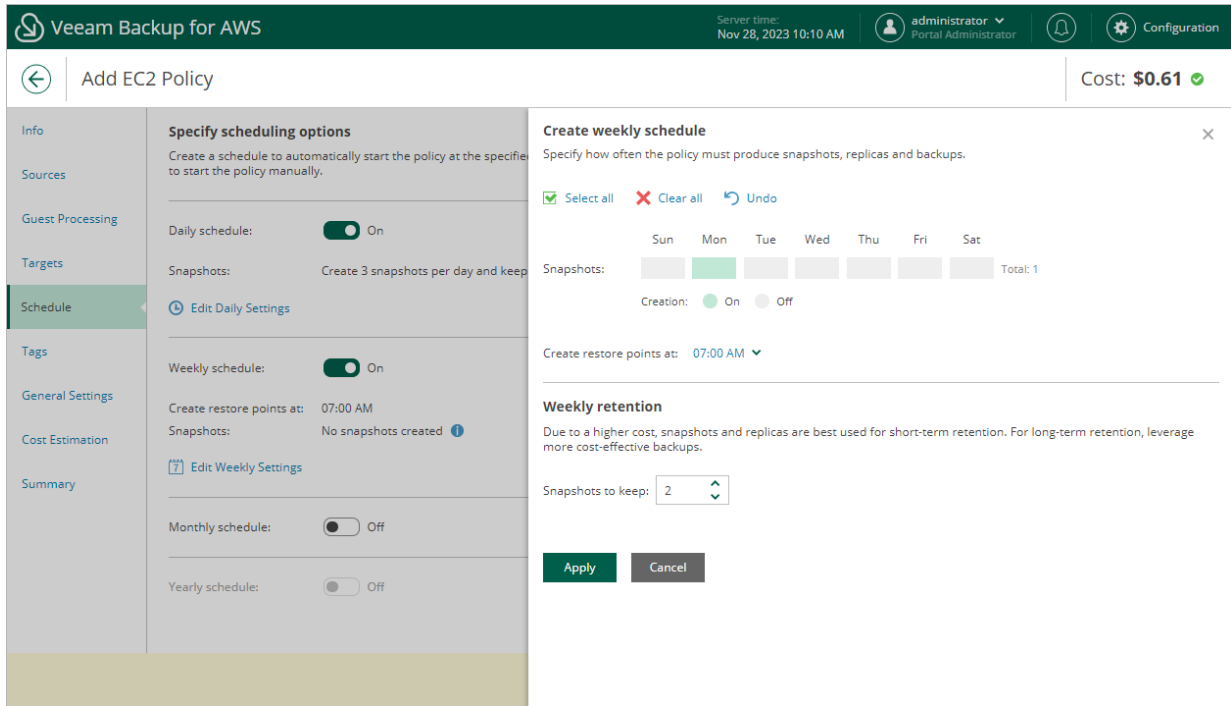
- In the daily scheduling settings, you select hours and days when snapshots will be created (for example, *7:00 AM, 9:00 AM, and 11:00 AM; Working Days*), and specify a number of daily restore points to retain (for example, *3*).

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 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.

You must consider that restoring from an archived backup will take more time to complete and cost more than restoring from a standard backup, as archived data is not available for real-time access and it is required to retrieve the data from the archive backup repository before performing the operation. For more information, see [Retrieving EC2 Data From Archive](#).

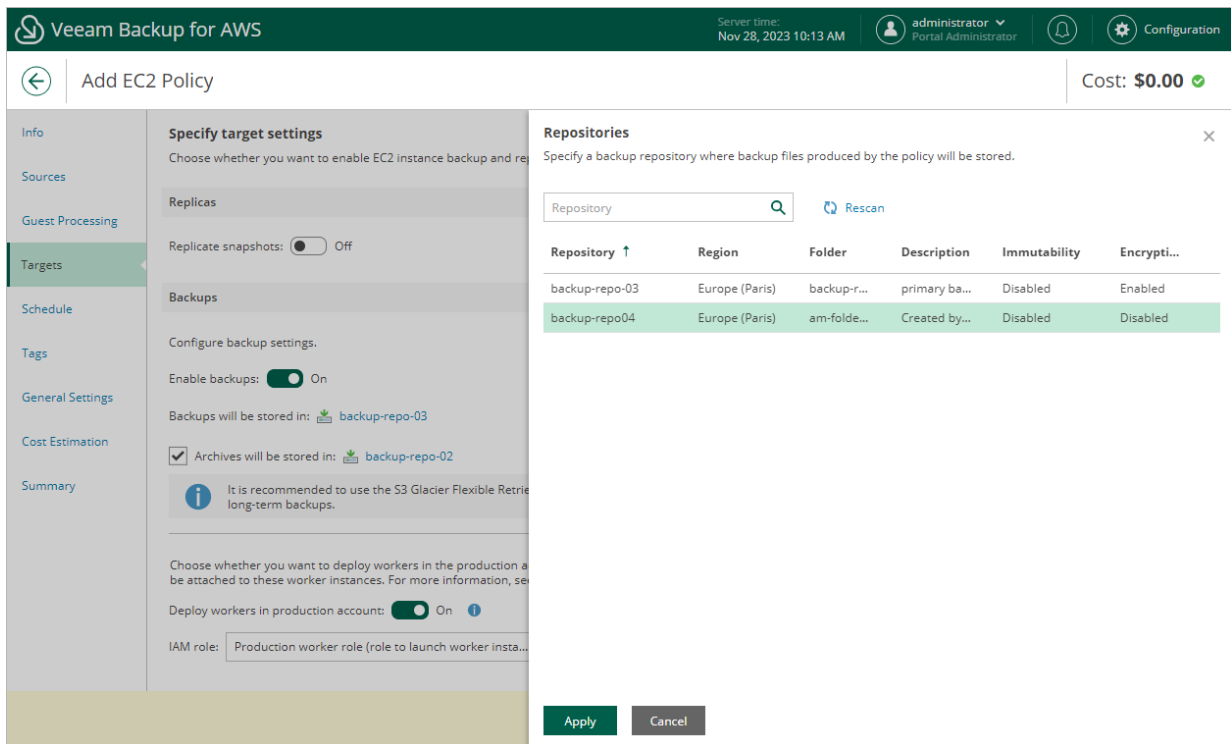
With backup archiving, Veeam Backup for AWS can retain backup files created according to a daily, weekly or monthly schedule for longer periods of time:

- 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](#).

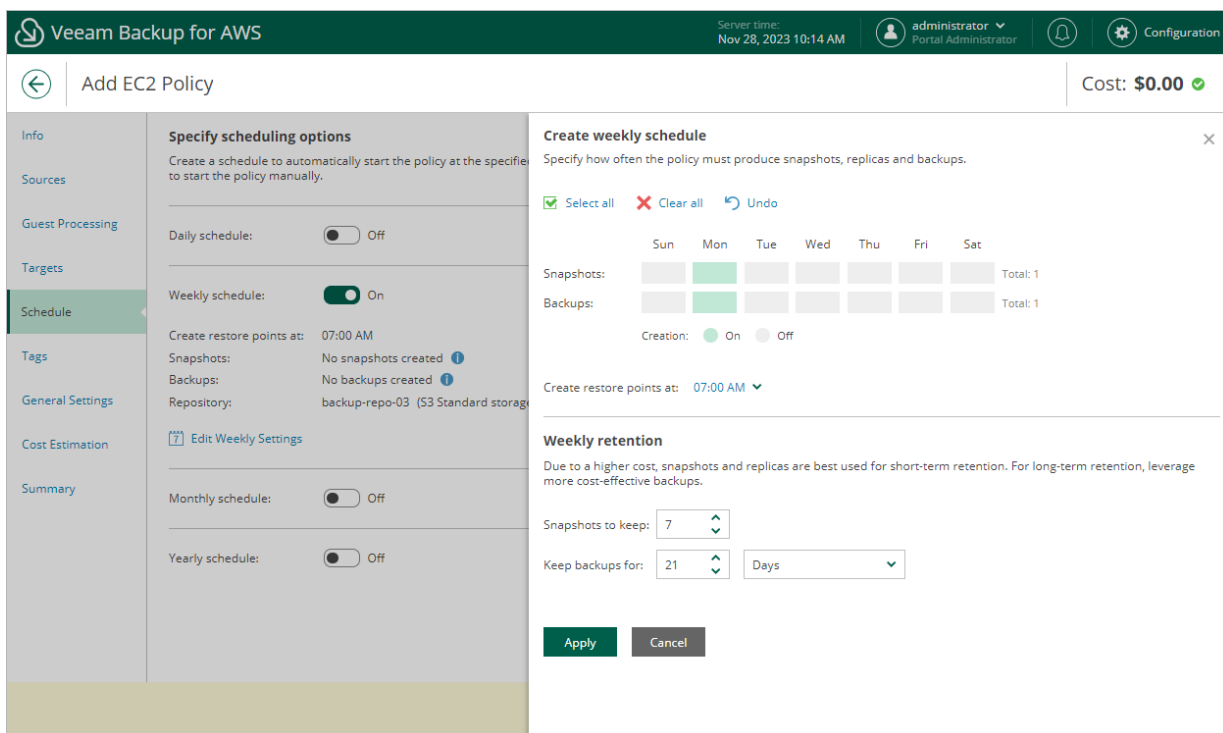
Consider the following example. You want a backup policy to create image-level backups of your critical workloads once a week, to keep the backed-up data in a standard backup repository for 3 weeks, and also to keep backups created once in 2 months in an archive backup repository for a year. In this 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.



- 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).



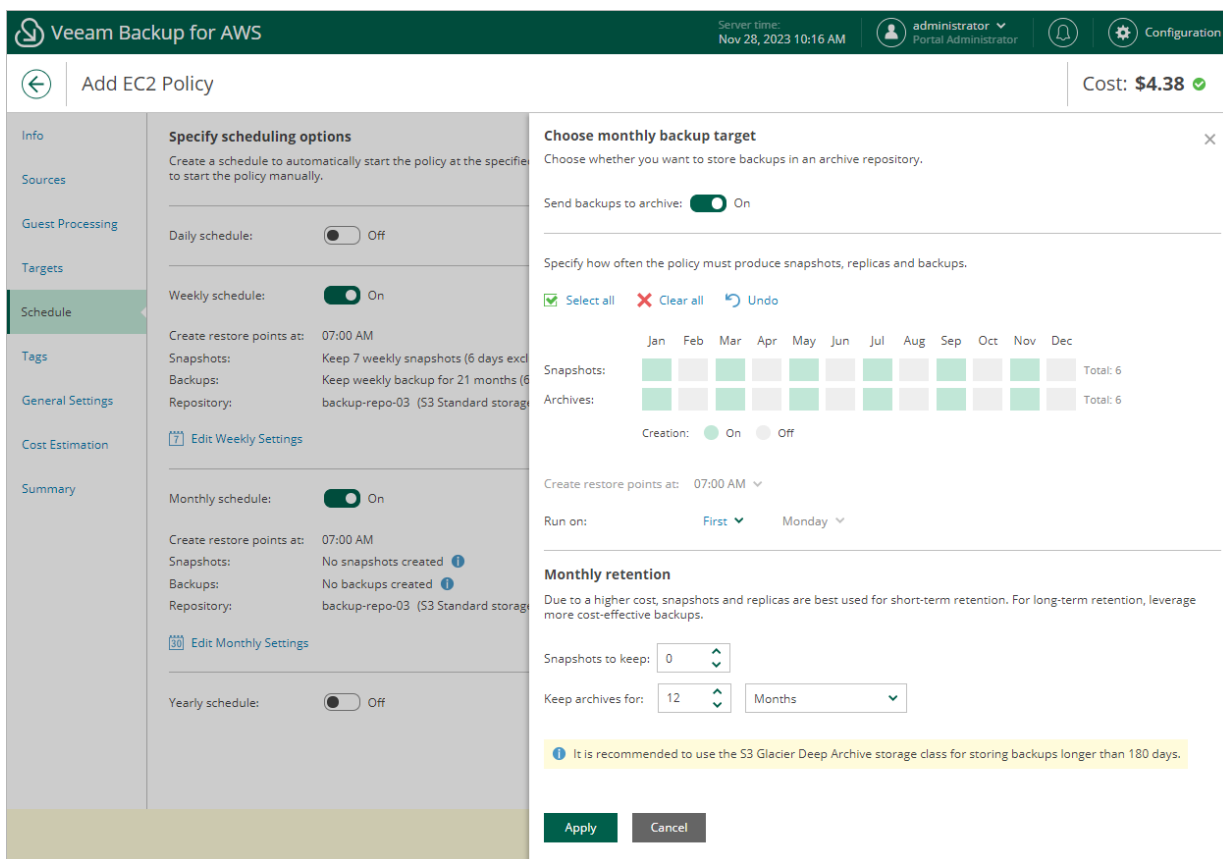
- In the monthly scheduling settings, you enable the archiving mechanism by setting the **Send backups to archive** toggle to *On*, specify when Veeam Backup for AWS will create archive backup files, and choose for how long you want to keep the created backups in the archive backup repository.

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.



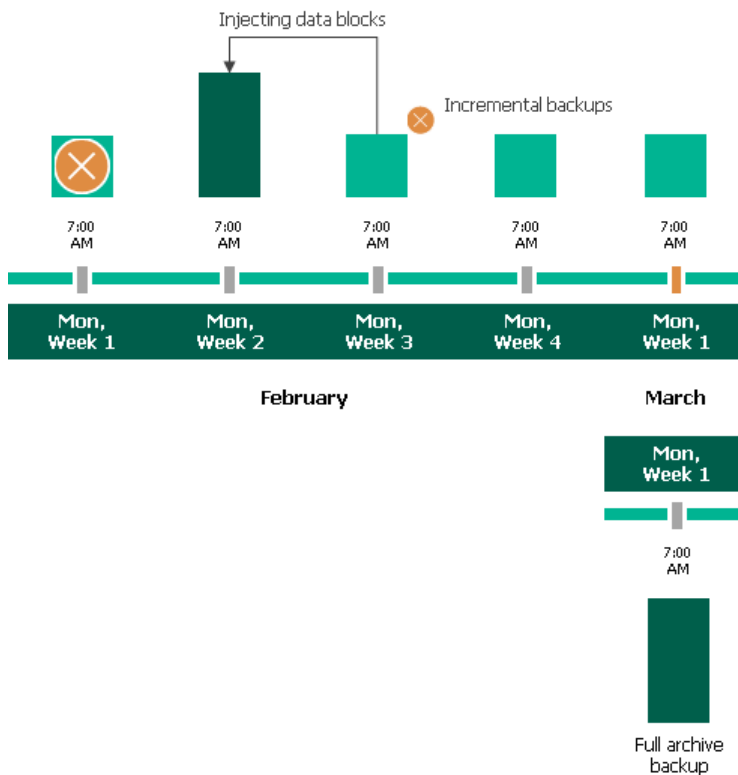
- On the fourth Monday of February, Veeam Backup for AWS will create a new restore point at 7:00 AM. By the moment the backup session completes, the earliest restore point in the standard backup chain will get older than the specified retention limit. That is why Veeam Backup for AWS will rebuild the 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.

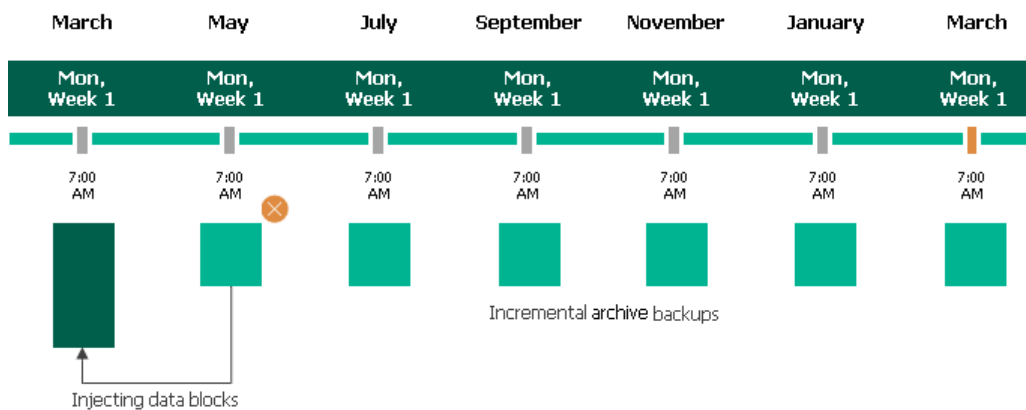
On the first Monday of May, an archive session will create a restore point with only that data that has changed since the previous archive session in March. Veeam Backup for AWS will copy this restore point as an incremental archive backup file to the archive backup repository.



- 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

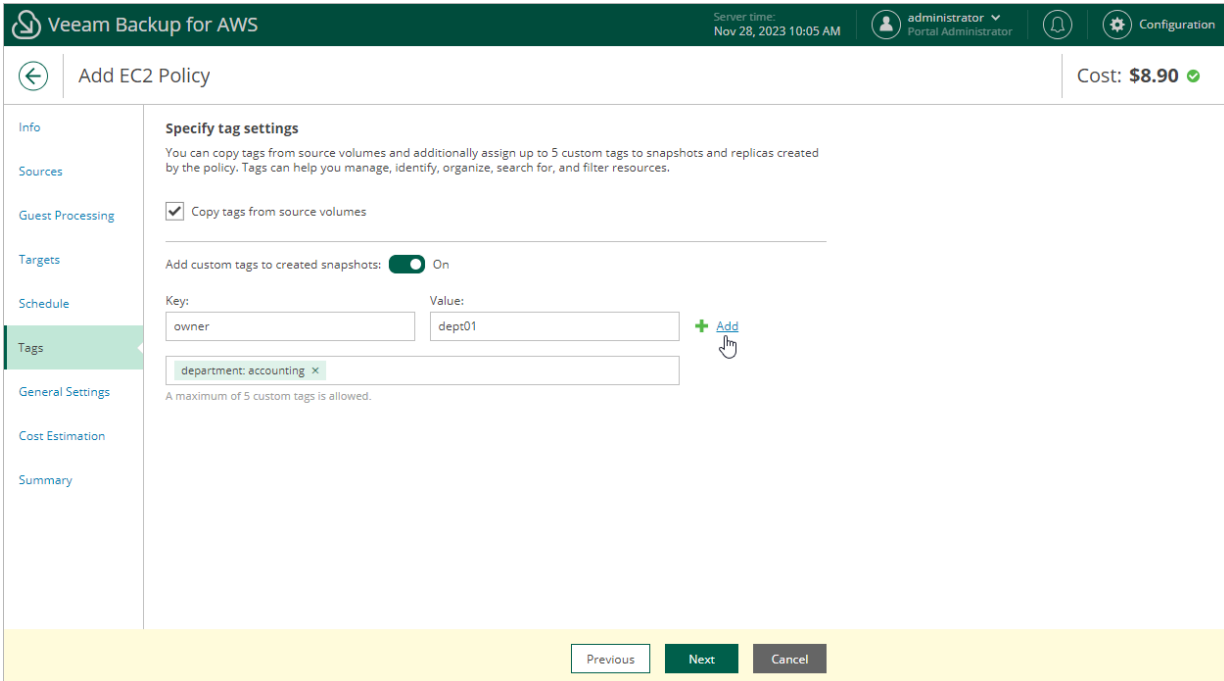
At the **Tags** step of the wizard, you can instruct Veeam Backup for AWS to assign AWS tags to snapshots and snapshots replicas:

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.



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 send each notification to this recipient twice.

The screenshot displays the 'Add EC2 Policy' configuration interface in Veeam Backup for AWS. The 'General Settings' tab is active, showing the following configuration:

- Configure retry and notification settings:** Specify how many times to retry the policy. You can also enable email notifications to receive policy results.
- Schedule:**
 - Automatically retry failed policy: 3 times
 - Automatic retry settings are only applicable on a scheduled run of the policy
- Health check:**
 - A health check includes an availability check for data blocks in backup files and a CRC check for metadata to verify its integrity. Scheduling options are based on the configured policy schedule.
 - Enable health check: On
 - Run on: First Sunday of every month
- Notifications:**
 - Enabled: On
 - Email: donnaortiz@company.com
 - Notify on:
 - Failure
 - Warning
 - Success
 - Suppress notifications until the last retry

At the bottom of the page, there are three buttons: 'Previous', 'Next', and 'Cancel'.

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 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).
If you get a warning message regarding additional costs associated with cross-region data transfer, you can click **View details** to see available cost-effective options.
- 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](#).

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](#).

\$3.97 Snapshots	\$3.23 Replicas	\$0.77 Backups	\$0.05 Archives	\$0.72 Traffic	\$0.17 Transactions
----------------------------	---------------------------	--------------------------	---------------------------	--------------------------	-------------------------------

Estimated monthly cost: \$8.90

Instance Export to...

Instance ↑	Snapshot	Replica	Bac
amroz-vm03	\$1.98	\$1.61	\$
amroz-vm04	\$1.98	\$1.61	\$

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:

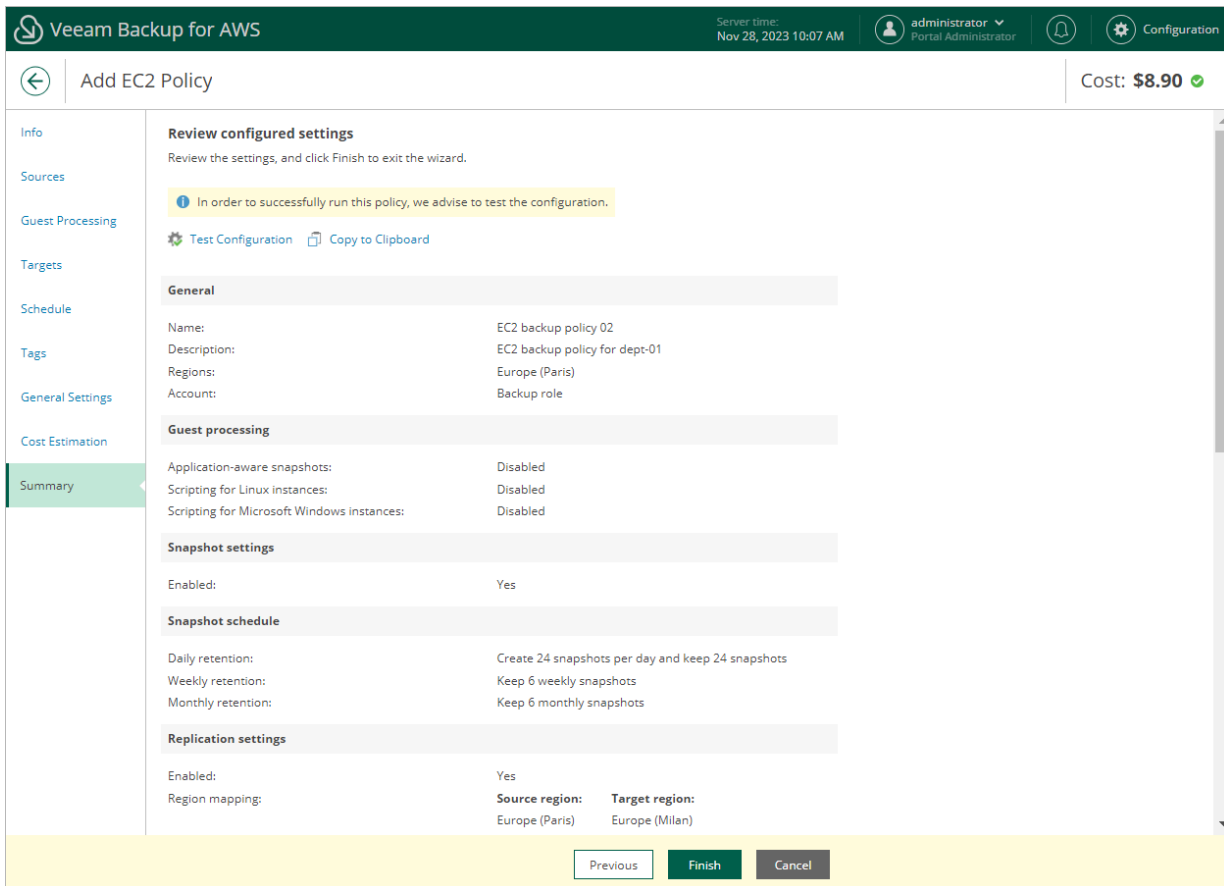
```
"iam:AttachRolePolicy",  
"iam:CreatePolicy",  
"iam:CreatePolicyVersion",  
"iam:CreateRole",  
"iam:GetAccountSummary",  
"iam:GetPolicy",  
"iam:GetPolicyVersion",  
"iam:GetRole",  
"iam:ListAttachedRolePolicies",  
"iam:ListPolicyVersions",  
"iam:SimulatePrincipalPolicy",  
"iam:UpdateAssumeRolePolicy"
```

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.



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 launch 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 details, see [Disabling and Enabling Backup 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 [System Requirements](#) section.
4. After network issues are fixed, you can enable the backup policy. For details, see [Disabling and Enabling Backup 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.

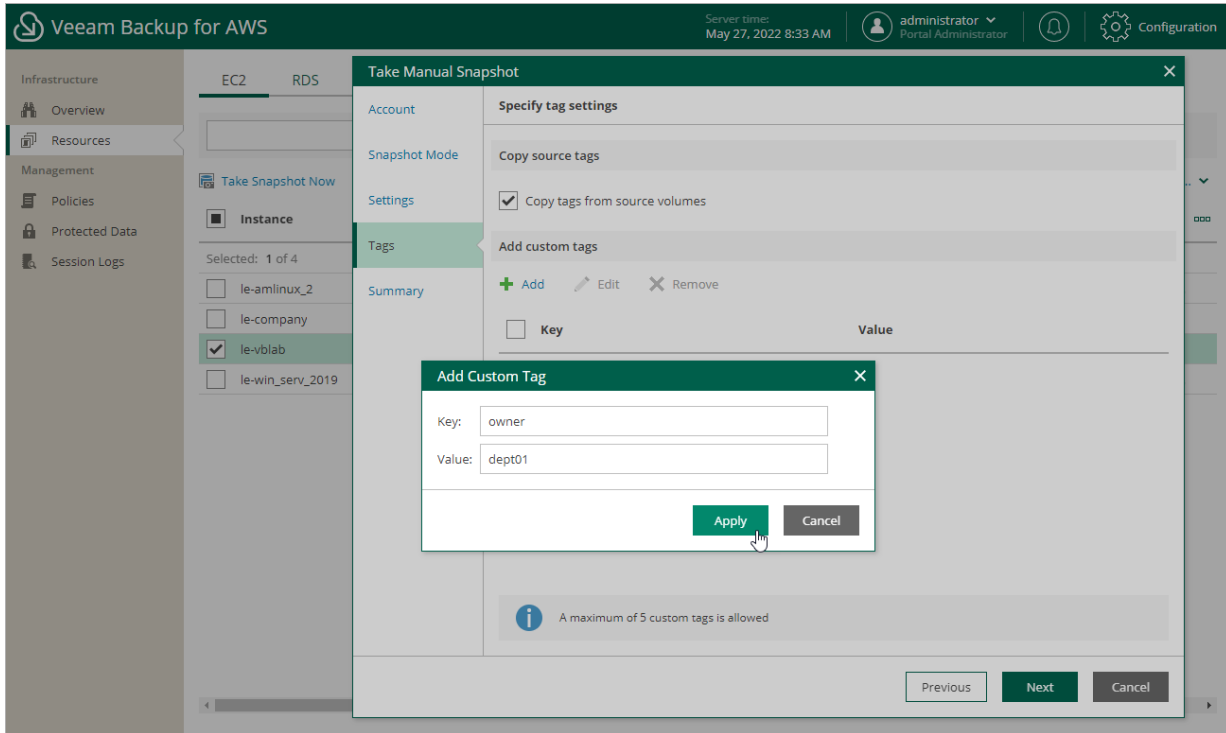
- To assign already existing AWS tags from the EBS volumes of the processed EC2 instance, select the **Copy tags from source volumes** check box.

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 [global notification settings](#) first.

For DB instances and Aurora DB clusters residing in any of the regions added to the backup policies, you can also [take a cloud-native snapshot manually](#) when needed.

Creating RDS Backup Policies

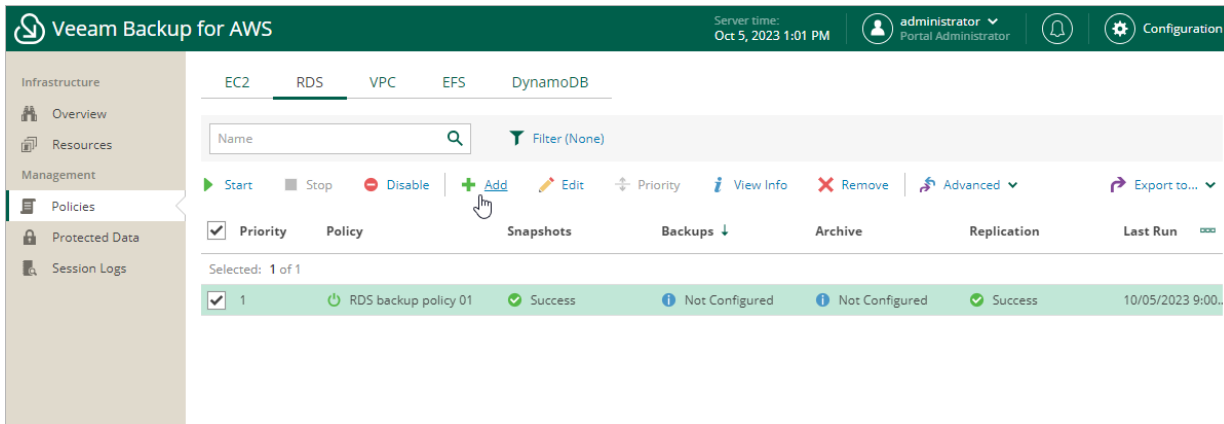
To create an RDS backup policy, do the following:

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 for protecting 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 > 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 backup policy 02' and the 'Description' field contains 'Backup of Dept01 databases'. 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.

Veeam Backup for AWS Server time: Oct 5, 2023 1:03 PM 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 backup policy 02

Description:
Backup of Dept01 databases

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.](#)

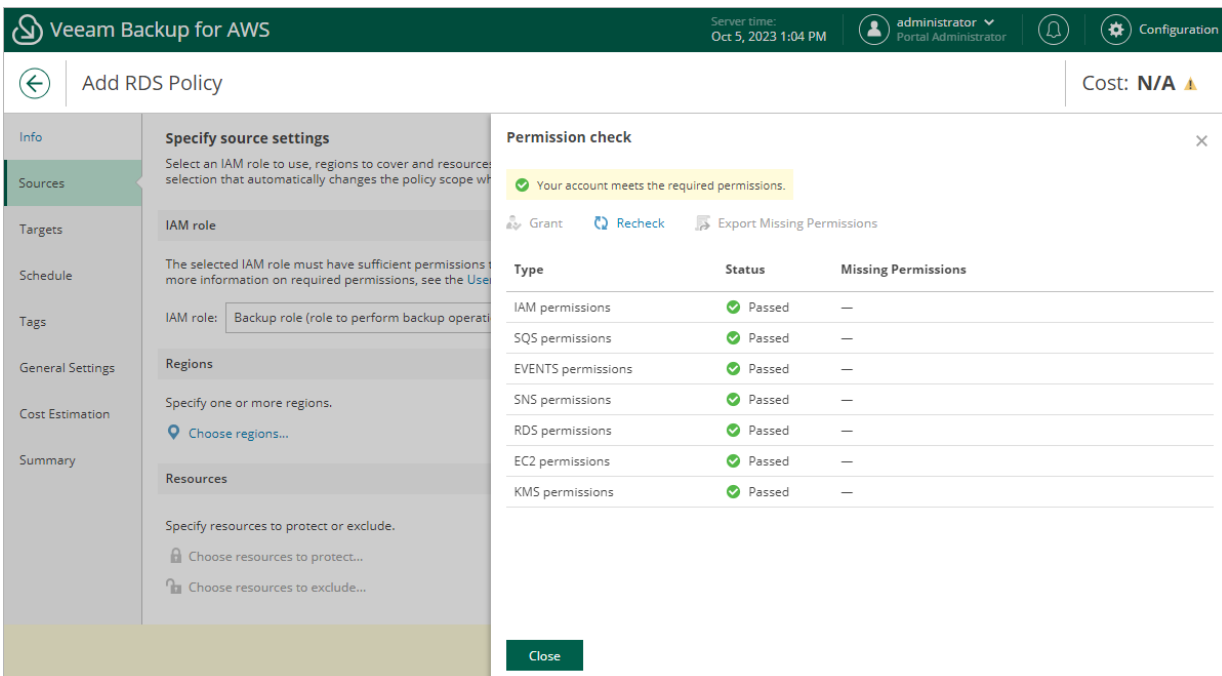
Step 3.1 Specify IAM Role

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 add an IAM role, click **Add** and complete the [Add IAM Role](#) wizard.

IMPORTANT

It is recommended that you check whether the selected IAM role has all the required permissions 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](#).



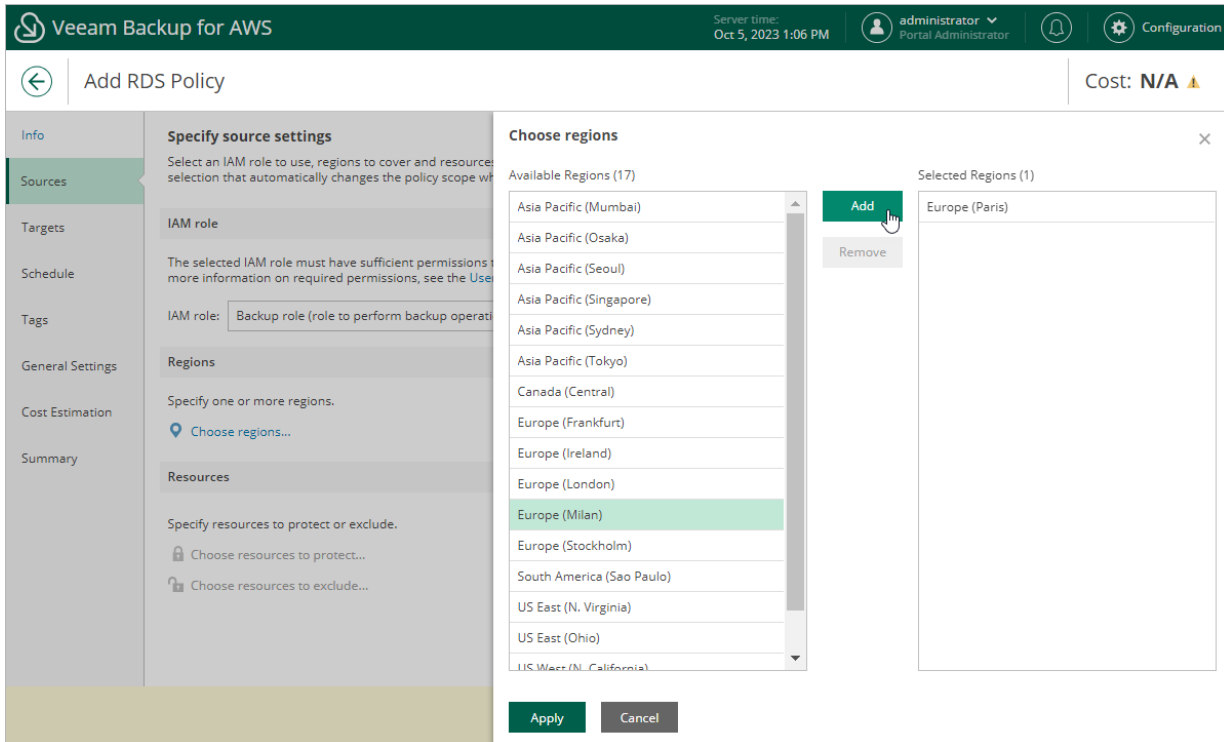
The screenshot shows the 'Add RDS Policy' wizard in Veeam Backup for AWS. The 'Sources' step is active, and the 'IAM role' is set to 'Backup role (role to perform backup operati...'. A 'Permission check' dialog box is open, displaying a table of permissions and their status.

Type	Status	Missing Permissions
IAM permissions	✓ Passed	—
SQS permissions	✓ Passed	—
EVENTS permissions	✓ Passed	—
SNS 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**.



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' modal is open, with the 'Protect only following resources' option selected. The 'Type' is set to 'Database' and the 'Database ID' is 'db01 (db-2c4mz2osktmujfkl4nsylew...)'. A table below shows two resources selected for protection:

Item	ID	Value	Region
db01	db-2c4mz2osktmujfkl4ns...	—	Europe (Paris)
db02	db-2fa6nhcfynux7t7jeaba...	—	Europe (Paris)

The 'Apply' button is highlighted in green.

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.

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

Consider that several limitations are applied to Aurora DB clusters:

- Snapshot replication is not supported for Aurora multi-master clusters.
- 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

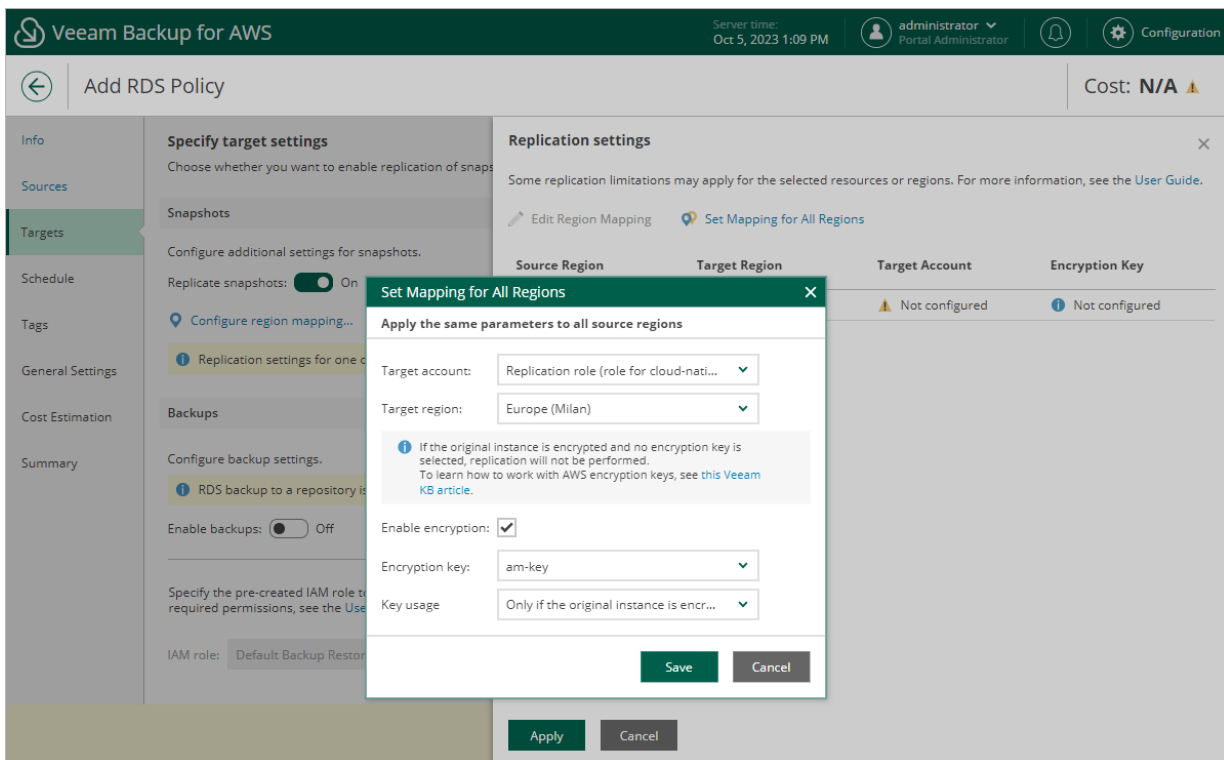
Consider the following:

- 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, 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 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 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 *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 add an IAM role, 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 create RDS image-level backups, Veeam Backup for AWS launches worker instances in a production account – that is, the same AWS account to which the processed resources belong. By default, Veeam Backup for AWS uses the most appropriate network settings of AWS Regions in production accounts to launch worker instances for RDS image-level backup operations. 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 launch worker instances for the RDS backup operation. For Veeam Backup for AWS to be able to launch a worker instance used to create an image-level backup:

- 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](#).

- The VPC to which the DB instance is connected must have at least one security group that allows outbound access on port **443**. This port is used by worker instances to communicate with [AWS services](#).

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 shows the Veeam Backup for AWS console interface. The main window is titled "Add RDS Policy" and displays a "Permission check" dialog box. The dialog box contains a message: "Your account does not meet the required permissions." Below this message are buttons for "Grant", "Recheck", and "Export Missing Permissions". A table below the message shows the results of the permission check:

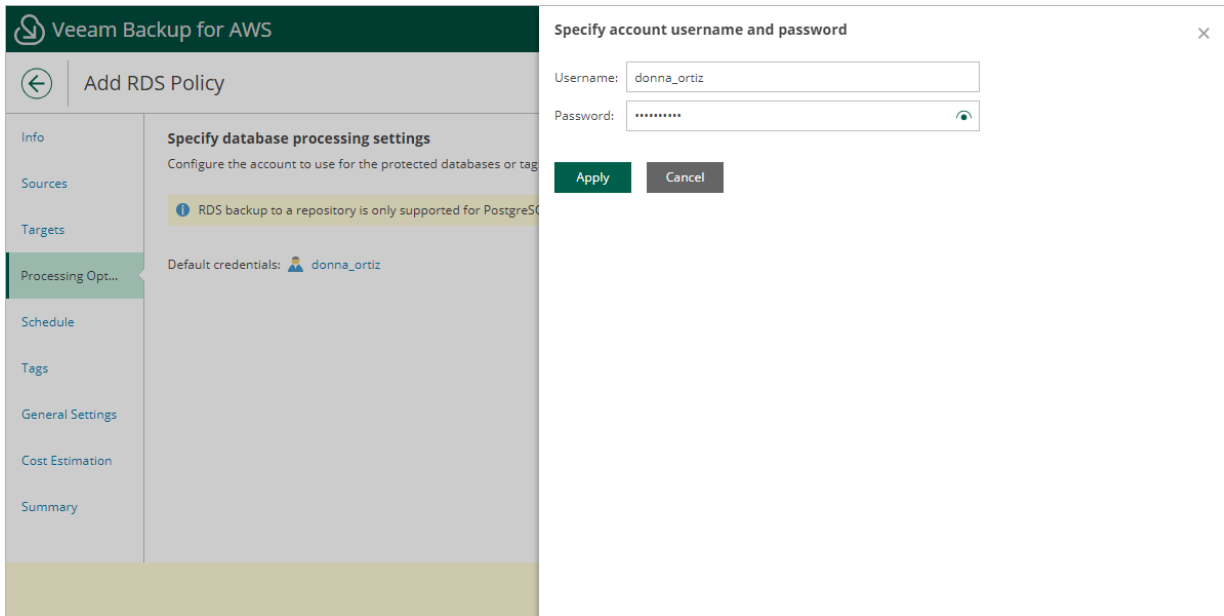
Type	Status	Missing Permissions
Checking backup policy role ...	Passed	—
Checking worker role permis...	Failed	Trust Policy: The following service must be added to t...

A "Grant Permissions" dialog box is open in the foreground, titled "Grant Permissions". It contains a 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." Below the message are fields for "Access key:" (AKIAY4ZWOU4WMVRAGEVN) and "Secret key:" (masked with dots). There are "Apply" and "Cancel" buttons at the bottom of the dialog box.

Step 5. Specify Processing Settings

[This step applies only if you have enabled backups at the **Targets** step of the wizard]

At the **Processing Options** step of the wizard, specify credentials of a user that Veeam Backup for AWS will use to access the databases and perform the backup operation. The specified user must exist on all DB instances processed by the policy.



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.

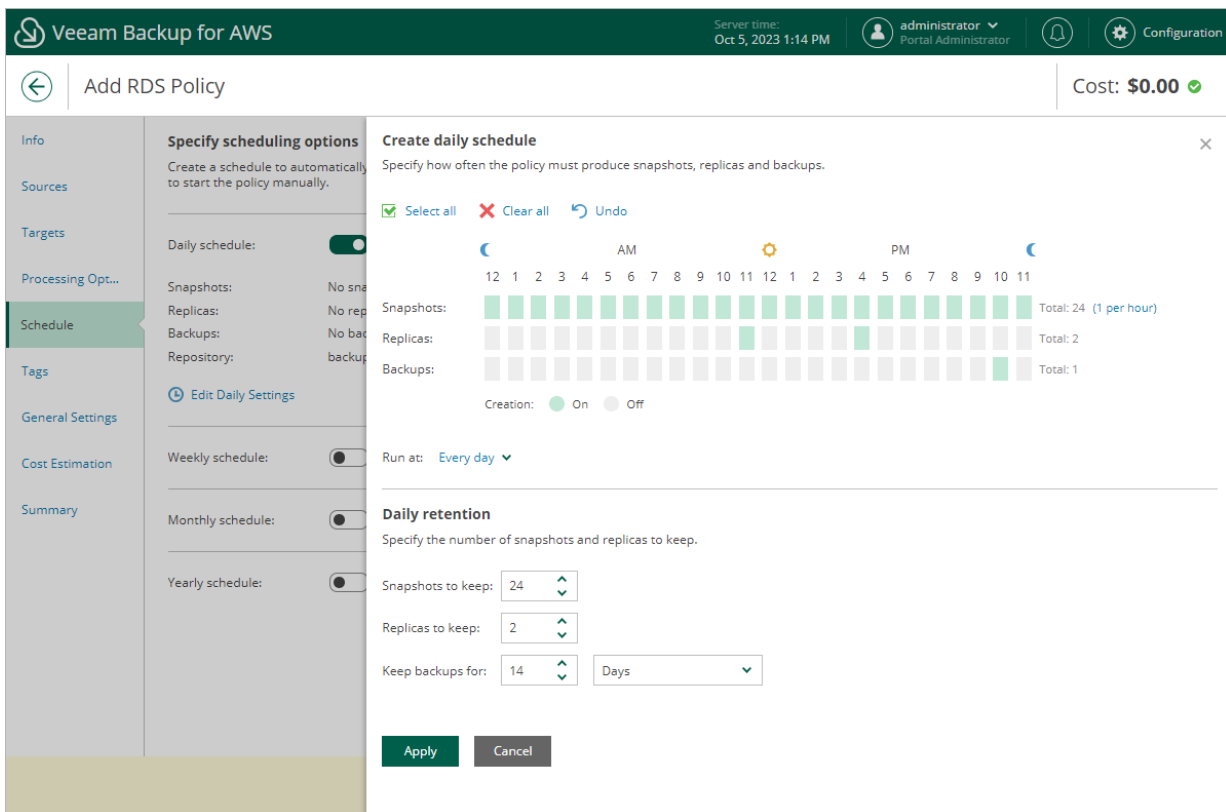
If you want to protect RDS resources data more frequently, you can instruct the backup policy to create multiple cloud-native snapshots per hour. To do that, click the link to the right of the **Snapshots** hour selection area, and specify the number of cloud-native snapshots that the backup policy must 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:

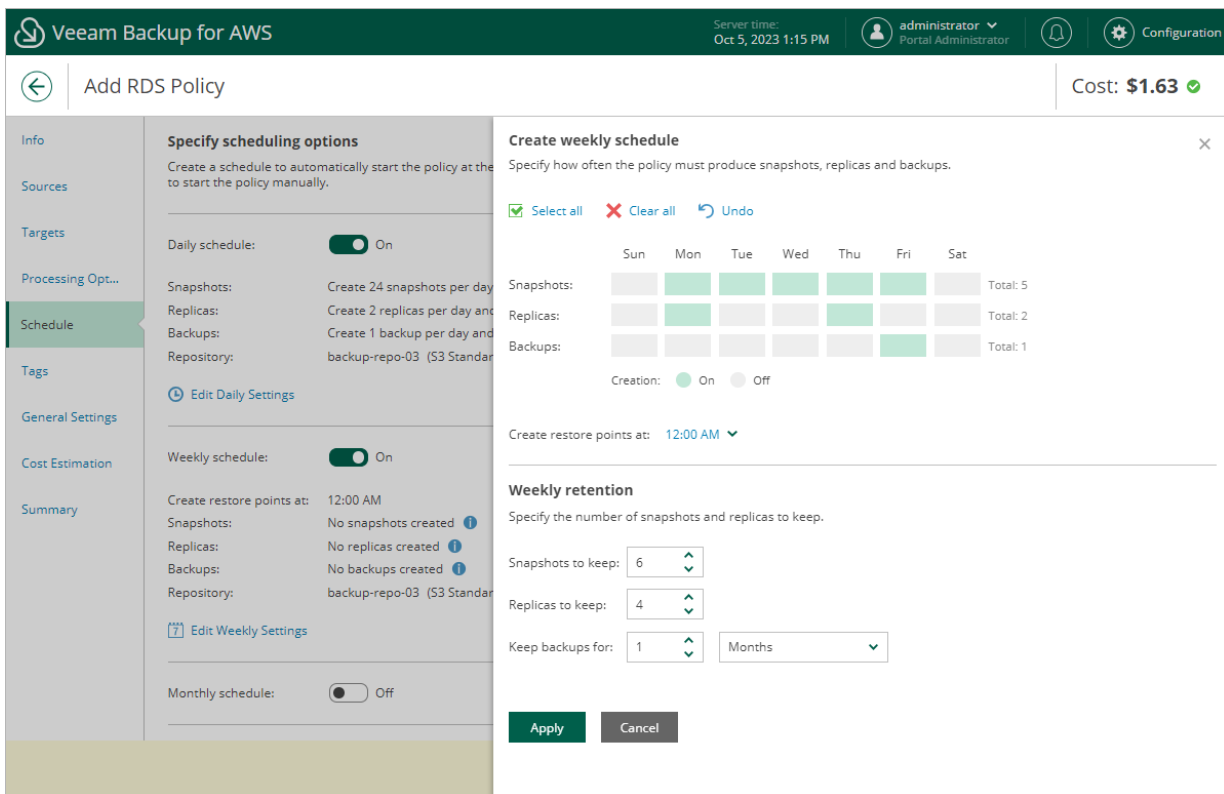
1. Set the **Weekly schedule** toggle to *On* and click **Edit Weekly Settings**.
2. In the **Create weekly schedule** window, select weekdays when the backup policy must create cloud-native snapshots, snapshot replicas or image-level backups.

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:

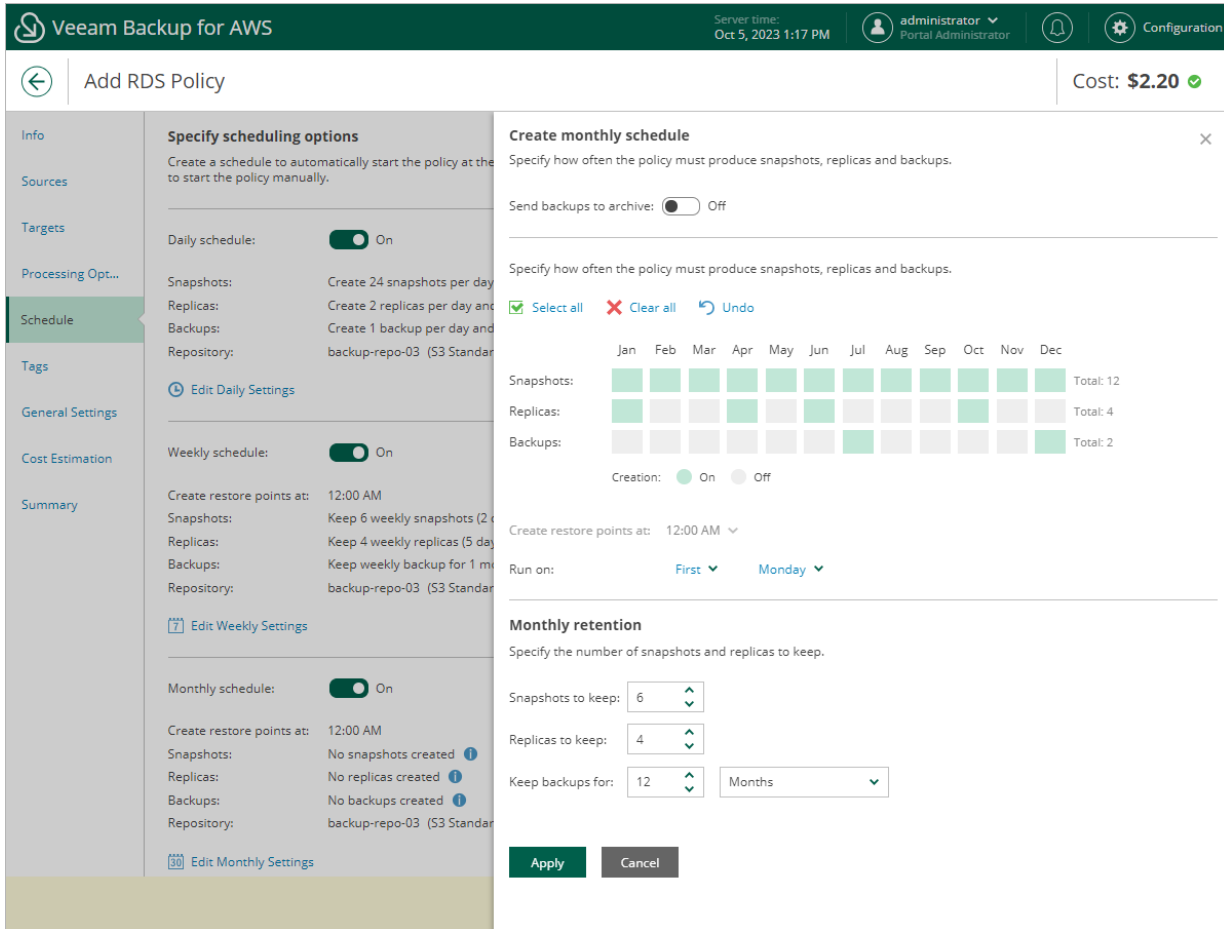
- If you have selected a specific time for the backup policy to run at the **Weekly schedule** section of the **Schedule** step of the wizard, you will not be able to change the time for the monthly schedule unless you select the **On Day** option from the **Run on** drop-down list.
- If you select the **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 monthly 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 Veeam Backup for AWS console interface. At the top, the header includes the Veeam logo, the text 'Veeam Backup for AWS', the server time 'Oct 5, 2023 1:19 PM', and user information 'administrator Portal Administrator'. The main content area is titled 'Add RDS Policy' and shows a 'Cost: \$8.50' indicator. A left-hand navigation pane lists various settings categories: Info, Sources, Targets, Processing Opt..., Schedule (highlighted), Tags, General Settings, Cost Estimation, and Summary. The 'Schedule' section is expanded, showing three schedule types: Daily, Weekly, and Yearly, each with a toggle switch set to 'On'. The 'Yearly schedule' section is currently selected, displaying settings for 'Create backup on: First Monday of July at 12:00', 'Keep archives for: 2 years', and 'Repository: backup-repo-02 (S3 Glacier)'. An 'Edit Yearly Settings' link is visible below. A modal dialog box titled 'Create yearly schedule' is open on the right, providing a detailed view of the yearly schedule configuration. It includes a description, a 'Send backups to archive' toggle (set to 'On'), and a section for specifying the schedule: 'Create restore points on: First Monday of July at 12:00 AM' and 'Keep archives for: 2 years'. 'Apply' and 'Cancel' buttons are at the bottom of the dialog.

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:

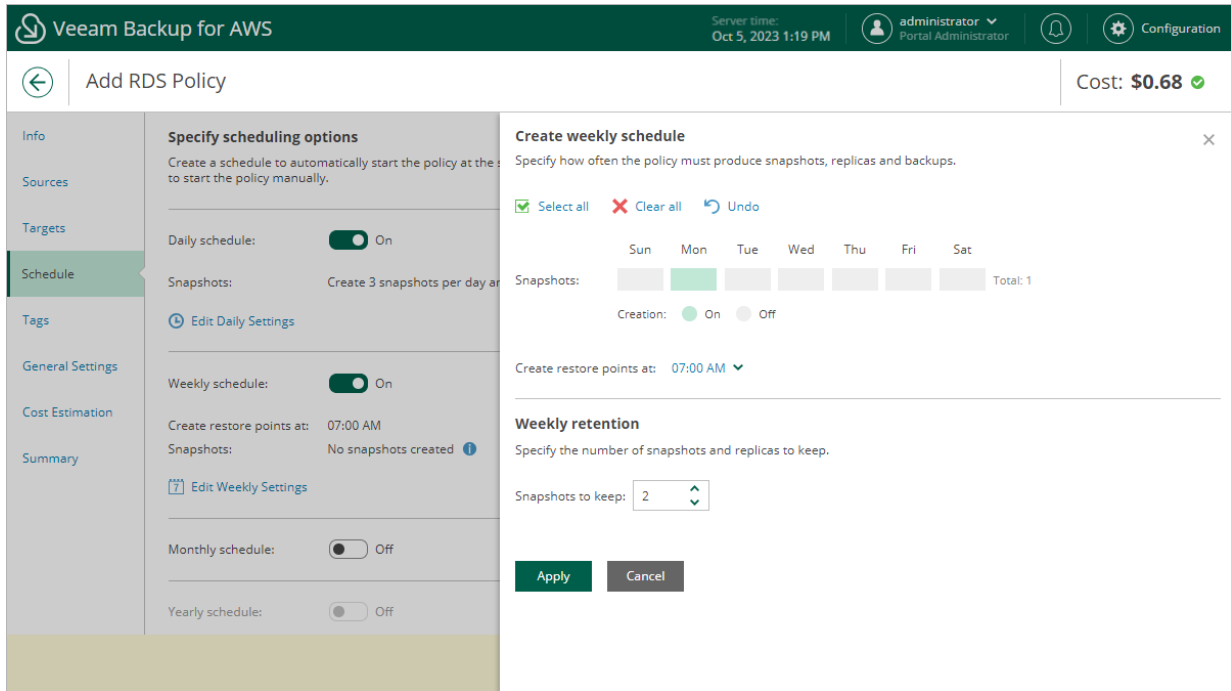
- In the daily scheduling settings, you select hours and days when snapshots will be created (for example, *7:00 AM, 9:00 AM, and 11:00 AM; Working Days*), and specify a number of daily restore points to retain (for example, *3*).

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 Veeam Backup for AWS configuration interface. At the top, the header shows 'Veeam Backup for AWS' and 'Server time: Oct 5, 2023 1:19 PM'. The user is logged in as 'administrator Portal Administrator'. The main area is titled 'Add RDS Policy' with a 'Cost: \$0.69' indicator. A sidebar on the left contains navigation options: Info, Sources, Targets, Schedule (selected), Tags, General Settings, Cost Estimation, and Summary. The 'Specify scheduling options' section is active, showing 'Daily schedule' enabled. Below it, 'Snapshots' is set to 'Create'. The 'Create daily schedule' section allows specifying how often the policy must produce snapshots, replicas and backups. It includes 'Select all', 'Clear all', and 'Undo' options. A 24-hour clock is shown with snapshots scheduled at 7 AM, 9 AM, and 11 AM. The 'Snapshots' section shows 'Creation' set to 'On' and 'Off'. The 'Run at' dropdown is set to 'Every day'. The 'Daily retention' section specifies the number of snapshots and replicas to keep, with 'Snapshots to keep' set to 3. The 'Apply' button is highlighted with a mouse cursor.

- 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.

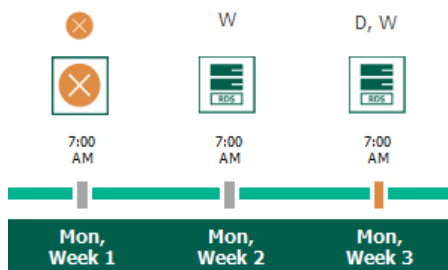
By the moment the backup session completes, the number of restore points with the (D) flag will exceed the retention limit specified in the daily schedule settings. However, Veeam Backup for AWS will not remove the earliest restore point (7:00 AM, Monday) with the (D) flag from the snapshot chain as 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.

You must consider that restoring from an archived backup will take more time to complete and cost more than restoring from a standard backup, as archived data is not available for real-time access and it is required to retrieve the data from the archive backup repository before performing the operation. For more information, see [Performing Database Restore](#).

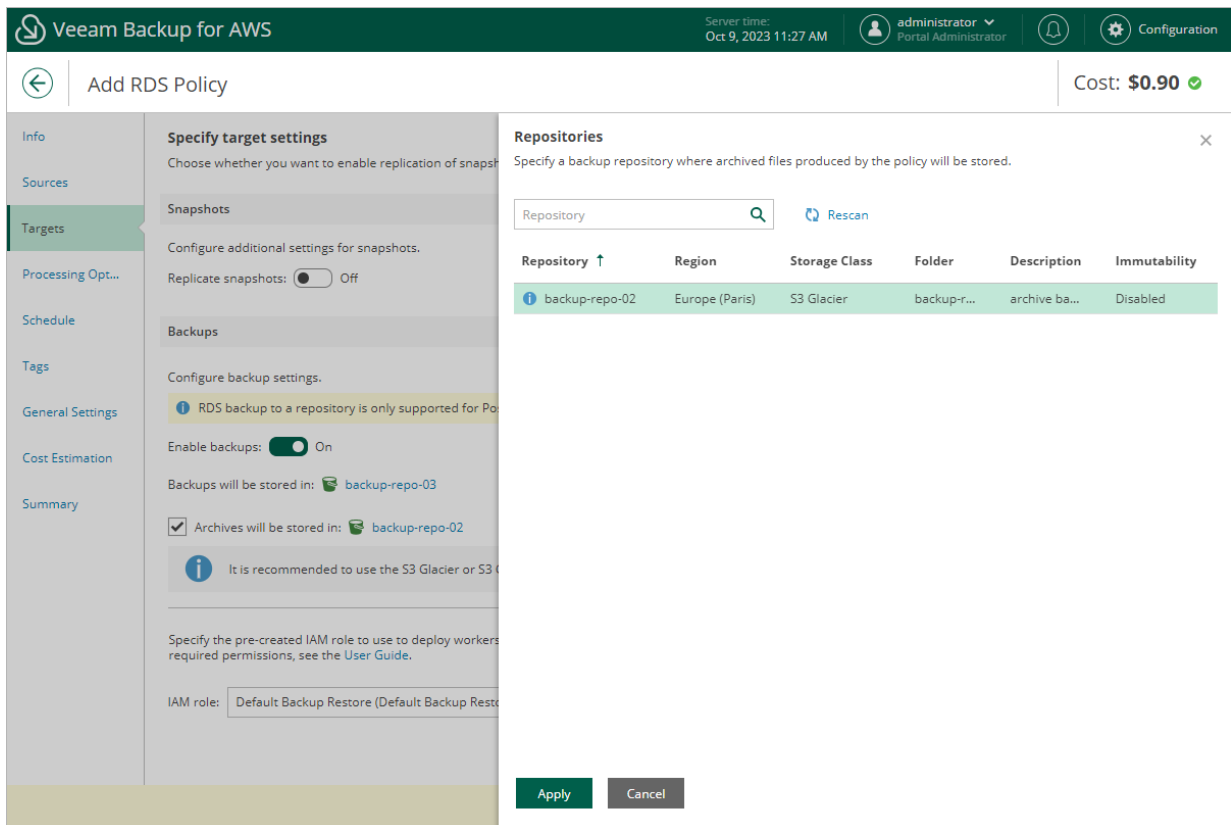
With backup archiving, Veeam Backup for AWS can retain backup files created according to a daily, weekly or monthly schedule for longer periods of time:

- 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](#).

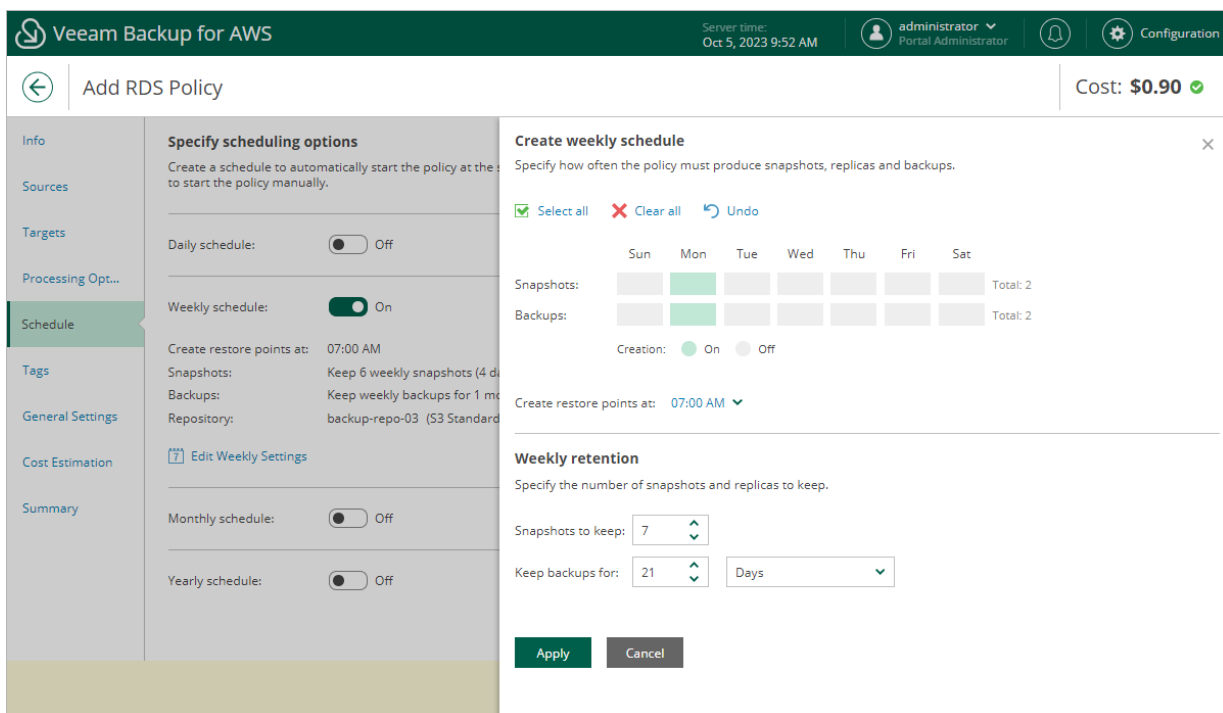
Consider the following example. You want a backup policy to create image-level backups of your critical workloads once a week, to keep the backed-up data in a standard backup repository for 3 weeks, and also to keep backups created once in 2 months in an archive backup repository for a year. In this 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.



- 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).



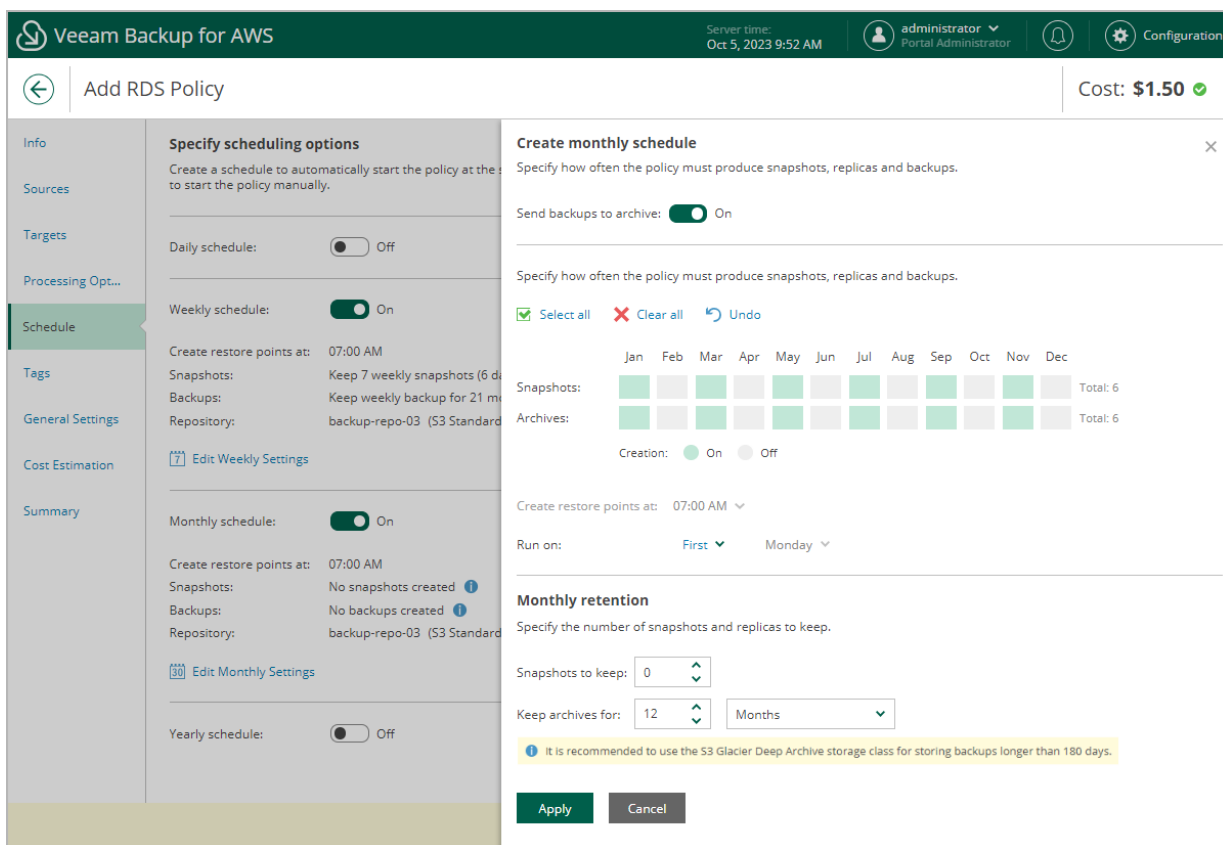
3. In the monthly scheduling settings, you enable the archiving mechanism by setting the **Send backups to archive** toggle to *On*, specify when Veeam Backup for AWS will create archive backup files, and choose for how long you want to keep the created backups in the archive backup repository.

For example, *January, March, May, July, September, November, 12 months* and *First Monday*.

IMPORTANT

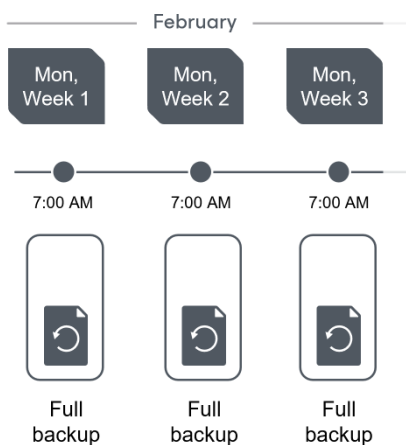
Consider the following:

- 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.



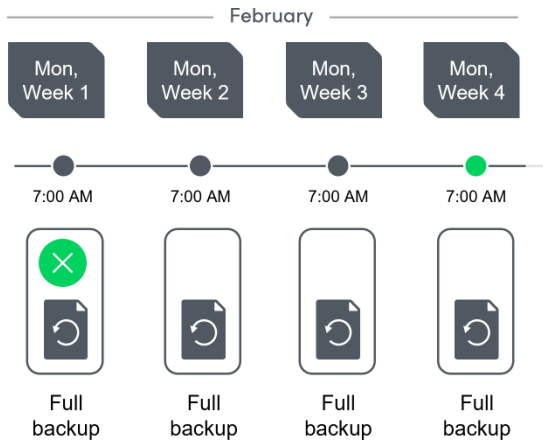
According to the specified scheduling settings, Veeam Backup for AWS will create image-level backups in the following way:

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](#).



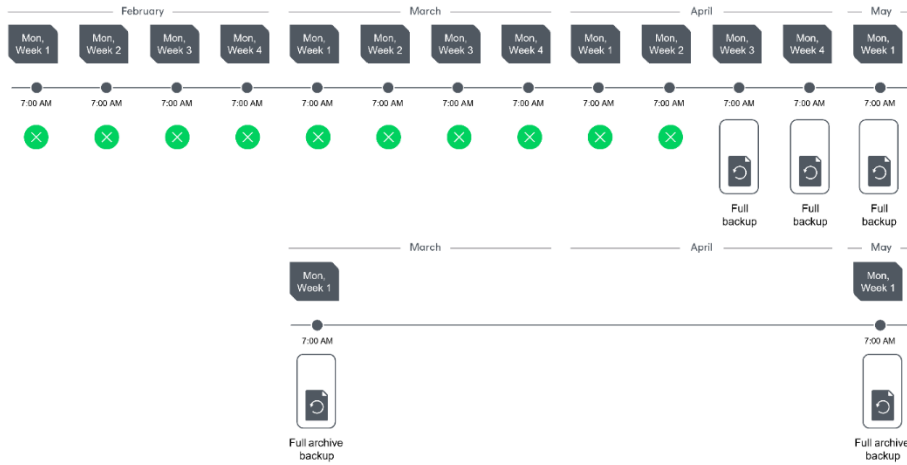
- On the first Monday of March, a backup session will start at 7:00 AM to create another restore point in the standard backup chain. At the same time, the earliest restore point in the standard backup chain will get older than the specified retention limit again. That is why Veeam Backup for AWS will 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 console. The 'Tags' step is active, showing 'Specify tag settings'. The 'Copy tags from source RDS instances' checkbox is checked. The 'Add custom tags to created snapshots' toggle is set to 'On'. Two custom tags are being added: 'user' with value 'donna_ortiz' and 'owner: dept01'. The 'Add' button is highlighted. The cost is estimated at \$8.50. Navigation buttons 'Previous', 'Next', and 'Cancel' are at the bottom.

Server time: Oct 5, 2023 1:20 PM | administrator Portal Administrator | Configuration

← Add RDS Policy | Cost: \$8.50 ✓

Specify tag settings
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.

Copy tags from source RDS instances

Add custom tags to created snapshots: On

Key: Value: [+ Add](#)

[×](#)

A maximum of 5 custom tags is allowed.

Previous Next Cancel

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 send each notification to this recipient twice.

The screenshot shows the 'Add RDS Policy' configuration page in Veeam Backup for AWS. The page is titled 'Add RDS Policy' and has a cost of '\$8.50'. The 'General Settings' section is active, showing the 'Configure retry and notification settings' configuration. The 'Schedule' section has 'Automatically retry failed policy' checked and 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' checked and set to 'On'. The 'Run on' section is set to 'First' of 'Sunday' of every month. The 'Notifications' section has 'Enabled' checked and set to 'On'. The 'Email' field contains 'donna_ortiz@companymail.com'. The 'Notify on' section has 'Failure', 'Warning', and 'Success' checked. The 'Suppress notifications until the last retry' checkbox is also checked. At the bottom, there are '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 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 instance class, 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 class, the number of restore points to be kept in the snapshot 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).
If you get a warning message regarding additional costs associated with cross-region data transfer, you can click **View details** to see available cost-effective options.
 - 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](#).

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](#).

\$8.89 Snapshots	\$7.23 Replicas	\$0.88 Traffic	\$0.00 Transactions
----------------------------	---------------------------	--------------------------	-------------------------------

Estimated monthly cost: \$16.99

Instance Export to...

Instance ↑	Snapshot	Replica	Traffic
db01	\$4.44	\$3.61	\$0.44
db02	\$4.44	\$3.61	\$0.44

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. 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' with a cost of \$16.99. The 'Summary' step is selected in the left sidebar. The 'Test policy configuration' window is open, displaying a table of test results:

Type	Status	Action	Result
Checking policy configuration...	Success	—	—
Checking policy role permissions	Success	—	—
Checking replication role permissions	Success	—	—
Checking backup repository role permissi...	Success	—	—

The 'Test policy configuration' window also includes a 'Recheck' button and a 'Close' button at the bottom.

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 > 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.

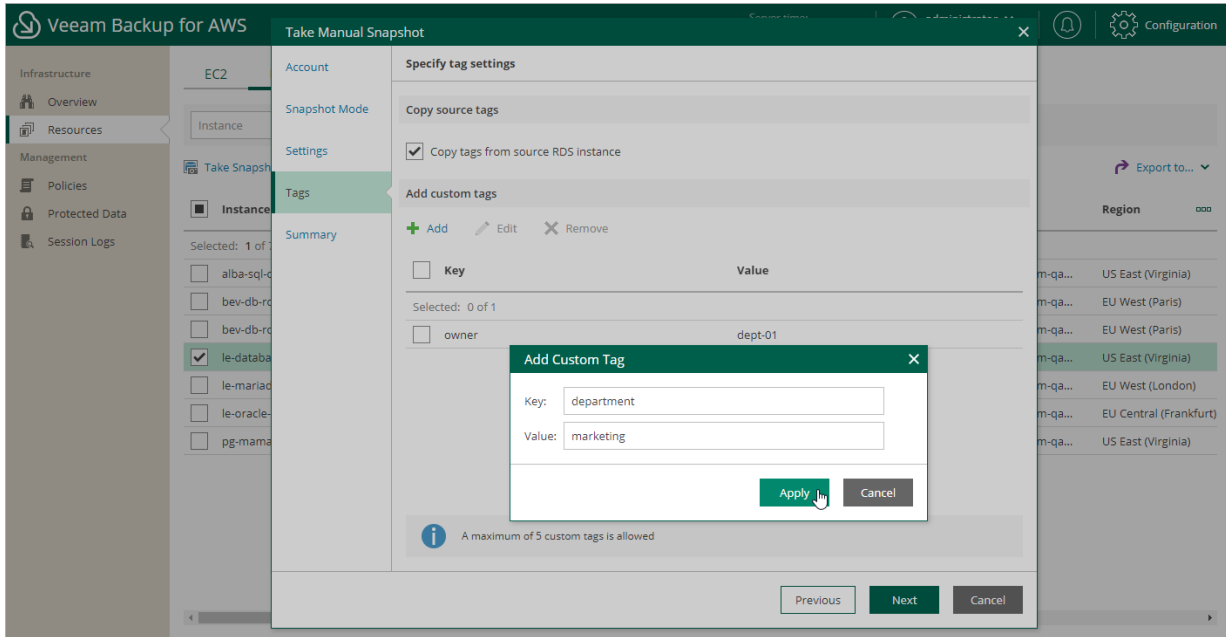
- To assign already existing AWS tags from the source DB instance and Aurora DB cluster, select the **Copy tags from source RDS instance** check box.

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 global notification settings before creating a DynamoDB backup policy. For more information, see [Configuring Global Notification Settings](#).

For DynamoDB tables residing in any of the regions added to the backup policies, 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

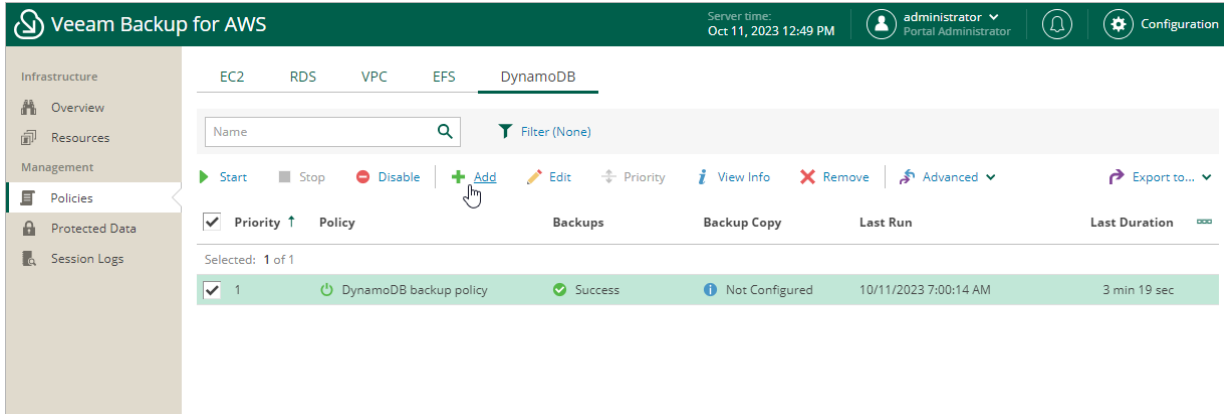
To create a DynamoDB backup policy, do the following:

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 for protecting 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 > 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 Veeam Backup for AWS. The interface includes a top navigation bar with the Veeam logo, server time (Oct 11, 2023 12:51 PM), user information (administrator, Portal Administrator), and a Configuration icon. The main area is titled 'Add DynamoDB Policy' and shows a 'Cost: N/A' warning. A left sidebar contains navigation options: Info (selected), Sources, Targets, Schedule, Tags, General Settings, Cost Estimation, and Summary. The main content area is titled 'Specify policy name and description' and prompts the user to 'Enter a name and description for the policy.' It features two input fields: 'Name:' with the value 'DynamoDB backup policy 02' and 'Description:' with the value 'Created by administrator at 10/11/2023 12:50 PM'. 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 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 cloud-native snapshots 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 add an IAM role, click **Add** and complete the [Add IAM Role](#) wizard.

IMPORTANT

It is recommended that you check whether the selected IAM role has all the required permissions 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 Veeam Backup for AWS. The 'Sources' step is active, showing the 'Specify source settings' section. The 'IAM role' dropdown is open, displaying a list of roles, with 'Backup role (role to perform backup operations and to launch workers in production a...)' selected. The 'Regions' section shows 'Default Backup Restore (Default Backup Restore)' selected. The 'Resources' section is empty. The 'Previous', 'Next', and 'Cancel' buttons are visible at the bottom.

Veeam Backup for AWS

Server time: Oct 11, 2023 12:59 PM

administrator Portal Administrator

Configuration

← Add DynamoDB Policy Cost: N/A

Info

Sources

Targets

Schedule

Tags

General Settings

Cost Estimation

Summary

Specify source settings

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.

IAM role

The selected IAM role must have sufficient permissions to create backups of tables protected by the policy. For more information on required permissions, see the [User Guide](#).

IAM role: Backup role (role to perform backup operations an... + Add Check Permissions

Backup role (role to perform backup operations and to launch workers in production a...

Default Backup Restore (Default Backup Restore)

Specify one or more regions.

Choose regions...

Resources

Specify resources to protect or exclude.

Choose resources to protect...

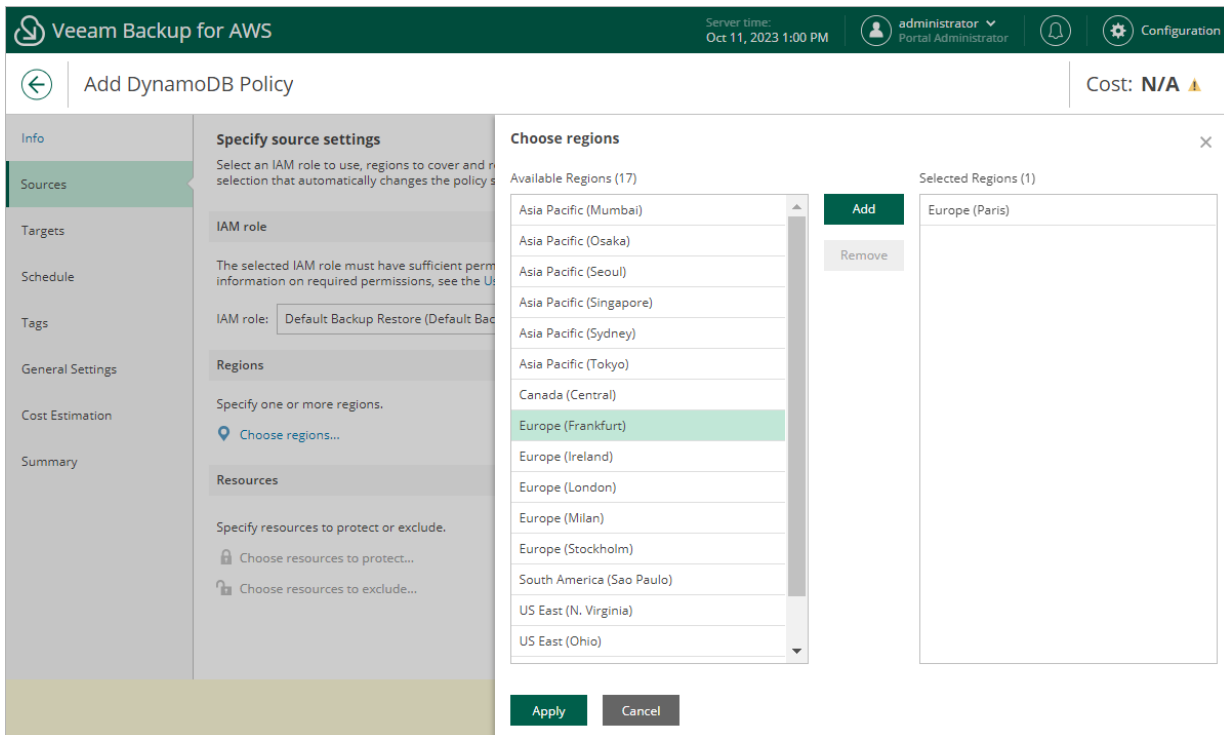
Choose resources to exclude...

Previous Next Cancel

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**.



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 options for IAM role, Regions (1 region selected), 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 "Table" and the "Name" dropdown is set to "DataTable". A "Protect" button is visible. Below the dropdowns is a link: "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	ab96bae2-2d5d-48be-a21...	—	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 that Veeam Backup for AWS will use to store backups of protected DynamoDB tables. 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.

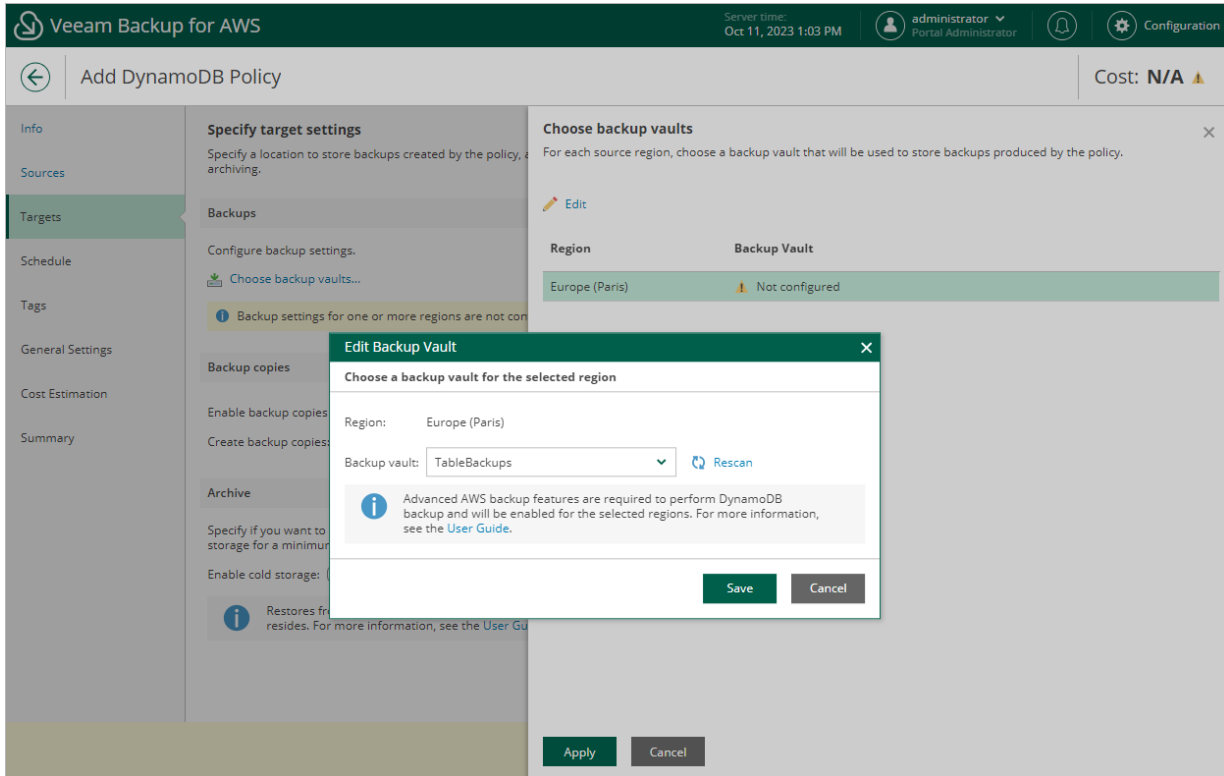
IMPORTANT

Consider the following:

- 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

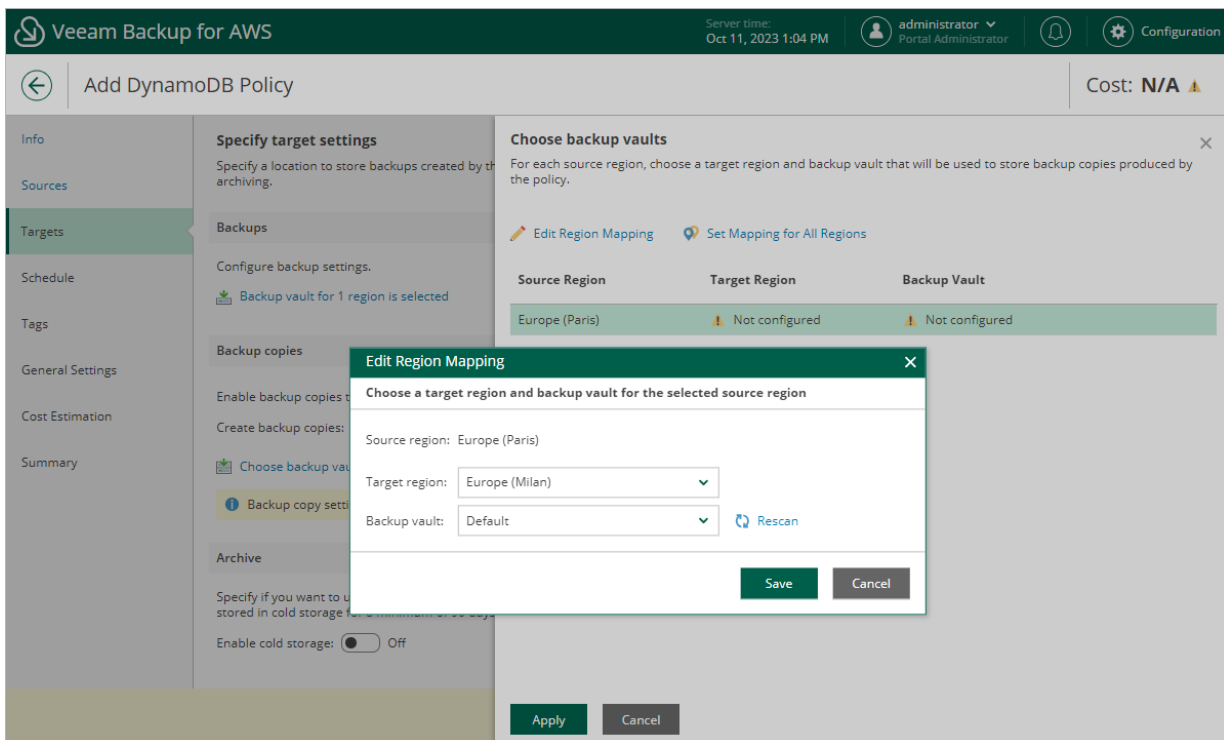
Consider the following:

- 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**.



The screenshot displays the Veeam Backup for AWS configuration interface. The main window is titled 'Add DynamoDB Policy' and shows a sidebar with navigation options like 'Info', 'Sources', 'Targets', 'Schedule', 'Tags', 'General Settings', 'Cost Estimation', and 'Summary'. The 'Targets' section is expanded, showing 'Backups' and 'Backup copies' settings. The 'Specify target settings' section is active, and the 'Choose backup vaults' section shows a table with 'Europe (Paris)' as the source region and 'Not configured' for target region and backup vault. An 'Edit Region Mapping' dialog box is open, allowing selection of a target region (Europe (Milan)) and a backup vault (Default) for the source region (Europe (Paris)).

Source Region	Target Region	Backup Vault
Europe (Paris)	Not configured	Not configured

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](#).

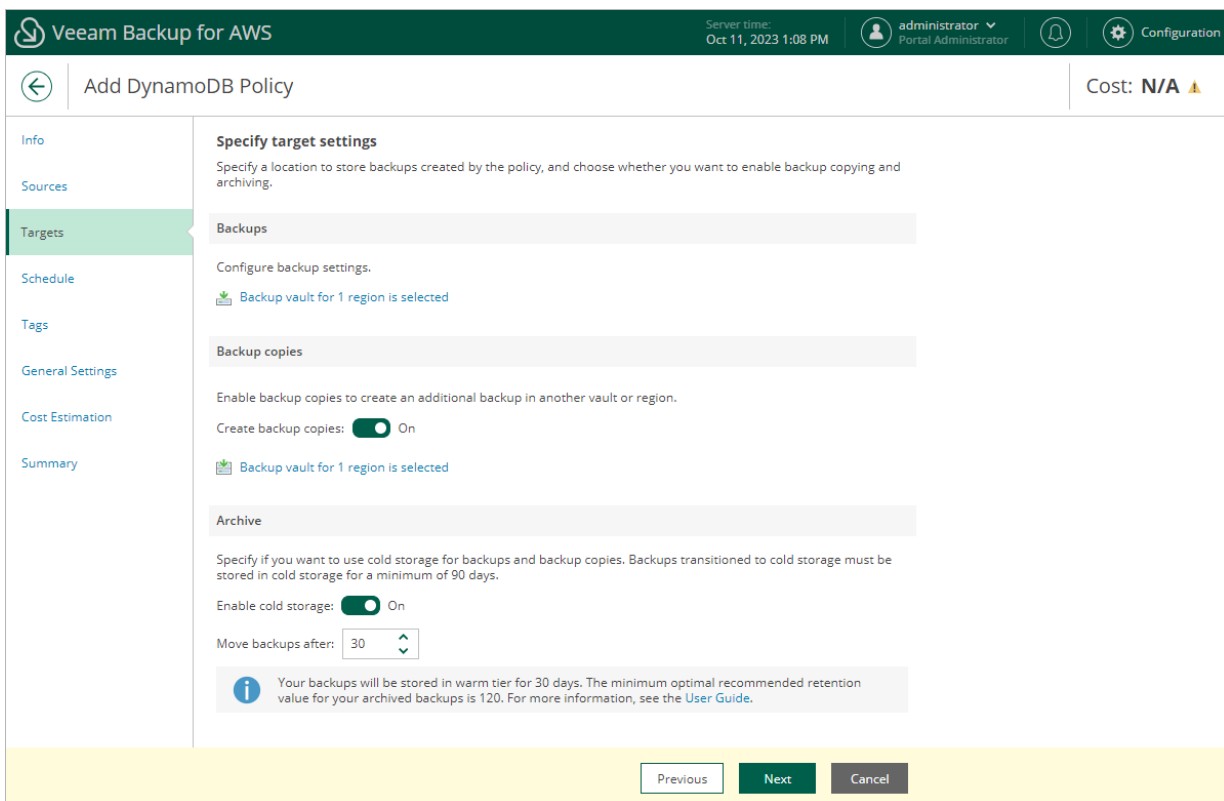
2. In the **Move backups after** field, specify the number of days for which you want to keep backups in a warm storage tier before moving them to a cold storage tier (the minimum value is 1; the maximum value is 36,135). As soon as the specified period is over, the backups will be moved to the cold 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.



The screenshot shows the Veeam Backup for AWS console interface. At the top, the header includes the Veeam logo, the text 'Veeam Backup for AWS', the server time 'Oct 11, 2023 1:08 PM', and user information 'administrator Portal Administrator'. The main navigation bar contains 'Add DynamoDB Policy' and 'Cost: N/A'. A left sidebar lists navigation options: Info, Sources, Targets (highlighted), Schedule, Tags, General Settings, Cost Estimation, and Summary. The main content area is titled 'Specify target settings' and includes instructions: 'Specify a location to store backups created by the policy, and choose whether you want to enable backup copying and archiving.' The 'Backups' section shows 'Backup vault for 1 region is selected'. The 'Backup copies' section has 'Create backup copies' set to 'On'. The 'Archive' section has 'Enable cold storage' set to 'On' and 'Move backups after' set to '30'. A warning icon and message at the bottom of the main area reads: '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.' At the bottom of the page are 'Previous', 'Next', and 'Cancel' buttons.

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

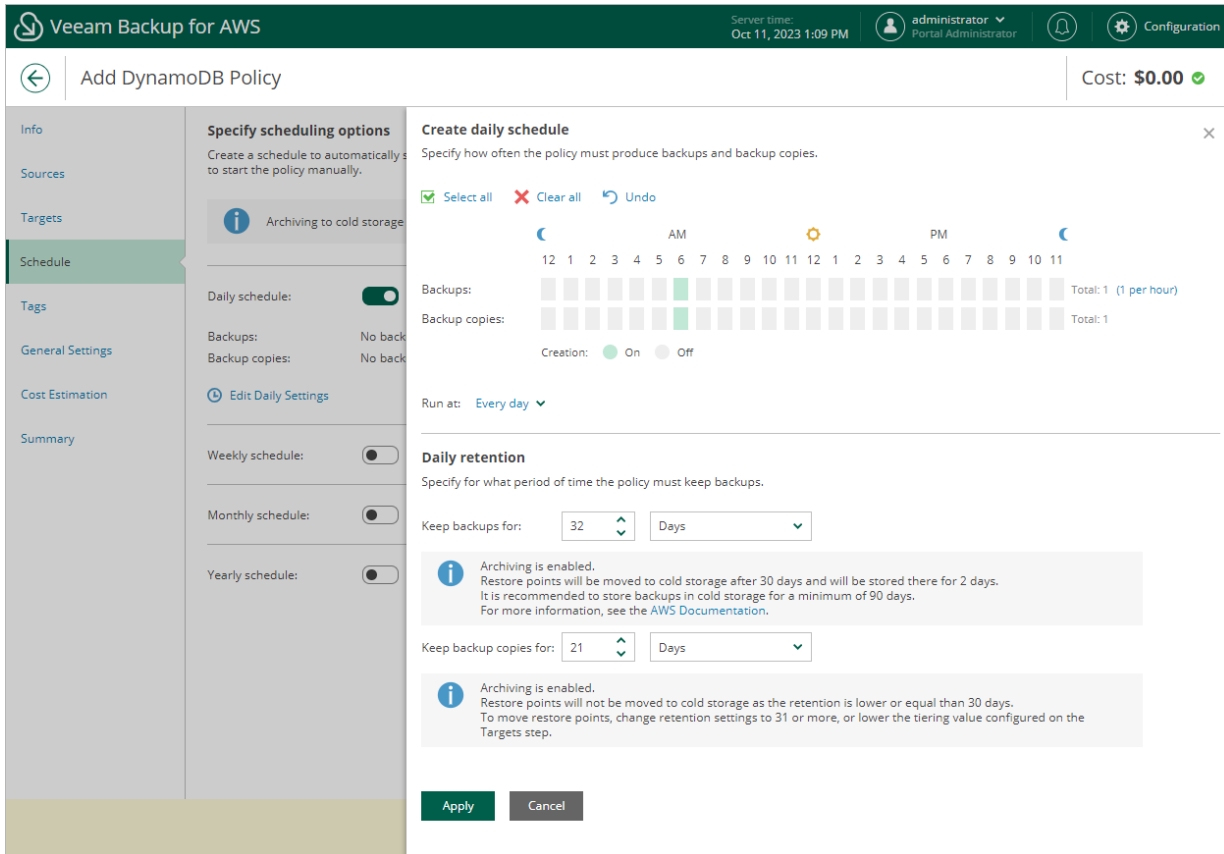
Veeam Backup for AWS does not create backup copies independently from table backups. That is why when you select hours for backup copies, the same hours are automatically selected for backups. To learn how Veeam Backup for AWS performs backup, see [DynamoDB Backup](#).

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.

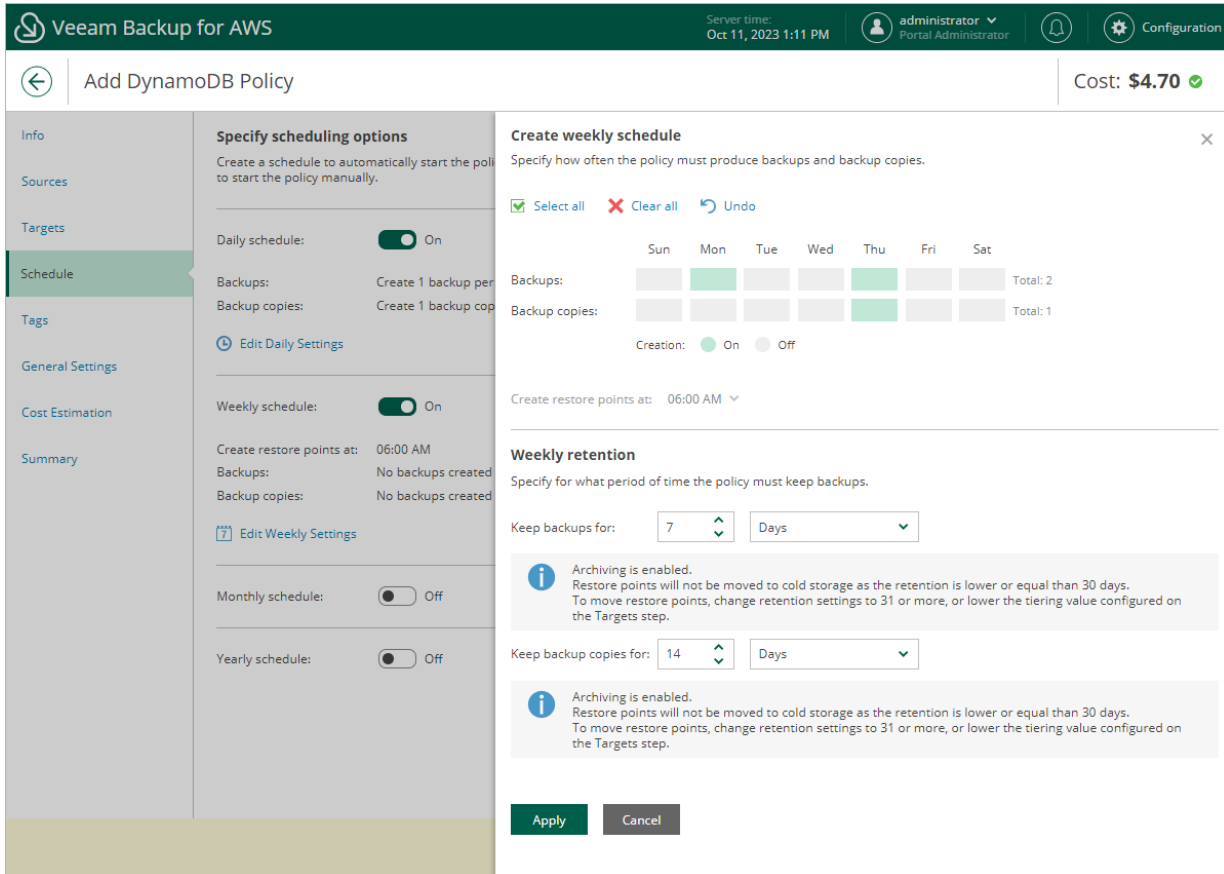
NOTE

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

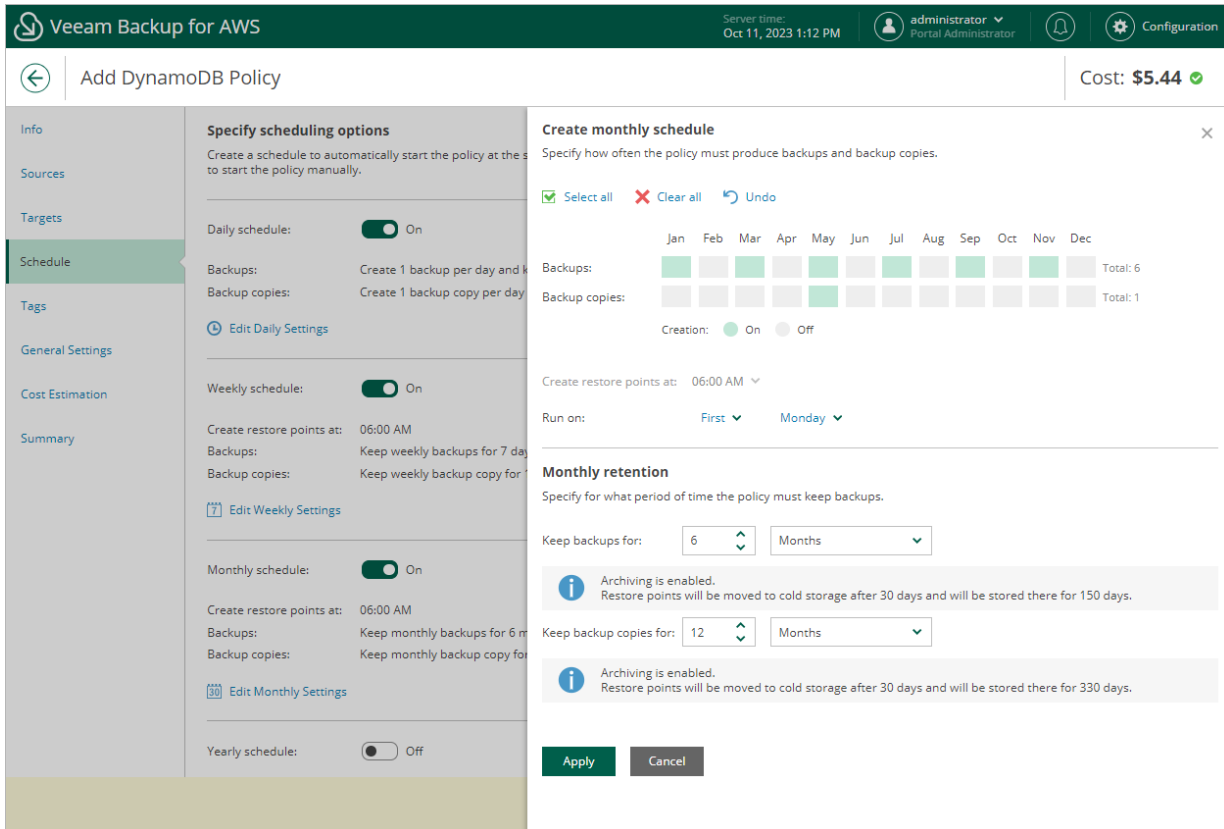
Consider the following:

- 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

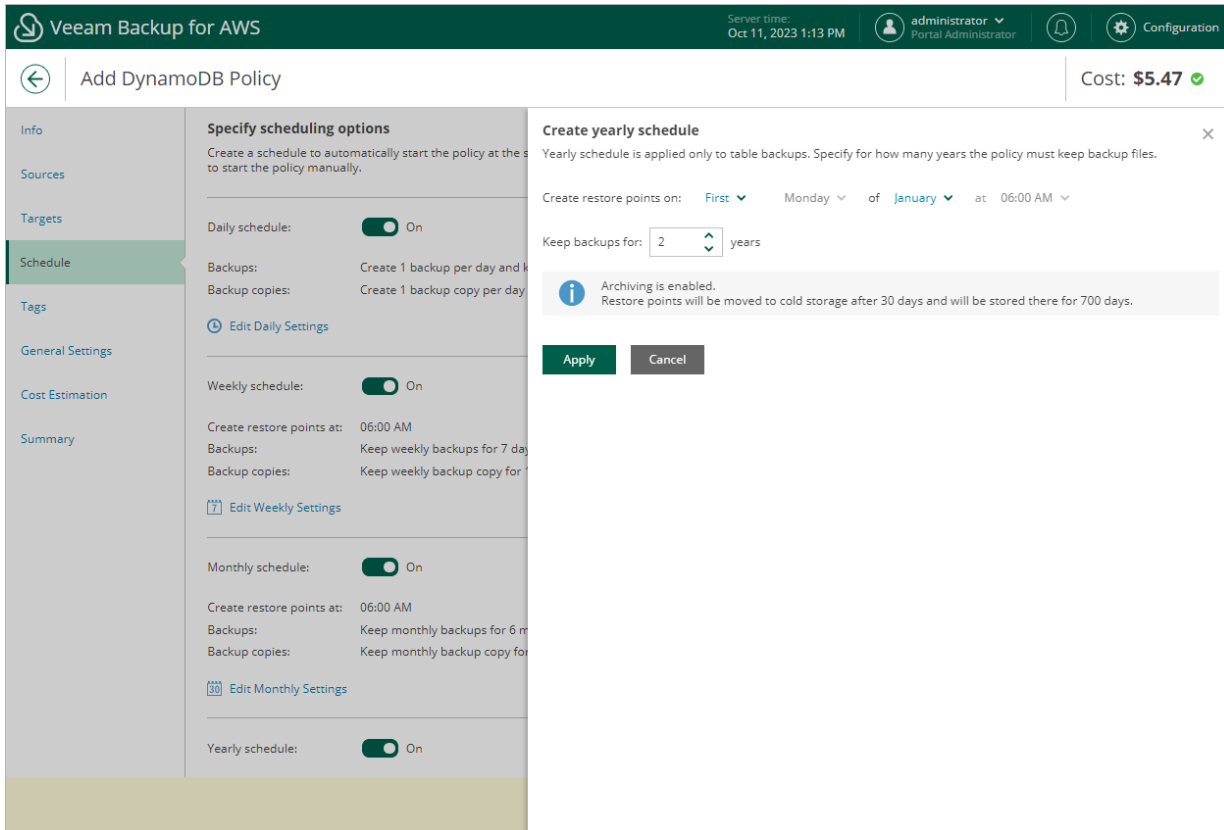
Consider the following:

- 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.

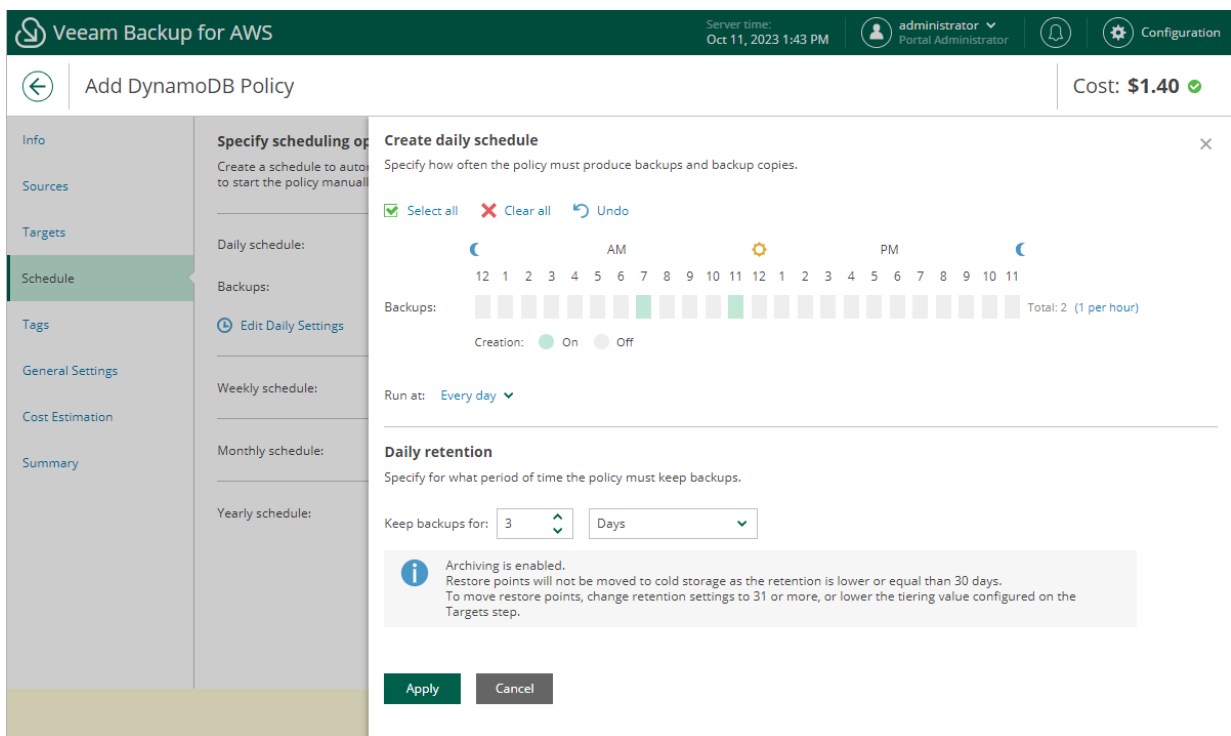
With harmonized scheduling, Veeam Backup for AWS can keep restore points created according to a daily, weekly or monthly schedule for longer periods of time: DynamoDB backups and backup copies can be kept for weeks, months and years.

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:

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 configuration area is titled "Add DynamoDB Policy" and shows a cost of \$0.60. The "Specify scheduling options" section is active, with the following settings:

- Daily schedule:** On
- Backups:** Create 2 backups per day
- Weekly schedule:** On
- Create restore points at:** 07:00 AM
- Backups:** No backups created
- Monthly schedule:** Off
- Yearly schedule:** Off

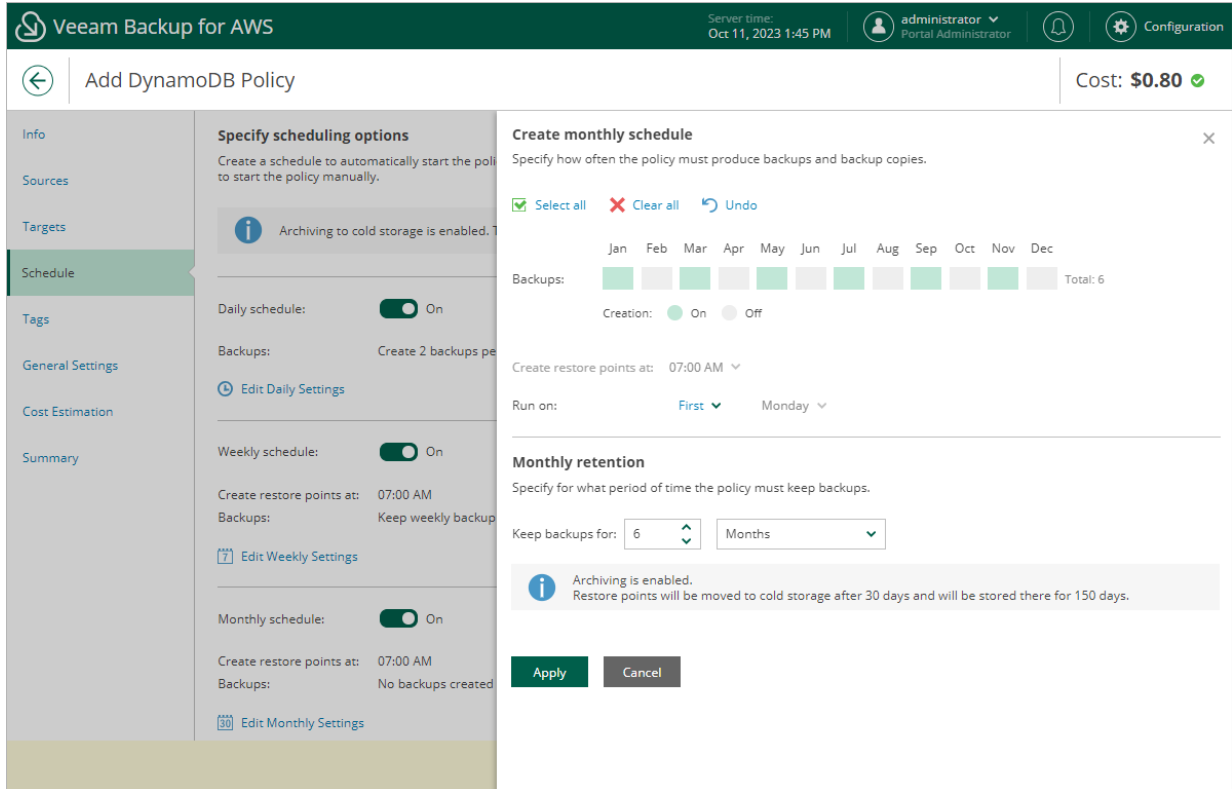
The "Create weekly schedule" dialog is open, showing the following settings:

- Backups:** Monday is selected (Total: 1)
- Create restore points at:** 07:00 AM
- Weekly retention:** Keep backups for 14 Days

A warning message is displayed: "Archiving is enabled. Restore points will not be moved to cold storage as the retention is lower or equal than 30 days. To move restore points, change retention settings to 31 or more, or lower the tiering value configured on the Targets step."

- 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).



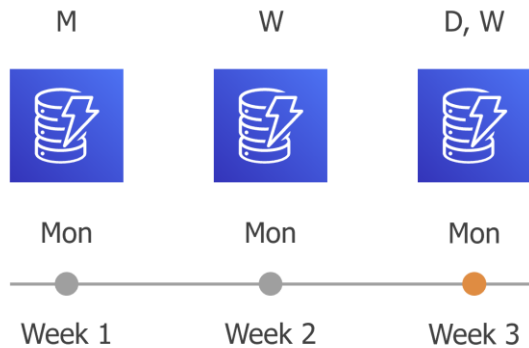
- On the fifth working day (Friday), after a backup session runs at 7:00 AM, the created restore point will be marked with the (D) flag.

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 snapshots, 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 the Veeam Backup for AWS interface. The 'Tags' step is active, showing options to copy tags from source tables and add custom tags. The 'Add custom tags to created backups' toggle is set to 'On'. A custom tag is added with the key 'owner' and value 'dept01'. A list of tags shows 'department: accounting' with a close button. The cost is \$0.72. Navigation buttons 'Previous', 'Next', and 'Cancel' are at the bottom.

Veeam Backup for AWS Server time: Oct 11, 2023 1:14 PM administrator Portal Administrator Configuration

← Add DynamoDB Policy Cost: \$0.72 ✓

Specify tag settings
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.

Copy tags from source tables

Add custom tags to created backups: On

Key: owner Value: dept01 + Add

department: accounting x

A maximum of 5 custom tags is allowed.

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 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 send each notification to this recipient twice.

The screenshot shows the 'Add DynamoDB Policy' configuration page in Veeam Backup for AWS. The interface includes a top navigation bar with the Veeam logo, server time (Oct 11, 2023 1:18 PM), user profile (administrator), and configuration settings. The main content area is divided into a left sidebar with navigation options (Info, Sources, Targets, Schedule, Tags, General Settings, Cost Estimation, Summary) and a main panel for 'Policy settings'. The 'General Settings' tab is active, showing options for 'Automatically retry failed policy' (set to 3 times), a notification email address (donna_ortiz@company.com), and checkboxes for 'Notify on' (Failure, Warning, Success) and 'Suppress notifications until the last retry'. A cost estimation of \$5.54 is displayed in the top right corner. At the bottom, there are 'Previous', 'Next', and 'Cancel' buttons.

Veeam Backup for AWS Server time: Oct 11, 2023 1:18 PM administrator Portal Administrator Configuration

← Add DynamoDB Policy Cost: **\$5.54** ✓

Policy settings
Specify retry times for the policy and e-mail notifications

Schedule

Automatically retry failed policy: 3 times

Notifications

Enabled: On

Email: donna_ortiz@company.com

Notify on

Failure

Warning

Success

Suppress notifications until the last retry

Previous Next 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](#). 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 transition 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.

Veeam Backup for AWS Server time: Oct 11, 2023 1:21 PM administrator Portal Administrator Configuration

← Add DynamoDB Policy Cost: **\$5.50** ✓

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](#).

\$1.79 Backups **\$2.34** Backup Copies **\$1.37** Traffic

Estimated monthly cost: \$5.50

Table Export to... ▼

Table ↑	Backup	Backup Copy	Traffic
DataTable	\$1.79	\$2.34	\$1.37

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, 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.

Veeam Backup for AWS

Server time: Oct 11, 2023 1:23 PM administrator Portal Administrator Configuration

← Add DynamoDB Policy Cost: \$5.51 ✓

Info Sources Targets Schedule Tags General Settings Cost Estimation Summary

Review configured settings

Review the settings, and click Finish to exit the wizard.

! In order to successfully run this policy, we advise to test the configuration.

⚙️ Test Configuration 📄 Copy to Clipboard

General

Name: DynamoDB backup policy 02
Description: Created by administrator at 10/12/2023 1:27 PM
Regions: Europe (Paris)
Account: Backup role

Backup settings

Copy tags from source tables: Yes
Add custom tags: Yes
Custom tags: department:accounting
owner:dept01

Backup copy settings

Enabled: Yes
Region mapping: Source region: Europe (Paris) Target region: Europe (Milan)

Archive Settings

Move to cold storage: Yes
Move backups to cold storage after: 30 days

Previous Finish Cancel

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 > 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.
- 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 policy 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.

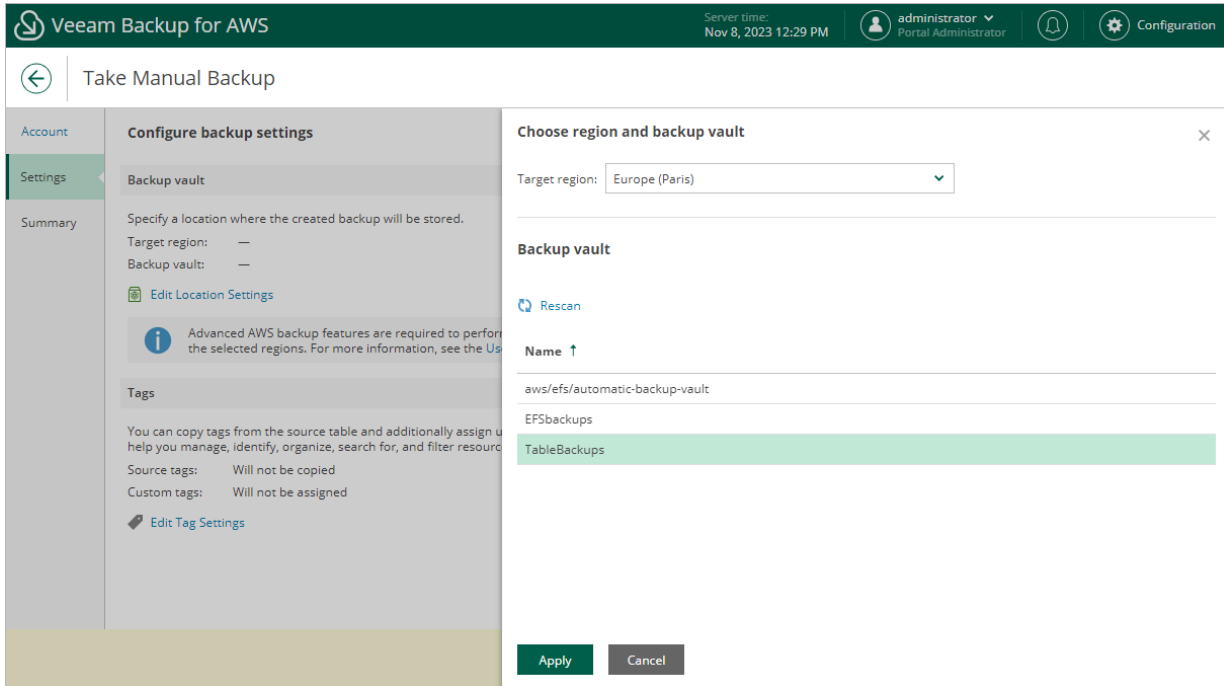
If you choose to copy tags from the source table, Veeam Backup for AWS will first create a backup of the DynamoDB table and assign to the created backup AWS tags with Veeam metadata, then Veeam Backup for AWS will copy tags from the processed table and, finally, assign the copied AWS tags to the backup.

- 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 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 global notification settings before creating an EFS backup policy. For more information, see [Configuring Global Notification Settings](#).

For EFS systems residing in any of the regions added to the backup policies, 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

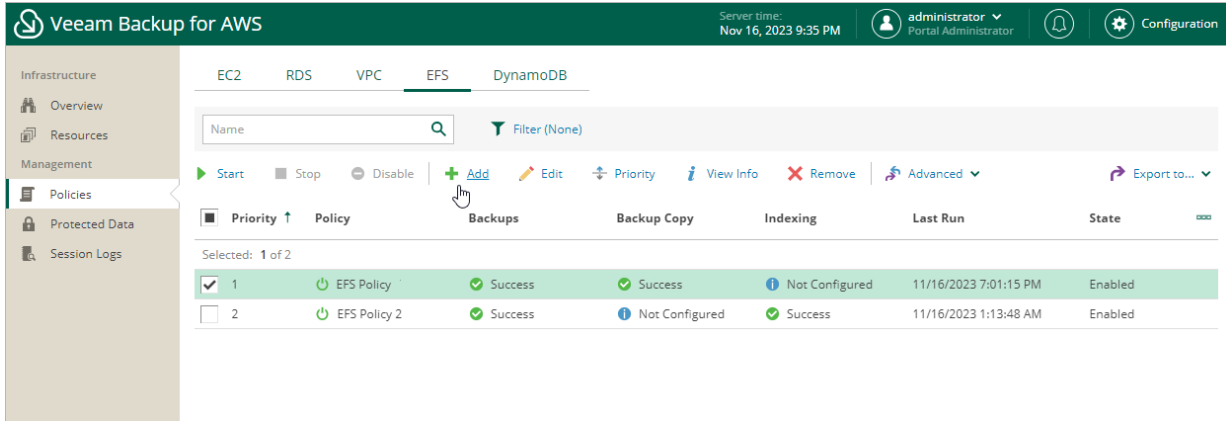
To create an EFS backup policy, do the following:

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 for protecting 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 > 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, server time (Nov 16, 2023 9:36 PM), user information (administrator, Portal Administrator), and a Configuration icon. The main header shows a back arrow, the title 'Add EFS Policy', and the cost 'N/A' with a warning icon. A left sidebar contains a navigation menu with 'Info' selected, and other options: Sources, Indexing, 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.' Below this are two input fields: 'Name:' with the value 'EFS Backup Policy' and 'Description:' with the value 'Backup of file system for D01'. At the bottom of the form 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 add an IAM role, click **Add** and complete the [Add IAM Role](#) wizard.

IMPORTANT

It is recommended that you check whether the selected IAM role has all the required permissions 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 'Sources' step. The 'IAM role' dropdown menu is open, displaying two options: 'Default Backup Restore (Default Backup Restore)' and 'EFS role (Created by bd-regress-2 at 11/14/2023 6:01 PM)'. The 'Add' button is highlighted in green. The 'Check Permissions' button is also visible. The 'Regions' section is currently empty, and the 'Resources' section is also empty. The 'Previous', 'Next', and 'Cancel' buttons are located at the bottom of the wizard.

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**.

The screenshot shows the 'Add EFS Policy' wizard in Veeam Backup for AWS. The interface is divided into several sections:

- Header:** Veeam Backup for AWS logo, server time (Nov 16, 2023 9:38 PM), user (administrator), and Configuration icon.
- Navigation:** A sidebar on the left with tabs for Info, Sources (selected), Indexing, Targets, Schedule, Tags, General Settings, Cost Estimation, and Summary.
- Specify source settings:** A section with an IAM role dropdown set to 'Default Backup Restore (Default Backup Restore)'. Below it, a 'Regions' section with the text 'Specify one or more regions.' and a 'Choose regions...' button.
- Resources:** A section with the text 'Specify resources to protect or exclude.' and two buttons: 'Choose resources to protect...' and 'Choose resources to exclude...'.
- Choose regions:** A modal window with two columns: 'Available Regions (22)' and 'Selected Regions (1)'. The 'Available Regions' list includes: Asia Pacific (Jakarta), Asia Pacific (Mumbai), Asia Pacific (Osaka), Asia Pacific (Seoul), Asia Pacific (Sydney), Asia Pacific (Tokyo), Canada (Central), Europe (Frankfurt) (highlighted), Europe (Ireland), Europe (London), Europe (Milan), Europe (Paris), Europe (Spain), and Europe (Stockholm). To the right of the list are 'Add' and 'Remove' buttons. The 'Selected Regions' list contains: Asia Pacific (Singapore). At the bottom of the modal are 'Apply' and 'Cancel' buttons.
- Cost:** A label 'Cost: N/A' with a warning icon.

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.

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.

Choose resources to protect

All resources

Protect only following resources

Type: Name or ID: Protect

[Browse to select specific resources from the global list...](#)

Protected resources (3)

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	—

[Apply](#) [Cancel](#)

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 launches 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 launch these worker instances. However, you can add [specific worker configurations](#) that will be used to launch 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 on license editions, see [Licensing of Standalone Backup Appliances](#).
- 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 launch worker instances for EFS indexing operations. For Veeam Backup for AWS to be able to launch 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](#). 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 [Indexing Worker IAM Role Permissions](#).

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 where 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' configuration window in Veeam Backup for AWS. The 'Indexing' tab is selected, with 'Enable indexing' turned on. The 'Repositories' dialog is open, showing a table of available repositories. The table has columns for Repository, Region, Folder, and Description. One repository is listed: 'Singapore repo' in the 'Asia Pacific (Singapor...)' region, with the folder 'Import' and description 'Created by bd-regres...'. The dialog also includes a search bar, a 'Rescan' button, and 'Apply' and 'Cancel' buttons at the bottom.

Repository	Region	Folder	Description
Singapore repo	Asia Pacific (Singapor...	Import	Created by bd-regres...

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.

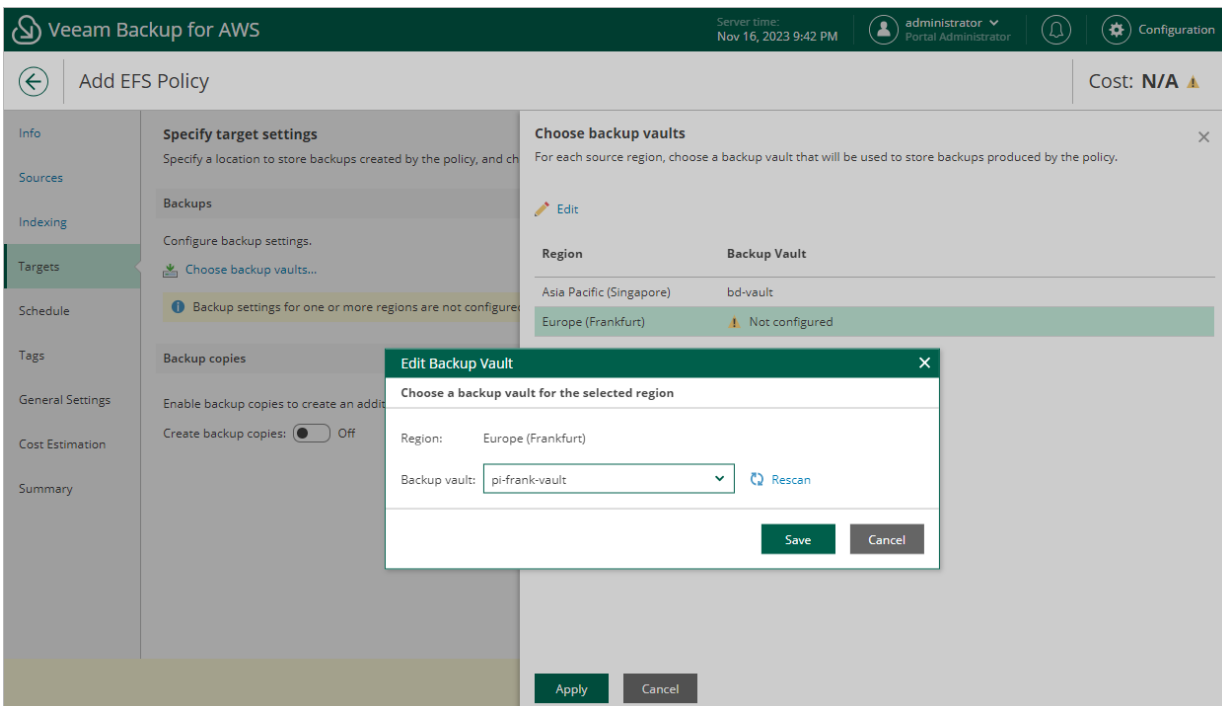
For a backup vault to be displayed in the **Backup vault** list, it must be created in the AWS Backup console as described in [AWS Documentation](#). If you have not created a backup vault for the selected AWS Region, Veeam Backup for AWS will display only the default backup vault existing in this region.

IMPORTANT

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:

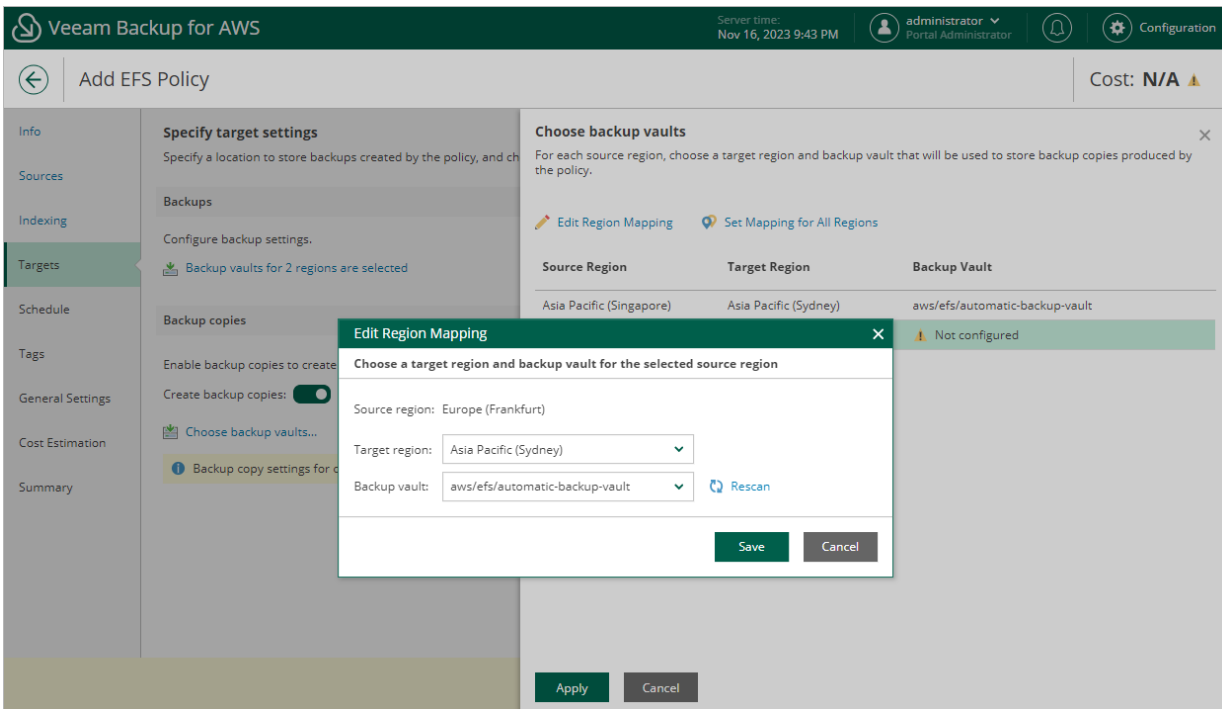
- i. From the **Target region** drop-down list, select the target AWS Region to which Veeam Backup for AWS must copy created backups of the selected file systems.
- ii. From the **Backup vault** drop-down list, select a backup vault that will be used to store the copied backups.

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.

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.

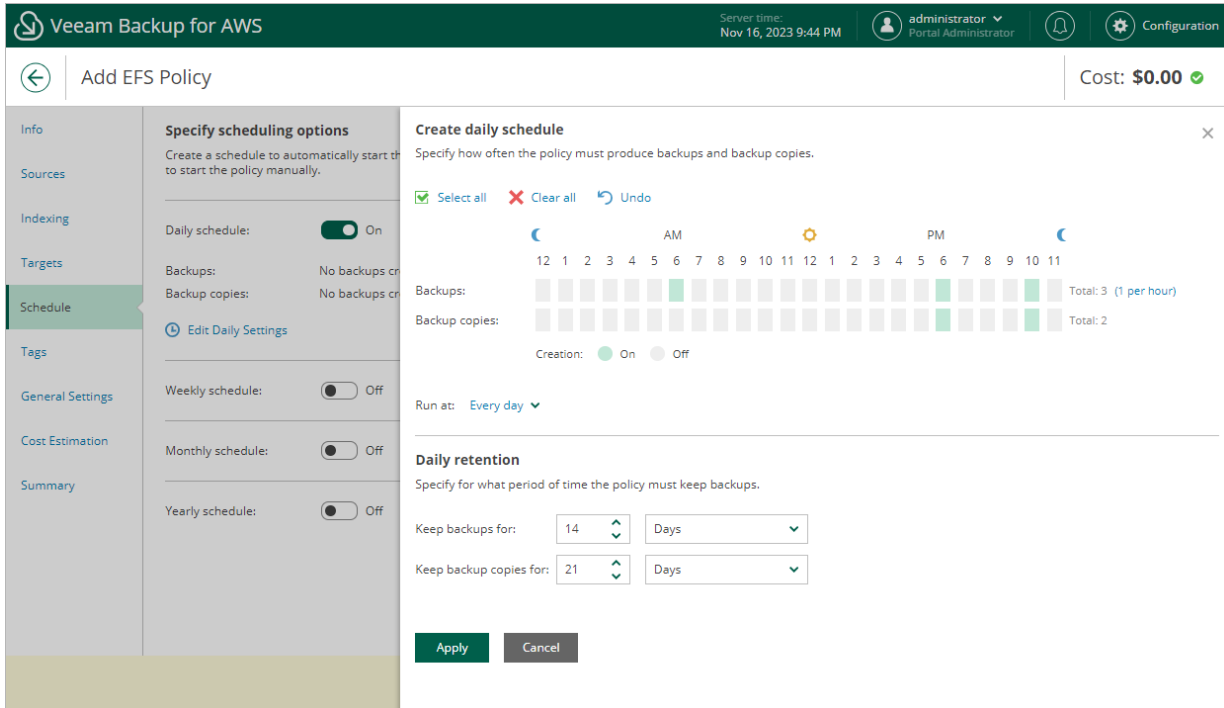
NOTE

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.

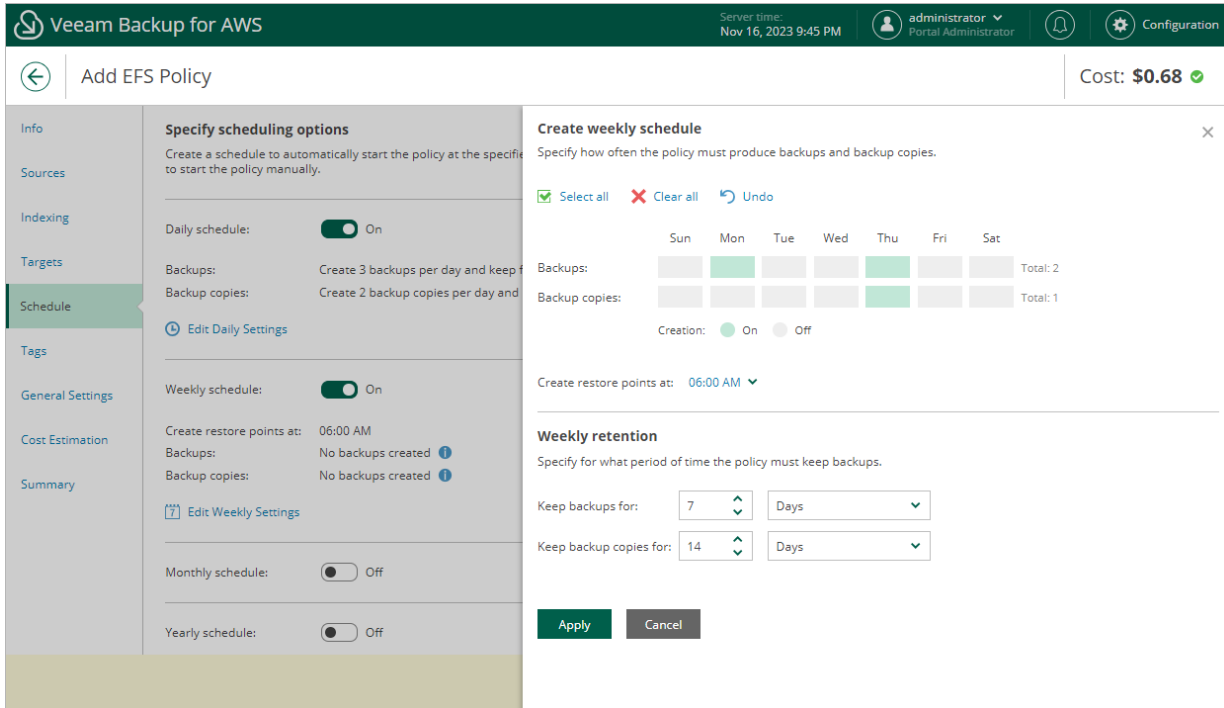
NOTE

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.

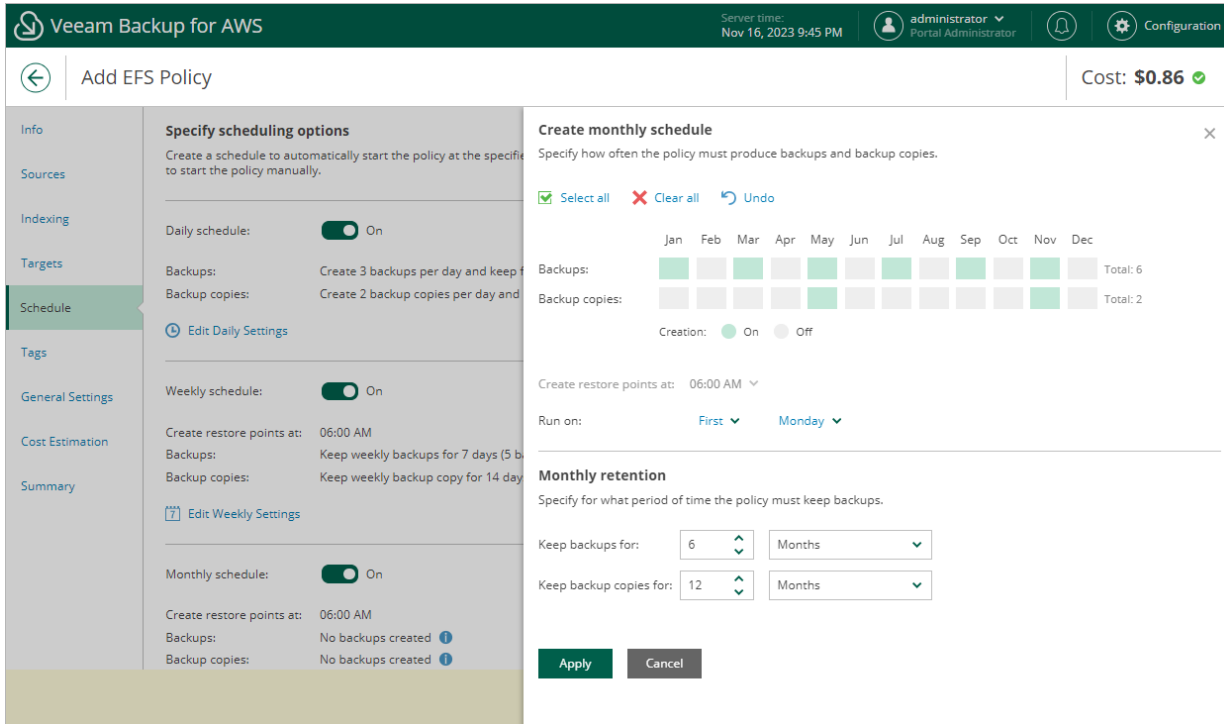
NOTE

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.

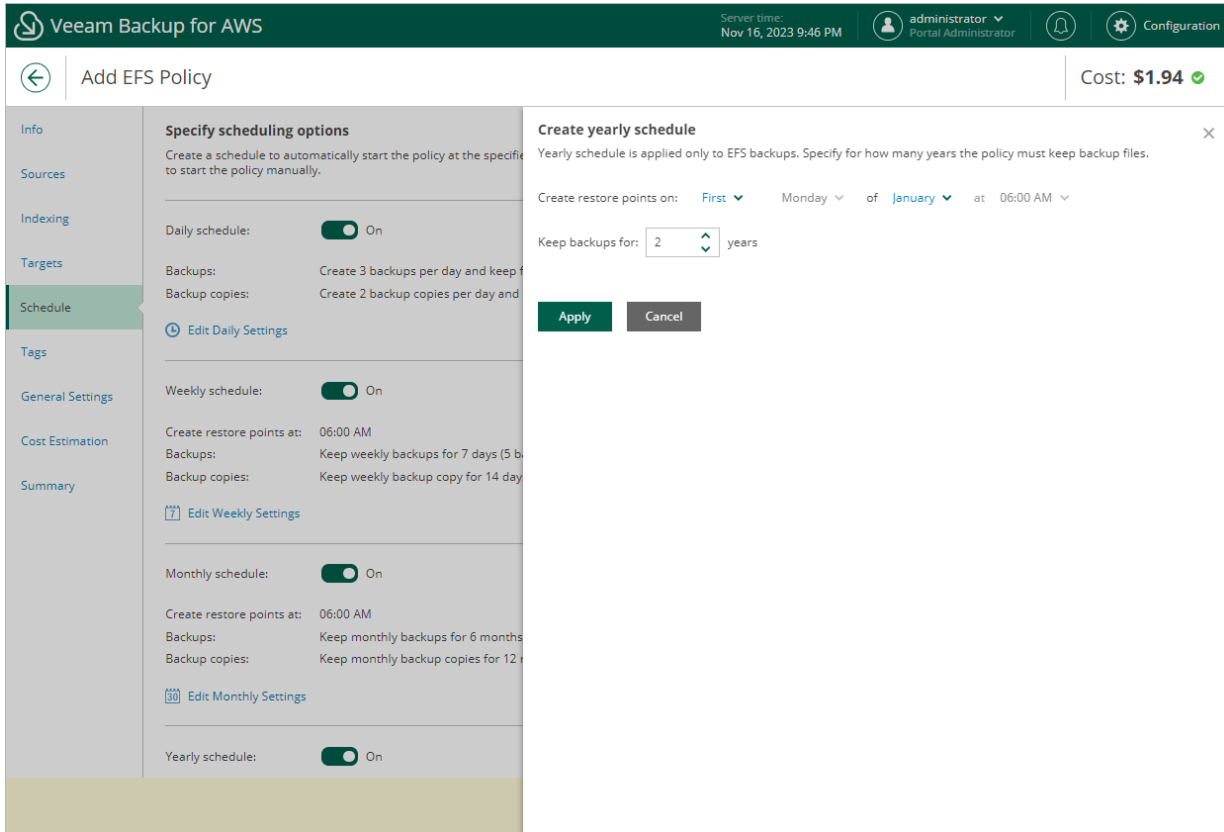
NOTE2

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.

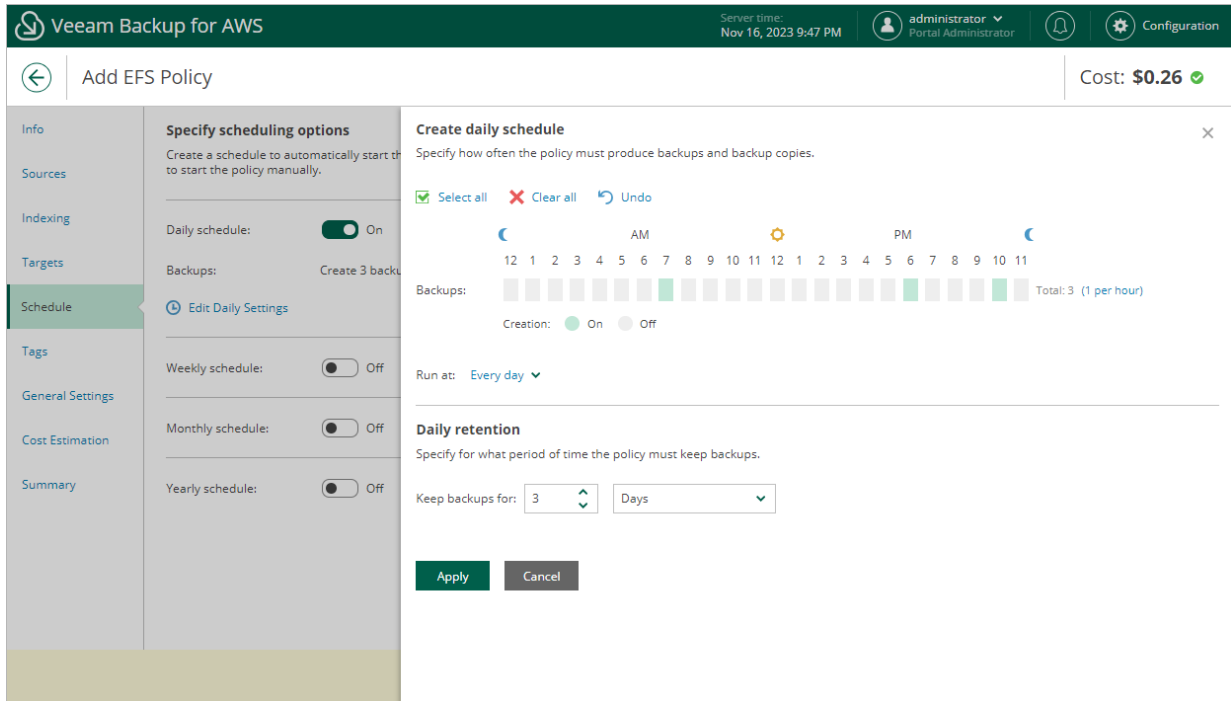
With harmonized scheduling, Veeam Backup for AWS can keep restore points created according to a daily, weekly or monthly schedule for longer periods of time: EFS backups and backup copies can be kept for weeks, months and years.

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:

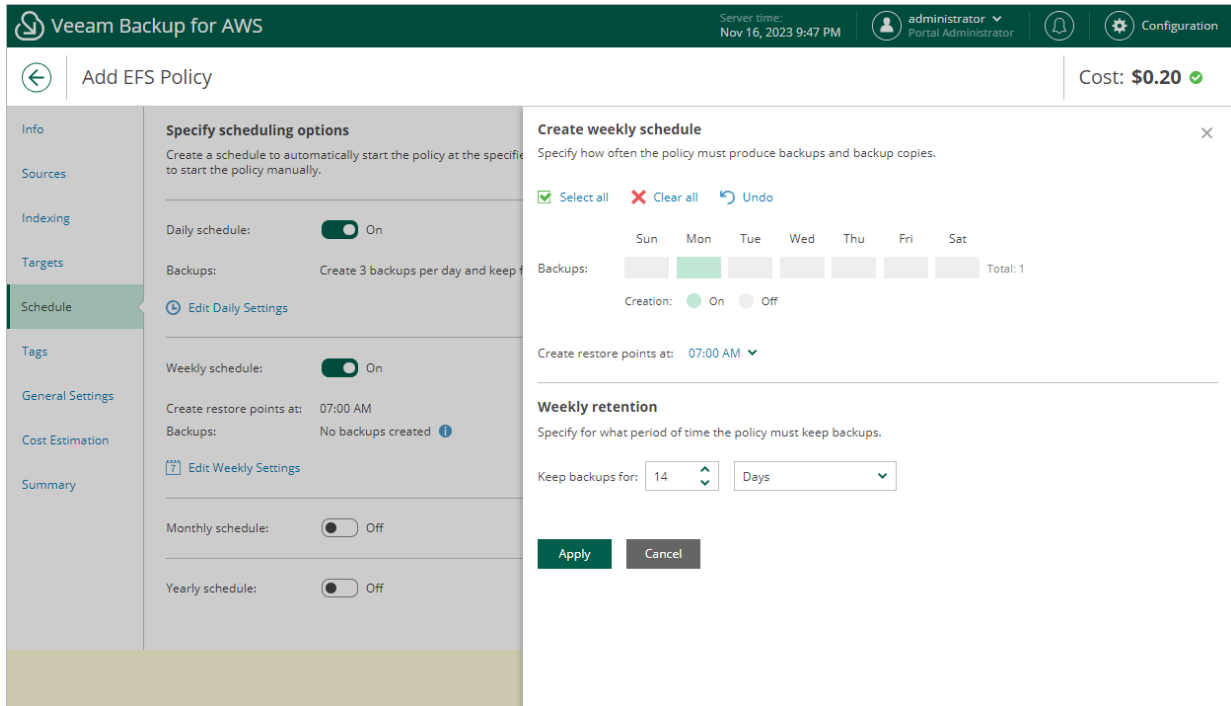
- In the daily scheduling settings, you select hours and days when backups will be created (for example, *7:00 AM; Working Days*), and specify a number of days for which you want to keep daily restore points in a backup chain (for example, 3).

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 the Veeam Backup for AWS console. The 'Tags' step is active, showing the 'Specify tag settings' section. The 'Copy tags from source file systems' checkbox is checked. The 'Add custom tags to created backups' toggle is set to 'On'. There are two custom tags defined: one with key 'user' and value 'donna_ortiz', and another with key 'owner' and value 'dept01'. An 'Add' button is visible next to the first tag. The interface includes a navigation sidebar on the left, a top header with user information and server time, and a bottom navigation bar with 'Previous', 'Next', and 'Cancel' buttons. A cost indicator of '\$2.10' is shown in the top right corner.

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 send each notification to this recipient twice.

The screenshot shows the 'Add EFS Policy' configuration page in Veeam Backup for AWS. The top navigation bar includes the Veeam logo, 'Veeam Backup for AWS', server time 'Nov 16, 2023 9:51 PM', user 'administrator Portal Administrator', and a 'Configuration' gear icon. The page title is 'Add EFS Policy' with a back arrow and a cost indicator 'Cost: \$2.10'. A left sidebar lists navigation options: Info, Sources, Indexing, Targets, Schedule, Tags, General Settings (highlighted), Cost Estimation, and Summary. The main content area is titled 'Configure retry and notification settings' and includes the instruction: 'Specify how many times to retry the policy. You can also enable email notifications to receive policy results.' Under the 'Schedule' section, there is a checkbox for 'Automatically retry failed policy:' set to '3' times, with a note: 'Automatic retry settings are only applicable on a scheduled run of the policy'. Under the 'Notifications' section, 'Enabled:' is a toggle switch set to 'On', and the 'Email:' field contains 'donna_ortiz@company.mail'. The 'Notify on' section has checkboxes for 'Failure', 'Warning', and 'Success', all of which are checked. There is also a checkbox for 'Suppress notifications until the last retry' which is checked. At the bottom, there are 'Previous', 'Next', and 'Cancel' buttons.

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.

To calculate the estimated cost, Veeam Backup for AWS uses capabilities of the [AWS Pricing Calculator](#).

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 EFS Policy' configuration page in Veeam Backup for AWS. The top navigation bar includes the Veeam logo, 'Veeam Backup for AWS', server time (Nov 16, 2023 9:51 PM), user (administrator), and a Configuration icon. The page title is 'Add EFS Policy' with a back arrow and a cost indicator 'Cost: \$2.10'. A left sidebar lists navigation options: Info, Sources, Indexing, Targets, Schedule, Tags, General Settings, Cost Estimation (highlighted), and Summary.

The main content area is titled 'Review cost estimation'. It includes a note: 'The estimated cost takes into account the configured target settings, the specified scheduling options and the number of resources to protect.' and another note: '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.'

Three cost breakdown cards are shown: Backups (\$1.18), Backup Copies (\$0.71), and Traffic (\$0.21). A large green box displays the 'Estimated monthly cost: \$2.10'. Below this is a search bar for 'File system' and an 'Export to...' dropdown menu.

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	

At the bottom, there are 'Previous', 'Next', and 'Cancel' buttons.

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: Nov 16, 2023 9:52 PM administrator Portal Administrator Configuration

← Add EFS Policy Cost: \$2.10 ✓

Review configured settings
Review the settings, and click Finish to exit the wizard.

In order to successfully run this policy, we advise to test the configuration.

Test Configuration Copy to Clipboard

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 > 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.
 - 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 during the 'Take Manual Backup' wizard. The top navigation bar includes the Veeam logo, 'Veeam Backup for AWS', server time 'May 27, 2022 10:27 AM', and user information 'administrator Portal Administrator'. The main content area is divided into three sections: 'Account', 'Settings', and 'Summary'. The 'Summary' section is active, displaying the following information:

- Configure backup settings**
- Backup vault**: Specify a location where the created backup will be stored. Target region: —, Backup vault: —. [Edit Location Settings](#)
- Tags**: You can instruct Veeam Backup for AWS to copy tags from the source resources to the created backup. Each tag consists of a user-defined key and value. You can identify, organize, search for, and filter resources. Source tags: Will not be copied. Custom tags: Will not be assigned. [Edit Tag Settings](#)

On the right, a modal window titled 'Choose region and backup vault' is open. It shows 'Target region' set to 'Canada (Central)'. Under 'Backup vault', there is a 'Rescan' button and a list of vaults. The vault 'aws/efs/automatic-backup-vault' is selected and highlighted in green. Below the list is a 'Default' label. At the bottom of the modal are 'Apply' and 'Cancel' buttons.

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

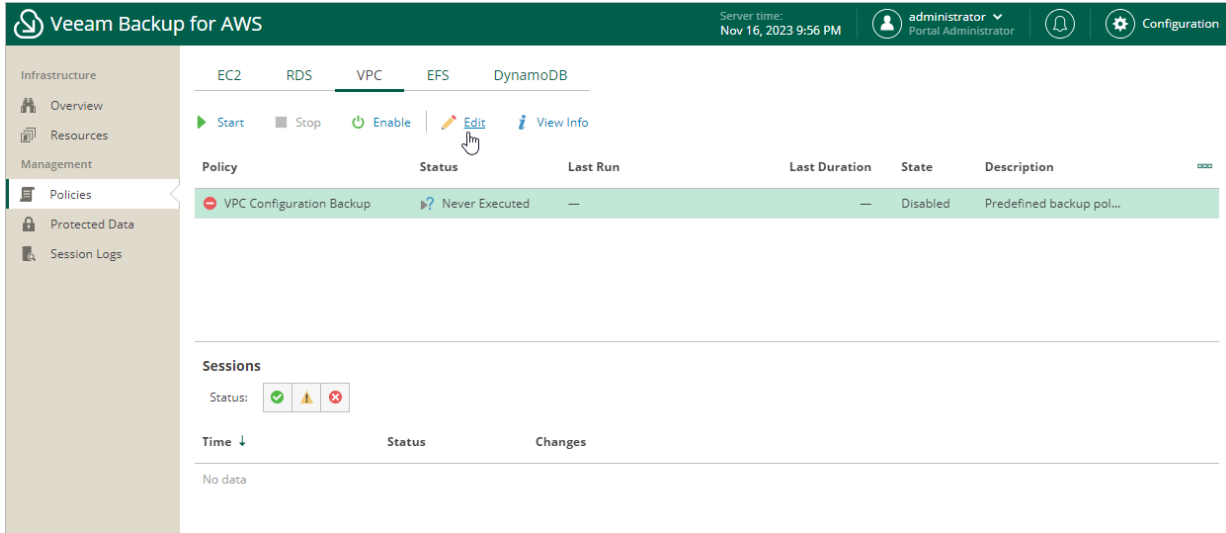
To configure the VPC Configuration Backup policy settings, do the following:

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 and EFS 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 and EFS 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 and EFS backup policies use to perform data protection operations have all the required permissions 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:

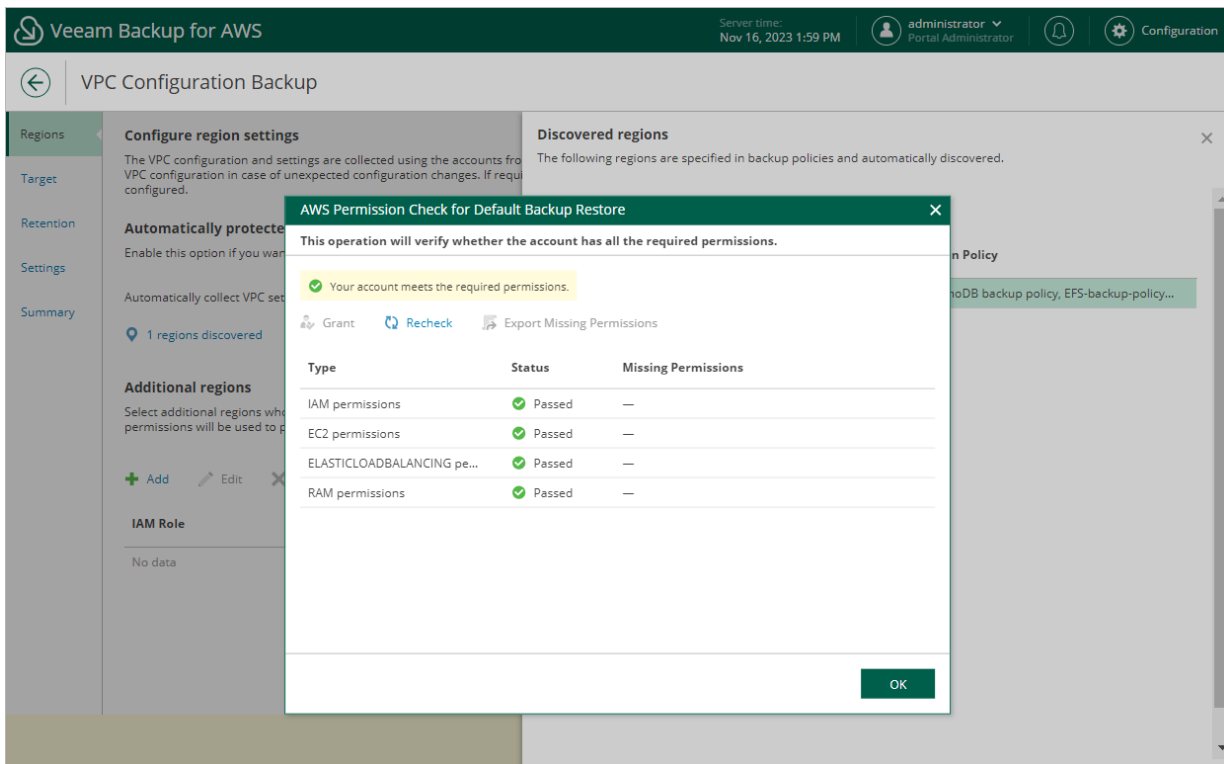
```

"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 add an IAM role, 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: Nov 16, 2023 2:00 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](#). The list shows only backup repositories of the *S3 Standard* storage class.

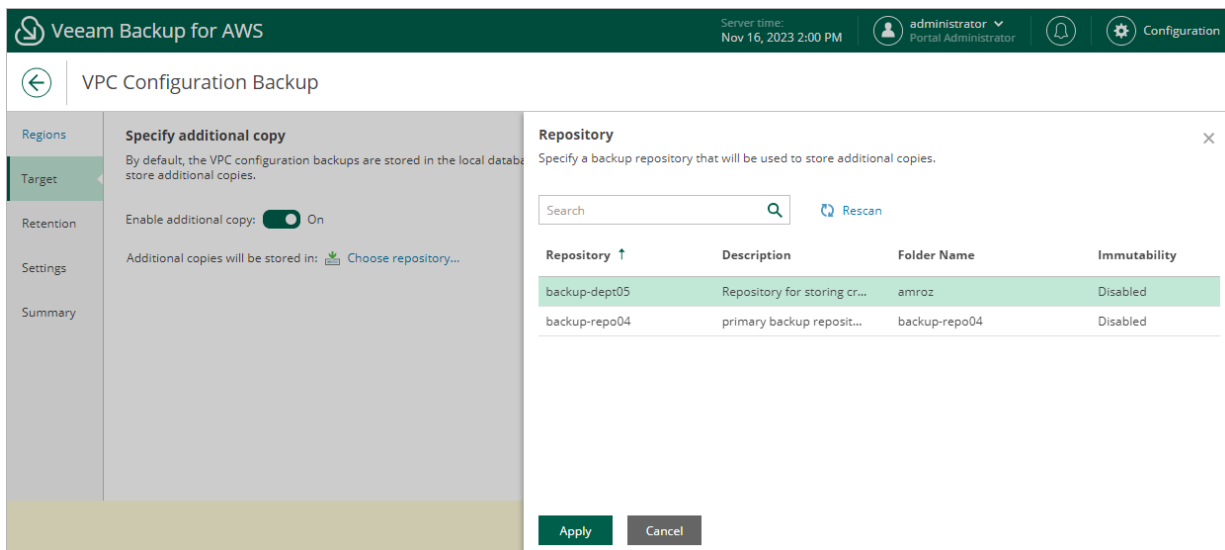
3. To save changes made to the backup policy settings, click **Apply**.

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](#).



The screenshot shows the Veeam Backup for AWS interface. The main window is titled 'VPC Configuration Backup' and is in the 'Target' step. A sidebar on the left shows navigation options: Regions, Target, Retention, Settings, and Summary. The 'Specify additional copy' section is active, showing a toggle for 'Enable additional copy' set to 'On' and a 'Choose repository...' button. A 'Repository' dialog box is open, displaying a search bar and a 'Rescan' button. Below is a table of repositories:

Repository ↑	Description	Folder Name	Immutability
backup-dept05	Repository for storing cr...	amroz	Disabled
backup-repo04	primary backup reposit...	backup-repo04	Disabled

At the bottom of the dialog are 'Apply' and 'Cancel' buttons.

Step 4. Configure Retention Settings

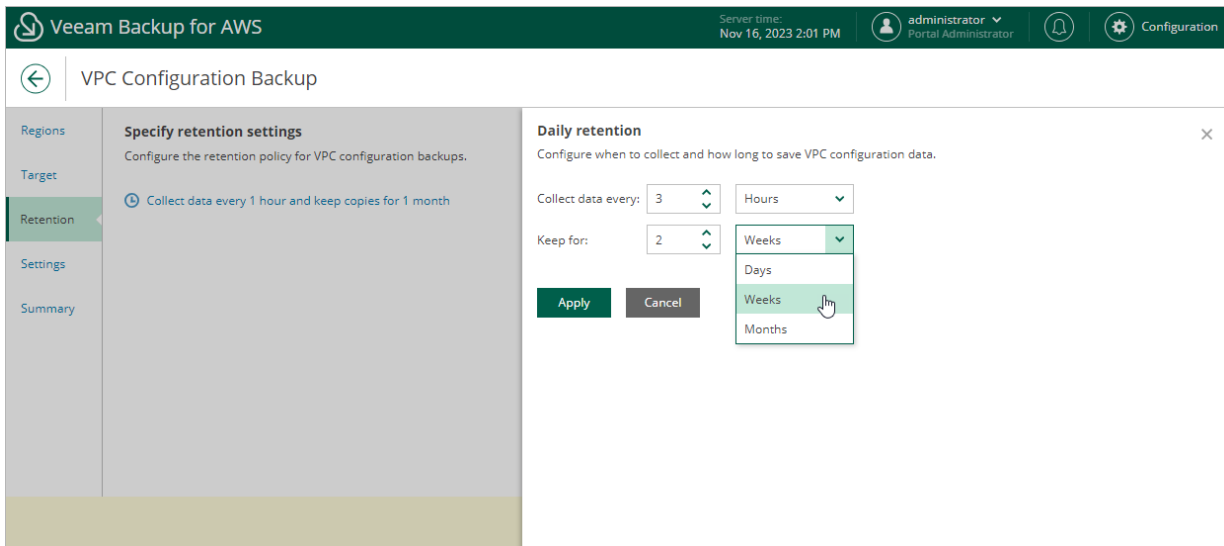
At the **Retention** step of the wizard, specify retention settings for VPC configuration backups.

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 send each notification to this recipient twice.

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, server time (Nov 16, 2023 2:02 PM), and user information (administrator, Portal Administrator). The main content area is titled 'VPC Configuration Backup' and features a left-hand navigation menu with options: Regions, Target, Retention, Settings (selected), and Summary. The 'Configure notification settings' section contains the following elements: a sub-header 'Configure notification settings' with the instruction 'Configure daily email notifications.', a 'Notifications' section with a toggle for 'Receive daily report' 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', and 'Cancel'.

Step 6. Finish Working with Wizard

At the **Summary** step of the wizard, review configuration information and click **Finish**.

The screenshot shows the 'VPC Configuration Backup' wizard in the 'Summary' step. The interface includes a navigation menu on the left with options: Regions, Target, Retention, Settings, and Summary (selected). The main content area is titled 'Review configured settings' and contains the following information:

- Regions:** Automatically protected regions: Enabled; Default Backup Restore: Europe (Milan), Europe (Stockholm).
- Target:** Additional copy: Enabled; Repository: backup-dept05.
- Notifications:** Enabled: Yes; Email: donna_ortiz@company.email; Notify on failure: Enabled; Notify on warning: Enabled; Notify on success: Enabled.

At the bottom of the wizard, there are three buttons: 'Previous', 'Finish' (highlighted in green), and 'Cancel'.

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**.

The screenshot shows the Veeam Backup for AWS console. The left sidebar contains navigation options: Infrastructure, Overview, Resources, Management, Policies (selected), Protected Data, and Session Logs. The main content area is titled 'VPC' and shows a table of policies. The 'VPC Configuration Backup' policy is highlighted in green and is in a 'Running' state.

Policy	Status	Last Run	Last Duration	State	Description
VPC Configuration Backup	Running	09/16/2021 5:01:03 PM	43 sec	Enabled	Predefined backup pol...

Below the table, there is a 'Sessions' section with a status bar showing a green checkmark, a yellow warning icon, and a red X icon. A table of sessions is displayed below:

Time	Status	Changes
09/16/2021 2:00:11 PM	Success	—
09/16/2021 2:00:11 PM	Success	3 Endpoint, 4 RouteTable, 6 SecurityGroup
09/16/2021 1:00:01 PM	Success	—
09/16/2021 12:00:04 PM	Success	21 Endpoint, 20 EndpointServices
09/16/2021 11:00:16 AM	Success	1 SecurityGroup, 4 RouteTable, 1 PeeringConnection

At the bottom of the sessions table, there is a pagination control showing 'Page 1 of 34'.

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 'Sep 16, 2021 5:01 PM', and the user 'administrator Portal Administrator'. The left sidebar shows a navigation menu with 'Policies' selected. The main content area is divided into tabs for 'EC2', 'RDS', 'VPC', and 'EFS', with 'VPC' currently active. Below the tabs, there are action buttons: 'Start', 'Stop', 'Disable', 'Edit', and 'View Info'. A table lists the backup policies, with one entry for 'VPC Configuration Backup' showing a status of 'Running'. Below this, a 'Sessions' section shows a list of backup sessions with columns for 'Time', 'Status', and 'Changes'. The sessions listed are all successful.

Policy	Status	Last Run	Last Duration	State	Description
VPC Configuration Backup	Running	09/16/2021 5:01:03 PM	43 sec	Enabled	Predefined backup pol...

Time	Status	Changes
09/16/2021 2:00:11 PM	Success	—
09/16/2021 2:00:11 PM	Success	3 Endpoint, 4 RouteTable, 6 SecurityGroup
09/16/2021 1:00:01 PM	Success	—
09/16/2021 12:00:04 PM	Success	21 Endpoint, 20 EndpointServices
09/16/2021 11:00:16 AM	Success	1 SecurityGroup, 4 RouteTable, 1 PeeringConnection

Managing EC2, RDS, DynamoDB and EFS Backup Policies

You can manage and edit created EC2, RDS, DynamoDB and EFS 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 interface. The top navigation bar includes the Veeam logo, server time (Oct 9, 2023 9:41 AM), and user information (administrator, Portal Administrator). The left sidebar shows the navigation menu with 'Policies' selected. The main content area is divided into tabs for EC2, RDS, VPC, EFS, and DynamoDB. The EC2 tab is active, showing a search bar, a filter dropdown, and a toolbar with actions like Start, Stop, Disable, Add, Edit, Priority, View Info, Remove, Advanced, and Export to... Below this is a table of backup policies. One policy is selected, showing its name, priority, status (Success), and state (Enabled). Below the table are two sections: 'Instances' and 'Sessions'. The 'Instances' section shows two instances (amroz-vm03 and amroz-vm04) with a 'Success' status. The 'Sessions' section shows three sessions (EC2 policy replication, EC2 policy backup, and EC2 policy snapshot) with a 'Success' status. The bottom of the page shows 'Page 1 of 3'.

Disabling and Enabling Policies

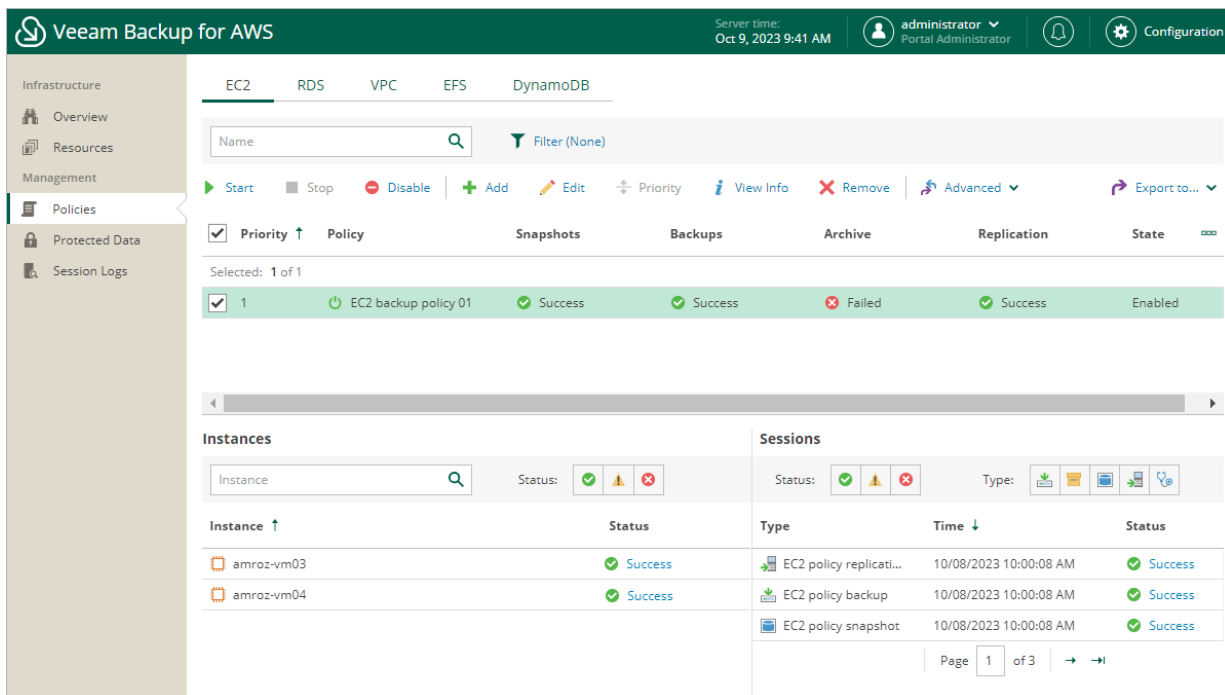
By default, Veeam Backup for AWS runs all created backup policies according to the specified schedules. However, you can temporarily disable a backup policy so that Veeam Backup for AWS does not run the backup policy automatically. You will still be able to [manually start](#) or enable the disabled 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 text 'Veeam Backup for AWS', the server time 'Oct 9, 2023 9:41 AM', and the user 'administrator Portal Administrator'. The left sidebar contains navigation options: Infrastructure, Overview, Resources, Management, Policies (selected), Protected Data, and Session Logs. The main content area is titled 'EC2' and shows a list of policies. A search bar and a 'Filter (None)' button are at the top. Below the search bar are action buttons: Start, Stop, Disable, Add, Edit, Priority, View Info, Remove, Advanced, and Export to... The policy list has columns for Priority, Policy, Snapshots, Backups, Archive, Replication, and State. One policy is selected: '1' with ID 'EC2 backup policy 01', showing Success for Snapshots, Backups, and Replication, and a Failed status for Archive. Below the policy list are two sections: 'Instances' and 'Sessions'. The 'Instances' section shows two instances: 'amroz-vm03' and 'amroz-vm04', both with a 'Success' status. The 'Sessions' section shows three sessions: 'EC2 policy replicati...', 'EC2 policy backup', and 'EC2 policy snapshot', all with a 'Success' status. The bottom right of the sessions table shows 'Page 1 of 3'.

Setting Policy Priority

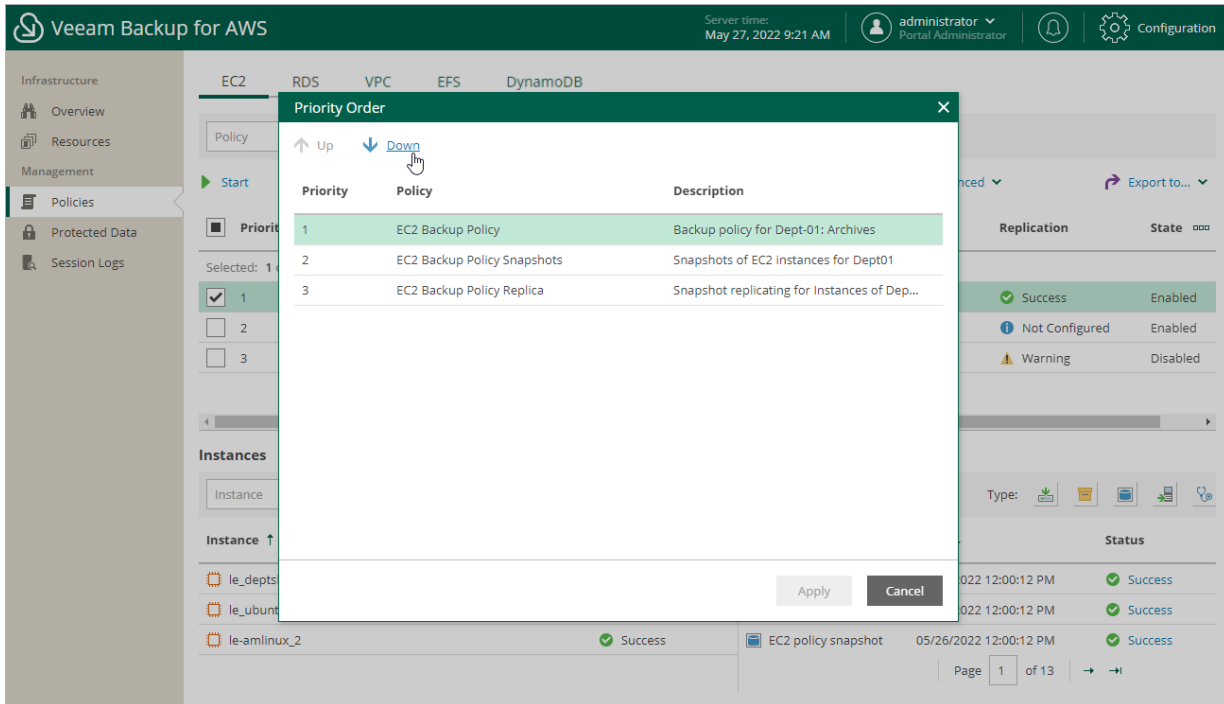
You can set priority for backup policies created in Veeam Backup for AWS. If a resource is included into several backup policies, it will be processed only by one backup policy that has the highest priority.

To set priority for backup policies:

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.



Editing Policy Settings

You can edit backup policies created in Veeam Backup for AWS. For example, you may want to add some resources to a backup policy, change a backup policy description and so on.

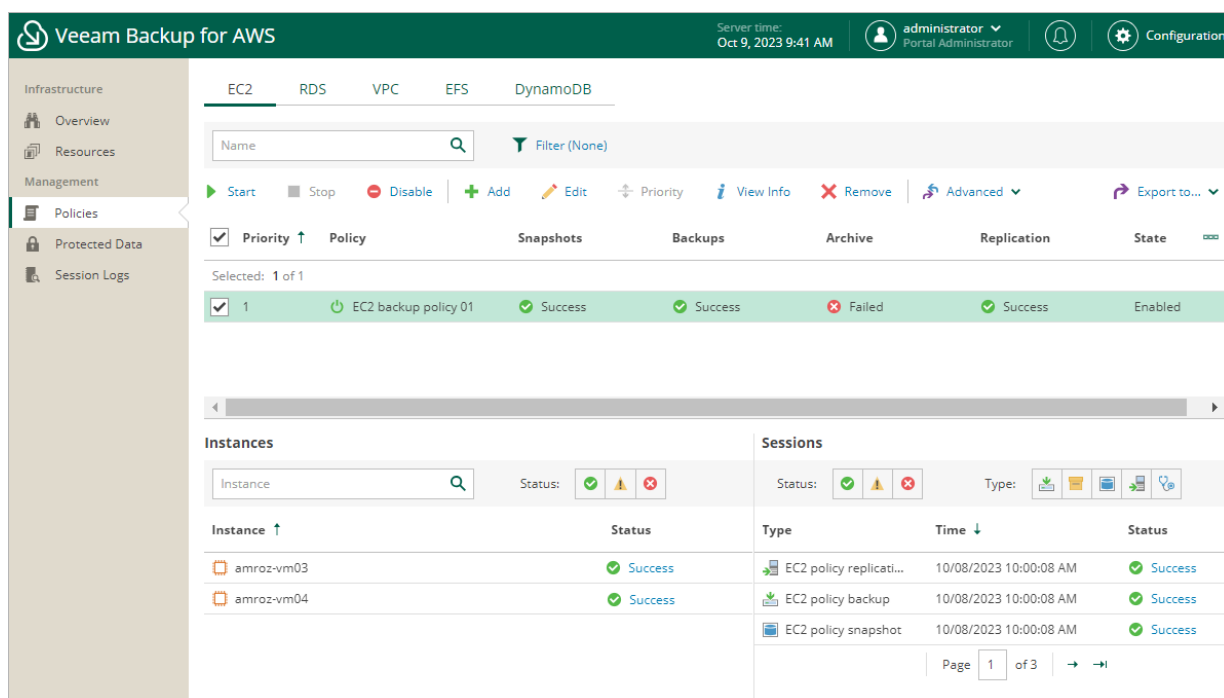
TIP

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](#).

To edit backup policy settings:

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](#) or [Creating EFS Backup Policies](#).



Exporting and Importing Policies

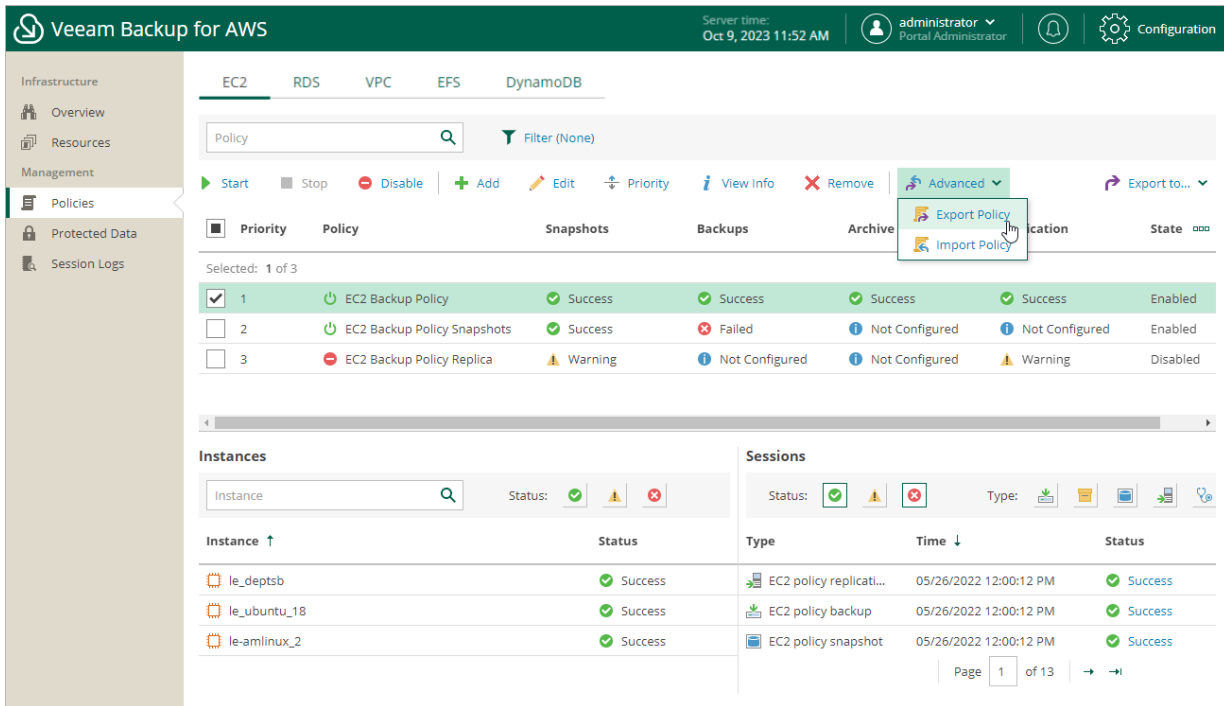
Veeam Backup for AWS allows you to use settings of an existing backup policy as a template for creating other backup policies. You can export a backup policy to a .JSON file, modify the necessary settings in the file, and then import the policy to the same or a different backup appliance.

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**.

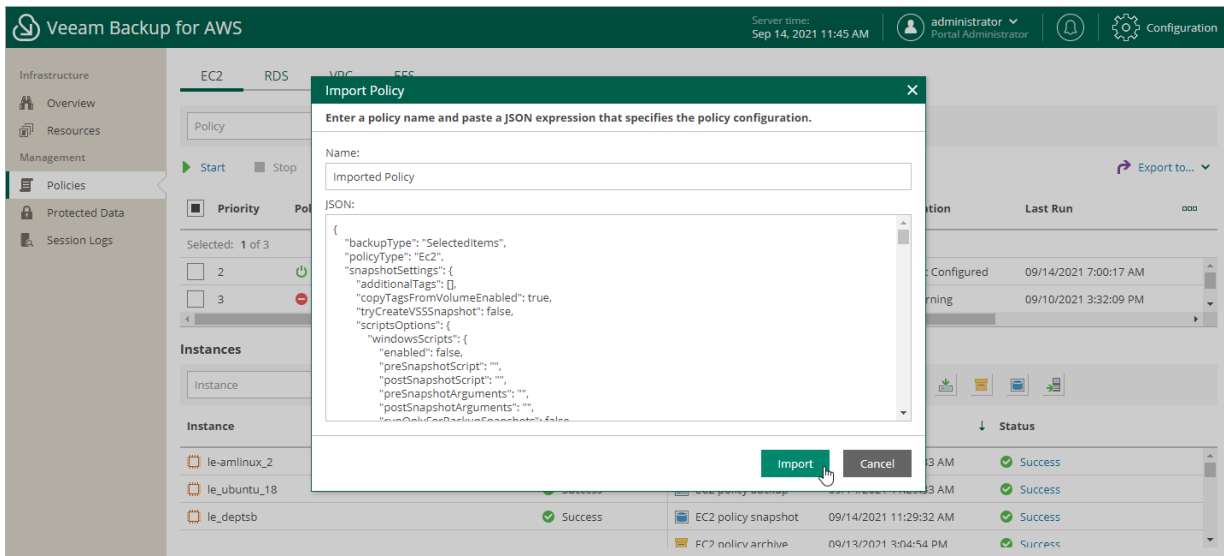
Veeam Backup for AWS will save the backup policy settings as a single .JSON file to the default download directory on the local machine.



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, EFS 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 VPC configuration.
	Indicates that the protected workload is an EFS file system.

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 DynamoDB tables and EFS 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.

NOTE

Veeam Backup & Replication displays all existing snapshots of RDS resources, not only snapshots created by the Veeam backup service. Amazon DB snapshots created for DB instances or Aurora DB clusters in AWS will have the **AWS Snapshot** type in the Veeam Backup & Replication console and the Veeam Backup for AWS Web UI.

- *<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 and EFS file systems, see sections [EC2 Backup](#), [RDS Backup](#), [DynamoDB Backup](#) and [EFS 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](#).

The screenshot shows the Veeam Backup and Replication console interface. The left sidebar displays a navigation tree with 'Home' selected, and 'Backups' expanded under 'External Repository (Encrypted)'. The main pane shows a table of backup objects with columns for Job Name, Creation Time, Restore Points, Repository, and Platform.

Job Name	Creation Time	Restore Points	Repository	Platform
db-snapshot-policy	11/3/2022 2:15 PM		Snapshot	AWS
dept-01-amroz-srv07	10/21/2022 5:36 PM		Snapshot	AWS
EC2-backup-policy01	10/17/2022 12:00 PM		backup-dept05	AWS
EC2-backup-policy01	12/5/2022 12:00 PM		Snapshot	AWS
EC2-backup-policy02	1/19/2023 5:16 PM		backup-dept06	AWS
EC2-backup-policy02	1/19/2023 5:16 PM		backup-dept05	AWS
EC2-backup-policy02	1/19/2023 5:16 PM		Snapshot	AWS
EC2-backup-policy01	10/17/2022 12:00 PM		backup-dept06	AWS
EFS-backup-policy	10/24/2022 11:03 AM		Snapshot	AWS
veeam-tw	1/19/2023 5:48 PM		dept-01-amroz-srv07	AWS

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.

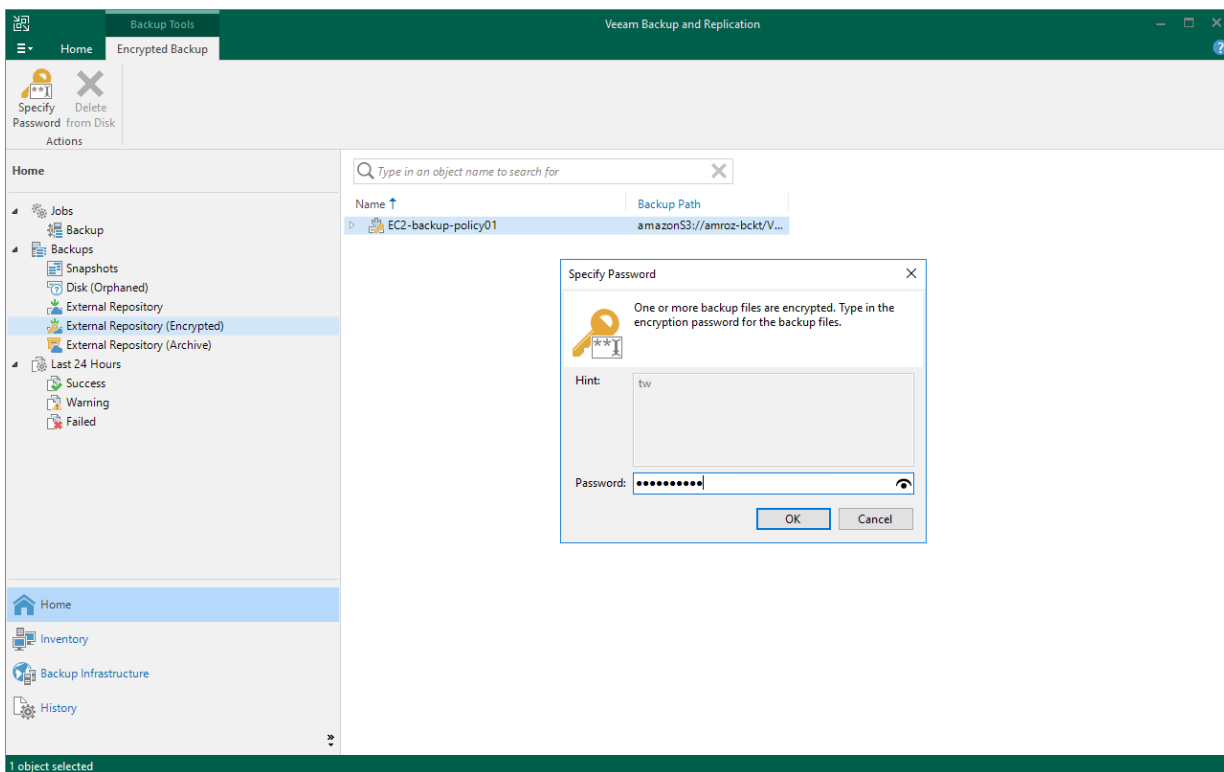
- To decrypt backup files encrypted using a KMS key, make sure that the IAM user specified when [creating a new repository](#) or [adding an existing repository](#) is assigned permissions required to access KMS keys. For more information on the required permissions, see [Plug-in Permissions](#).
- To 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](#) and [Removing EFS Backups Created Manually](#).

EC2 Data

To view and manage backed-up EC2 instance data, navigate to **Protected Data > EC2**. The **EC2** tab displays information on all protected EC2 instances and allows you to remove restore points of the instances if you no longer need them.

For each backed-up EC2 instance, Veeam Backup for AWS creates a record in the configuration database with the following set of properties:

- **Instance** – a name of an EC2 instance.
- **Policy** – a name of the backup policy that processed the EC2 instance.
- **Restore Points** – a number of restore points created for the EC2 instance.
- **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** – an 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** – an operating system running on the EC2 instance.
- **IAM Role** – an IAM Role used to back up the EC2 instance.
- **AWS Account** – an AWS account where the EC2 instance belongs.
- **Instance ID** – an AWS ID of the EC2 instance.

The screenshot shows the Veeam Backup for AWS console interface. The top navigation bar includes the Veeam logo, the text 'Veeam Backup for AWS', and the server time 'Oct 9, 2023 4:18 PM'. The user is logged in as '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 titled 'EC2' and contains a search bar for 'Instance', a filter dropdown set to 'None', and action buttons for 'Restore', 'Remove', and 'Extend Availability'. Below this is a table with columns: Instance, Policy, Restore Points, Latest Restore Point, Backup Size, Archive Size, Region, Operating System, IAM Role, and AWS Account. Three instances are listed: amroz-vm03, amroz-vm04, and amroz-vm05. A dropdown menu is open over the table, showing a list of columns with checkboxes: Instance (checked), Policy (checked), Restore Points (checked), Latest Restore Point (checked), Backup Size (checked), Archive Size (checked), Region (checked), Data Retrieval (unchecked), File-level Recovery URL (unchecked), Operating System (checked), IAM Role (checked), AWS Account (checked), and Instance ID (unchecked).

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

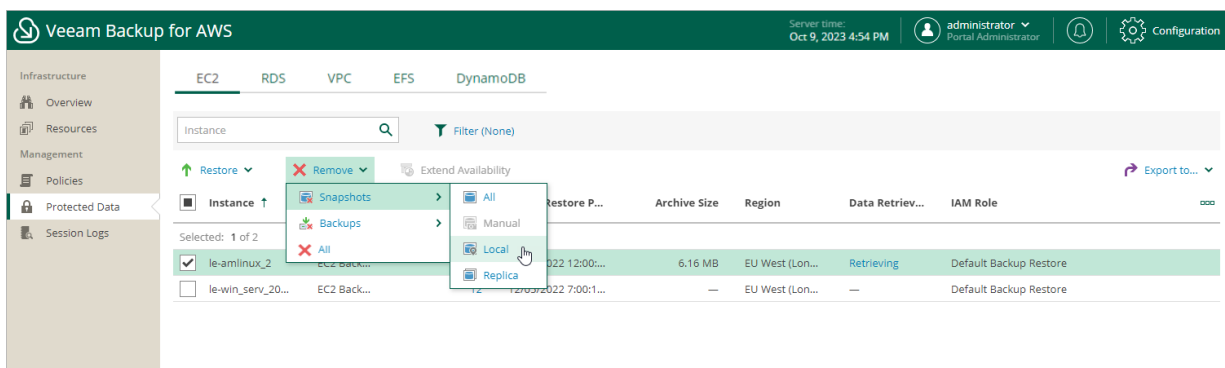
Do not delete backup files from Amazon S3 buckets in the AWS Management Console. If some file in a backup chain is missing, you will not be able to roll back EC2 instance data to the necessary state.

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](#).

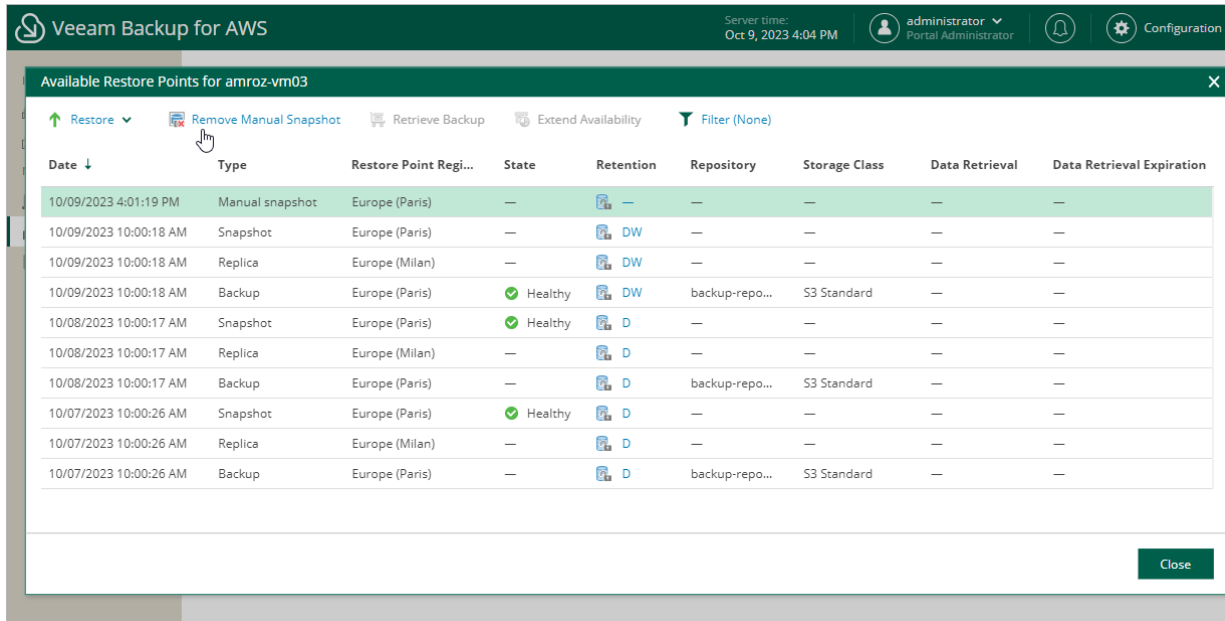
- **Snapshots > Local** – to remove cloud-native snapshots created for the selected EC2 instances by backup policies.
- **Snapshots > Replicas** – to remove snapshot replicas created for the selected EC2 instances by backup policies.
- **Backups > All** – to remove all backups created for the selected 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 details, 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 required permissions 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:

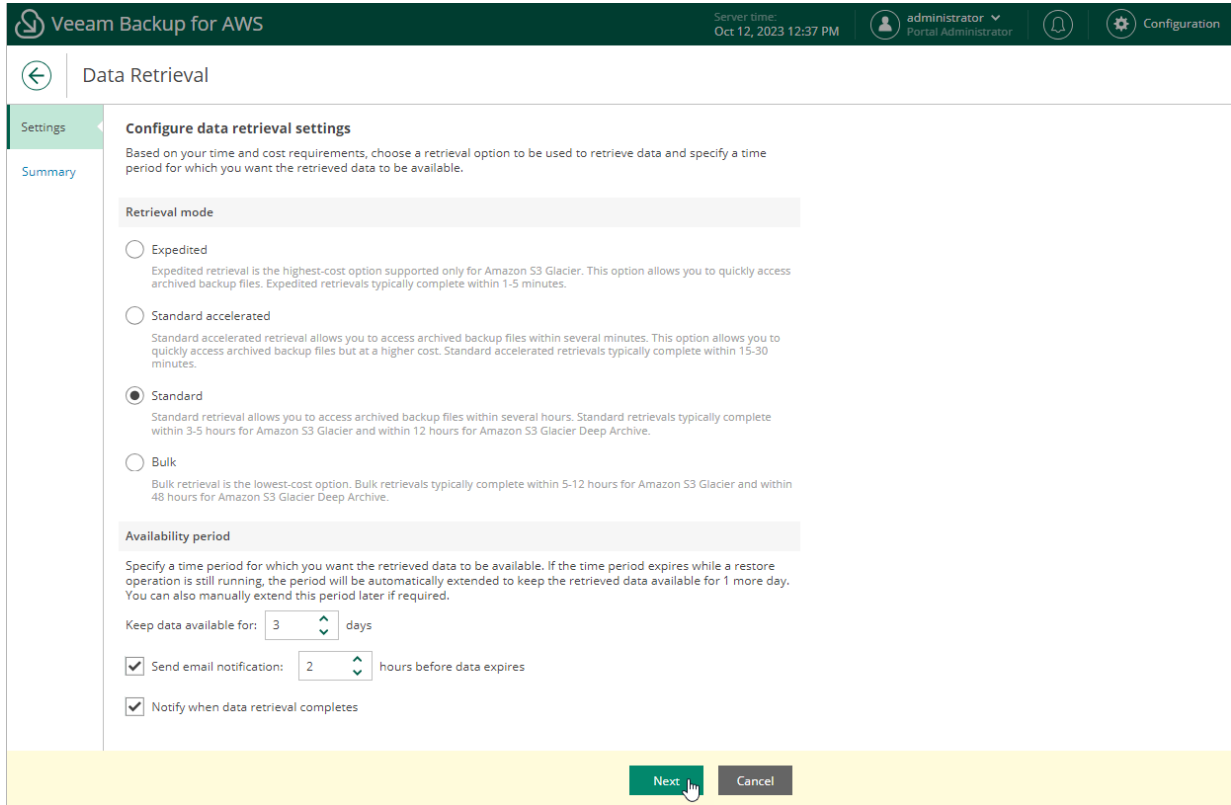
1. Navigate to **Protected Data > EC2**.
2. Select the necessary instance, and click the link in the **Restore Points** column.
3. In the **Available Restore Points** window, select a restore point that contains archived data you want to retrieve, and click **Retrieve Backup**. The **Data Retrieval** wizard will open.

Date ↓	Type	Restore Point Region	State	Retention	Repository	Storage Class	Data Retrieval	Data Retrieval Expirat
08/29/2023 5:13:01 PM	Backup	US East (N. Virginia)	Healthy	D	test-std-virg...	S3 Standard	—	—
08/29/2023 5:13:01 PM	Archive	US East (N. Virginia)	Healthy	D	virginia-arch...	S3 Glacier	—	—
08/29/2023 2:14:29 PM	Backup	US East (N. Virginia)	Healthy	D	test-std-virg...	S3 Standard	—	—
08/29/2023 2:14:29 PM	Archive	US East (N. Virginia)	Healthy	D	virginia-arch...	S3 Glacier	—	—
08/28/2023 6:05:17 PM	Backup	US East (N. Virginia)	Healthy	D	test-std-virg...	S3 Standard	—	—
08/28/2023 6:05:17 PM	Archive	US East (N. Virginia)	Healthy	D	khromeev-d...	S3 Glacier De...	—	—
08/07/2023 4:41:19 PM	Archive	US East (Ohio)	Healthy	D	arch-ohio-1	S3 Glacier	—	—
08/07/2023 2:17:19 PM	Archive	US East (Ohio)	Healthy	D	arch-ohio-1	S3 Glacier	—	—
08/07/2023 12:46:11 PM	Archive	US East (N. Virginia)	Healthy	D	arch-virginia...	S3 Glacier	—	—

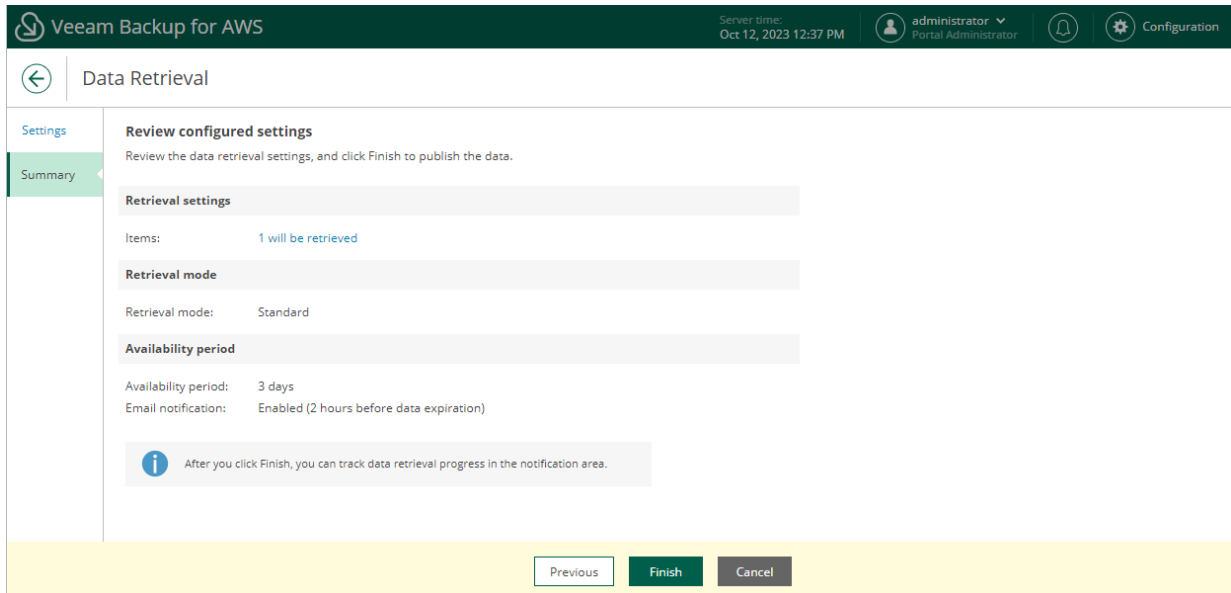
4. At the **Settings** step of the wizard, specify the following settings:
 - a. In the **Retrieval mode** section, select the [retrieval option](#) that Veeam Backup for AWS will use to retrieve the data.

b. In the **Availability period** section, specify the number of days for which you want to keep the data available for restore operations.

If you want to receive an email notification when the 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).



5. At the **Summary** step of the wizard, review configuration information and click **Finish**.



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.

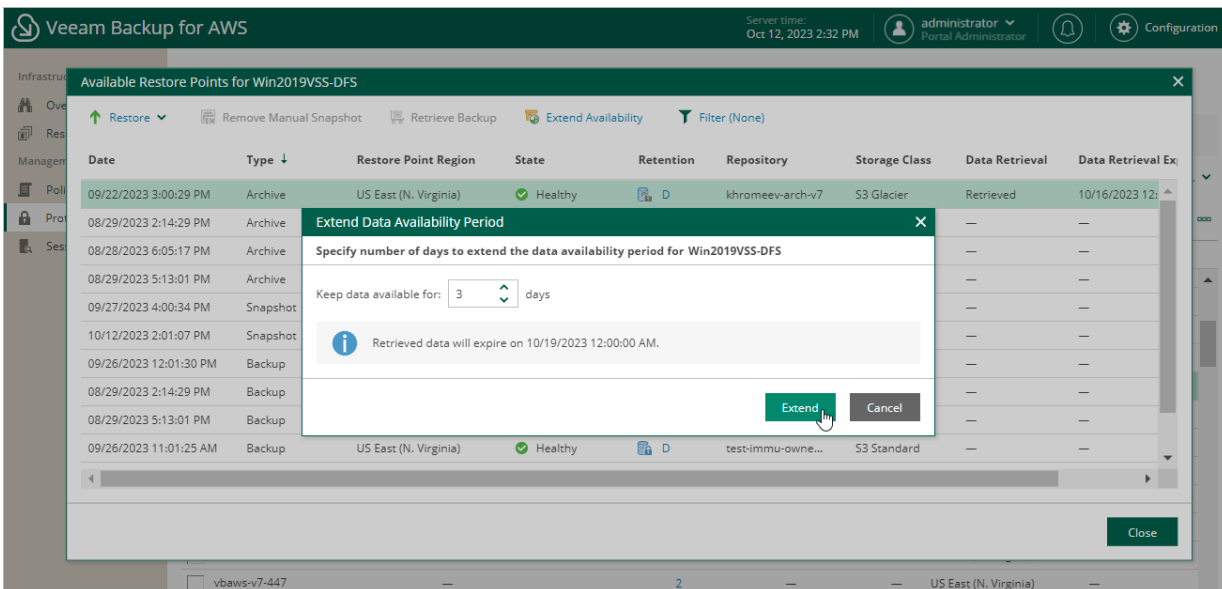
Extending Data Availability

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

To view and manage backed-up RDS data, navigate to **Protected Data > RDS**. The **RDS** tab displays information on all protected DB instances and Aurora DB clusters and allows you to remove restore points of the instances if you no longer need them.

For each backed-up RDS resource, Veeam Backup for AWS creates a record in the configuration database with the following set of properties:

- **Instance** – a name of a DB instance or an Aurora DB cluster.
- **Policy** – a name of the backup policy that processed the DB instance or Aurora DB cluster.
- **Restore Points** – a number of restore points created for the DB instance or Aurora DB cluster.

NOTE

Veeam Backup for AWS displays all existing snapshots of RDS resources, not only snapshots created by the Veeam backup service. Amazon DB snapshots created for DB instances or Aurora DB clusters in AWS have the **AWS Snapshot** type and cannot be deleted from the Veeam Backup for AWS Web UI.

- **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 DB instance or Aurora DB cluster stored in standard.
- **Archive size** – the size of all backups created for the DB instance or Aurora DB cluster stored in archive backup repositories.
- **Engine** – a database engine of the DB instance or Aurora DB cluster.
- **Instance Size** – a size of the DB instance storage.
- **AWS Account** – an AWS account where the DB instance or Aurora DB cluster belongs.
- **Instance ID** – an AWS ID of the DB instance or Aurora DB cluster.
- **Region** – an AWS Region in which the DB instance or Aurora DB cluster resides.

Instance	Policy	Restore Points	Latest Restore Point	Backup Size	Archive Size	Engine	Instance Size	AWS Account	Instance ID	Region
amroz-db-01	—	2	12/20/2022 11:29:52 AM	—	—	PostgreSQL	200 GB	611610175276	db-htugro7...	Europe (Paris)
db01	RDS backup policy 01	96	10/09/2023 3:23:15 PM	16.83 KB	8.26 KB	PostgreSQL	20 GB	61161017527...	db-2c4mz2...	Europe (Paris)
db02	RDS backup policy 02	14	10/09/2023 3:00:27 PM	613.55 MB	23.9 MB	PostgreSQL	20 GB	61161017527...	db-2f6nhc...	Europe (Paris)

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 > 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** – 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** – to remove cloud-native snapshots created for the selected RDS resources by backup policies.
- **Snapshots > Replicas** – 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.

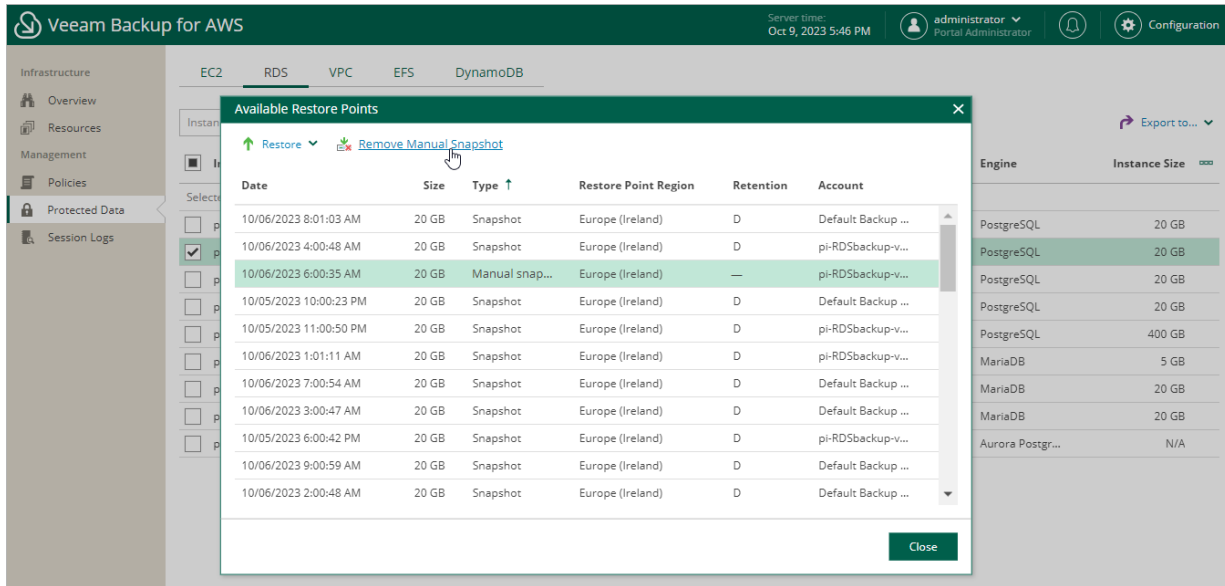
The screenshot shows the Veeam Backup for AWS interface. The top navigation bar includes 'EC2', 'RDS', 'VPC', 'EFS', and 'DynamoDB'. The left sidebar shows 'Protected Data' selected. The main area displays a table of RDS instances. A context menu is open over the table, showing options: Snapshots >, Backups >, All, Standard, and Archived. The 'All' option is highlighted.

Instance	Policy	Restore Point	Backup Size	Archive Size	Engine	Instance Size
pi-psql-empty-retest	RDS-archive	2	16.83 KB	8.26 KB	PostgreSQL	20 GB
pi-postgresql-non-empty	RDS-archive	43 10/06/2023 9:00:59 AM	598.47 MB	47.76 MB	PostgreSQL	20 GB
pi-postgres-ireland	RDS-archive	119 10/06/2023 3:12:41 PM	613.55 MB	23.9 MB	PostgreSQL	20 GB
pi-postgres-empty	RDS-archive	49 10/06/2023 9:01:15 AM	541.54 MB	7.97 MB	PostgreSQL	20 GB
pi-postgres-backup	—	1	—	—	PostgreSQL	400 GB
pi-maria-testsnap	—	6	—	—	MariaDB	5 GB
pi-maria-simple	—	2	—	—	MariaDB	20 GB
pi-maria	RDS-archive	37	—	—	MariaDB	20 GB
pi-aurora	Licensing	1	—	—	Aurora Postgr...	N/A

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 > 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

To view and manage backed-up DynamoDB table data, navigate to **Protected Data > DynamoDB**. The **DynamoDB** tab displays information on all protected DynamoDB tables and allows you to remove restore points of the tables if you no longer need them.

For each backed-up DynamoDB table, Veeam Backup for AWS creates a record in the configuration database with the following set of properties:

- **Name** – a name of a DynamoDB table.
- **Policy** – a name of the backup policy that processed the DynamoDB table.
- **Restore Points** – a number of restore points created for the DynamoDB table.
- **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** – an AWS Region in which the DynamoDB table resides.
- **AWS Account** – an AWS account where the DynamoDB table belong.

Name	Policy	Restore Points	Latest Restore Point	Backup Size	Region	AWS Account
DataTable	DynamoDB backup policy	5	11/21/2023 7:00:14 AM	1.23 KB	Europe (Paris)	61161017527...
DataTable02	—	3	11/21/2023 8:00:13 AM	429 Bytes	Europe (Paris)	61161017527...

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.

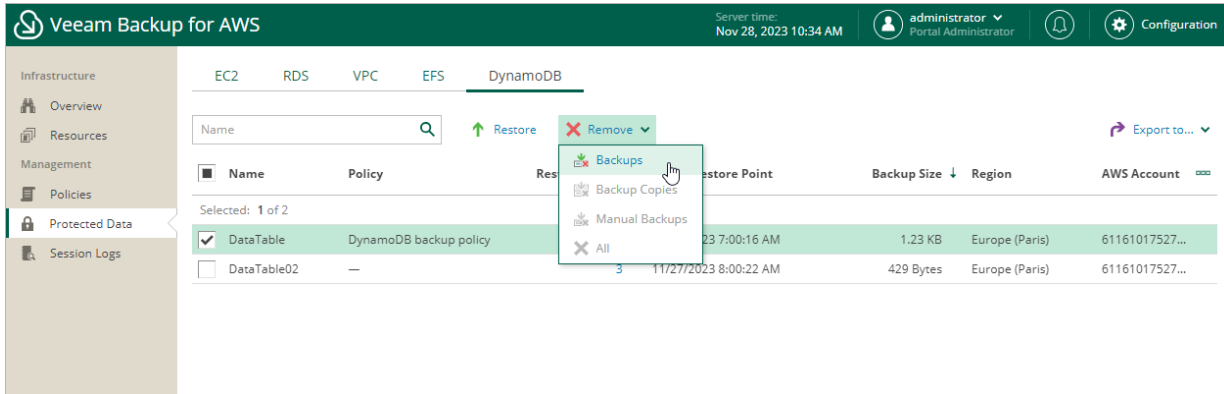
To remove backed-up data manually, do the following:

1. Navigate to **Protected Data > 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](#).

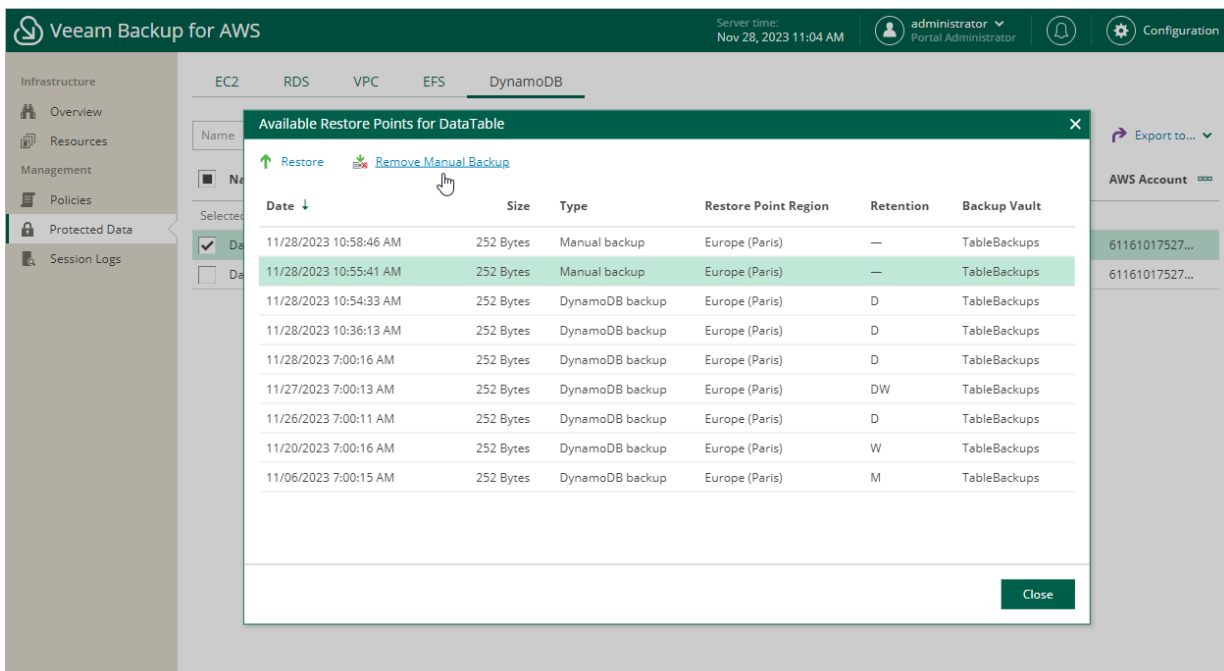
- **All** – to remove all backups and backup copies created for the selected tables both by backup policies and manually.



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 > 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**.



EFS Data

To view and manage backed-up EFS file system data, navigate to **Protected Data > EFS**. The **EFS** tab displays information on all protected EFS file systems and allows you to remove restore points of the file systems if you no longer need them.

For each backed-up Amazon EFS file system, Veeam Backup for AWS creates a record in the configuration database with the following set of properties:

- **Name** – a name of an EFS file system.
- **Policy** – a name of the backup policy that processed the EFS file system.
- **Restore Points** – a number of restore points created for the EFS file system.
- **Latest Restore Point** – the date and time of the latest restore point that was created for the EFS file system.
- **Total Size** – a size of the EFS file system storage.
- **Region** – an AWS Region in which the EFS file system resides.
- **AWS Account** – an AWS account where the EFS file system belong.
- **File System ID** – an 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.

Name	Policy	Restore Points	Latest Restore Point	Total Size	Region	AWS Account	File System ID	File-level Recovery U...
le-base1	—	2	09/14/2021 9:28:14 ...	6 KB	EU West (L...	359000203834 (...)	fs-3cc004cc	—
le-base64	EFS Back...	142	05/30/2022 9:00:29 ...	6 KB	EU West (L...	359000203834 (...)	fs-23c004d3	FLR
le-dept01-share	EFS Back...	157	05/30/2022 9:00:27 ...	6 KB	US East (Vl...	359000203834 (...)	fs-59ead0ec	—
le-mrkt-files	EFS Back...	156	05/30/2022 9:00:27 ...	6 KB	US East (Vl...	359000203834 (...)	fs-9eebd12a	—

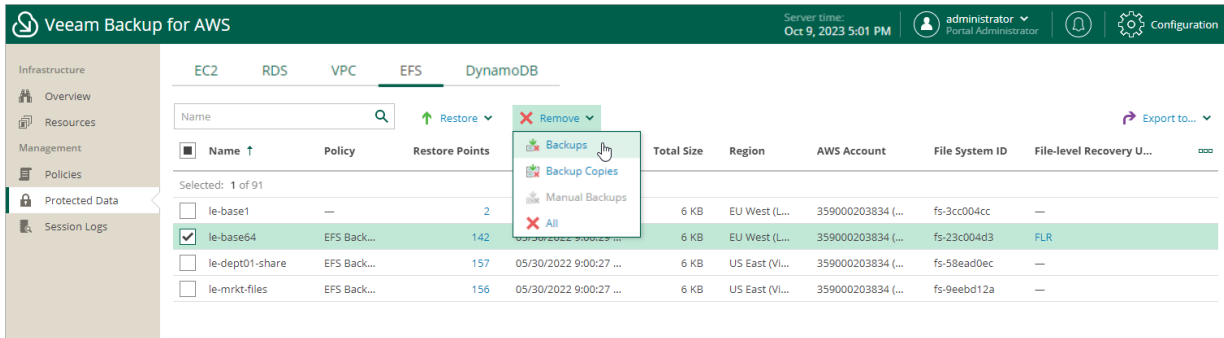
Removing EFS Backups

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 > 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.

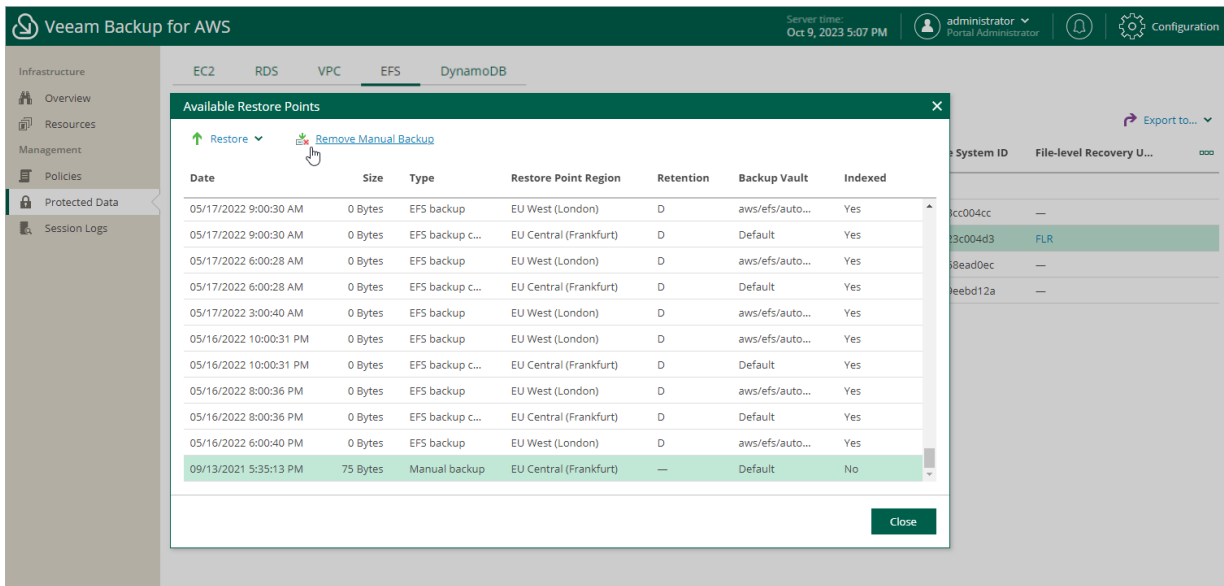
- **Backup Copies** – to remove backup copies created for the selected file systems by backup policies.
 - **Manual Backups** – to remove EFS backups created for the selected file systems manually.
- If you want to remove only specific manual backup, follow the instructions provided in section [Removing 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 > 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**.



VPC Configuration Data

To view and manage backed-up VPC configuration data, navigate to **Protected Data > VPC**. The **VPC** tab displays information on all saved VPC configurations, and allows you to export the configurations and to remove configuration restore points if you no longer need them.

For each protected AWS Region within the AWS account, Veeam Backup for AWS creates a configuration record in the database. To view all existing configuration records, navigate to **Protected Data > VPC**.

Each configuration record is described with a set of properties:

- **AWS Account** – a name of an AWS account whose IAM role was used to collect VPC configuration data.
- **Region** – an 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** – a number of restore points created for the VPC configuration.

In the **Configuration details** section, Veeam Backup for AWS displays the backed-up VPC configuration details for the selected configuration record.

You can [export](#), [compare](#) and [remove](#) backed-up Amazon VPC configuration data.

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 'Oct 9, 2023 5:12 PM', 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 divided into two sections. The top section, titled 'VPC', shows a table with columns for 'AWS Account', 'Region', 'Latest Backup', 'Latest Changes', and 'Restore Points'. A single row is visible with the account '611610175276 (veeam-tw)', region 'Europe (Paris)', latest backup '10/09/2023 5:00:16 PM', latest changes 'No changes detected', and 181 restore points. The bottom section, titled 'Configuration Details', shows a table with columns for 'Name', 'ID', 'Type', 'Modification Date', and 'State'. It lists several VPC resources such as 'ManagedPrefixList', 'SecurityGroup', and 'Endpoint' with their respective IDs and modification dates.

Removing VPC Configuration Backups

Veeam Backup for AWS applies the [configured retention policy settings](#) to automatically remove VPC configuration backups created by the VPC Configuration Backup policy. If necessary, you can also remove the backed-up data manually.

IMPORTANT

If you remove a configuration record for an AWS Region, all VPC configuration backups for the selected AWS Region will be removed.

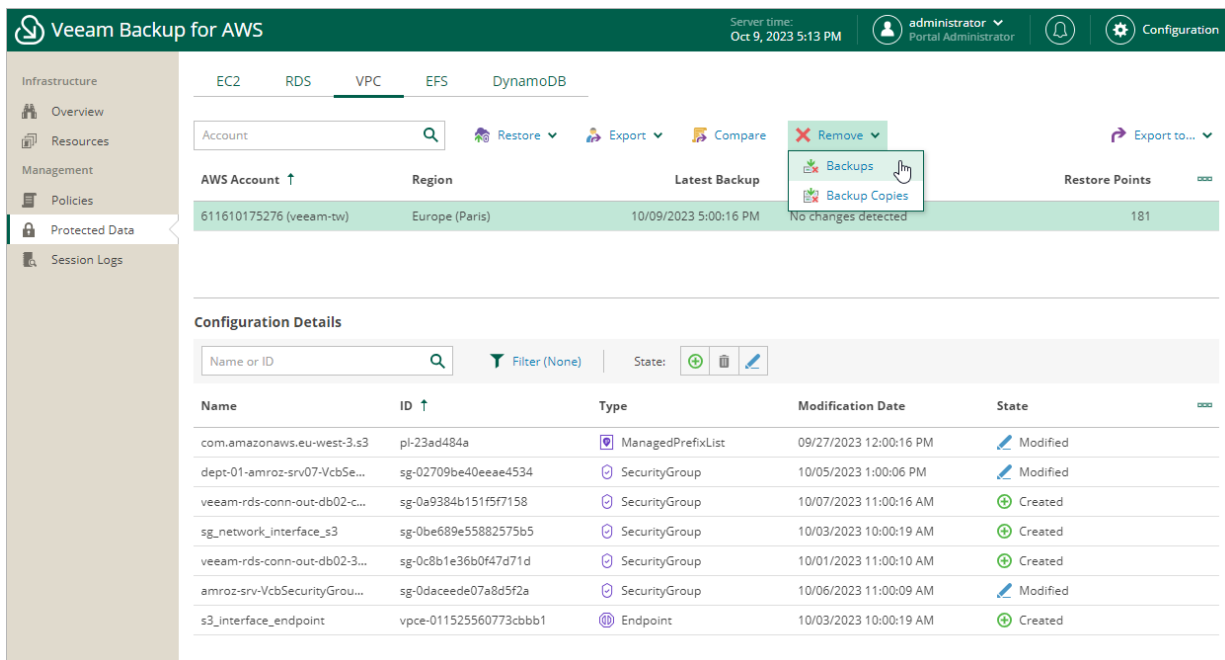
To remove backed-up data manually, do the following:

1. Navigate to **Protected Data > VPC**.
2. Select the configuration record for which you want to remove the backed-up data.
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.

NOTE

If you remove Amazon VPC configuration backups from the Veeam Backup for AWS database but leave their additional copies in a backup repository, you must to re-add the backup repository to Veeam Backup for AWS to be able to view the additional copies in this repository.

- **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.



The screenshot shows the Veeam Backup for AWS interface. The top navigation bar includes 'Infrastructure', 'Overview', 'Resources', 'Management', 'Policies', 'Protected Data', and 'Session Logs'. The main content area is divided into 'EC2', 'RDS', 'VPC', 'EFS', and 'DynamoDB'. The 'VPC' section is active, showing a table of configuration records. A dropdown menu is open over the 'Remove' button, with options for 'Backups' and 'Backup Copies'. Below the table, the 'Configuration Details' section is visible, showing a search bar and a table of configuration records.

Name	ID ↑	Type	Modification Date	State
com.amazonaws.eu-west-3.s3	pl-23ad484a	ManagedPrefixList	09/27/2023 12:00:16 PM	Modified
dept-01-amroz-srv07-VcbSe...	sg-02709be40eeae4534	SecurityGroup	10/05/2023 1:00:06 PM	Modified
veeam-rds-conn-out-db02-c...	sg-0a9384b151f5f7158	SecurityGroup	10/07/2023 11:00:16 AM	Created
sg_network_interface_s3	sg-0be689e55882575b5	SecurityGroup	10/03/2023 10:00:19 AM	Created
veeam-rds-conn-out-db02-3...	sg-0c8b1e36b0f47d71d	SecurityGroup	10/01/2023 11:00:10 AM	Created
amroz-srv-VcbSecurityGrou...	sg-0daceede07a8d5f2a	SecurityGroup	10/06/2023 11:00:09 AM	Modified
s3_interface_endpoint	vpce-011525560773cbbb1	Endpoint	10/03/2023 10:00:19 AM	Created

Comparing VPC Configuration Backups

You can compare the current Amazon VPC configuration of an AWS Region to the backed-up Amazon VPC configuration.

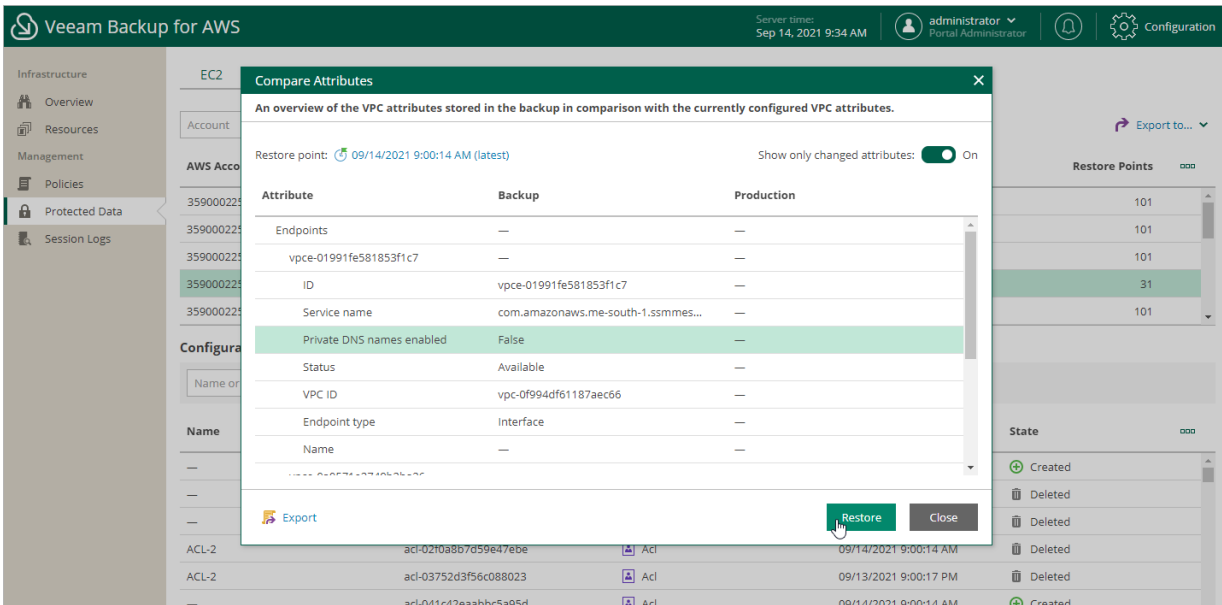
1. Navigate to **Protected Data > VPC**.
2. Select the configuration record for an AWS Region whose VPC configuration you want to compare.
3. Click **Compare**.

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. In the **Compare Attributes** window, click the link to the right of **Restore point** to select the necessary restore point.

If you want Veeam Backup for AWS to display only backed-up VPC configuration items that differs from the current VPC configuration items, set the **Show only changed attributes** toggle to *On*.

You can export or restore the VPC configuration using the selected restore point:

- To export the entire VPC configuration, click **Export** and follow the instructions provided in [Performing Entire Configuration Export](#).
- To restore the entire VPC configuration, click **Restore** and follow the instructions provided in [Performing Entire Configuration Restore](#).



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.

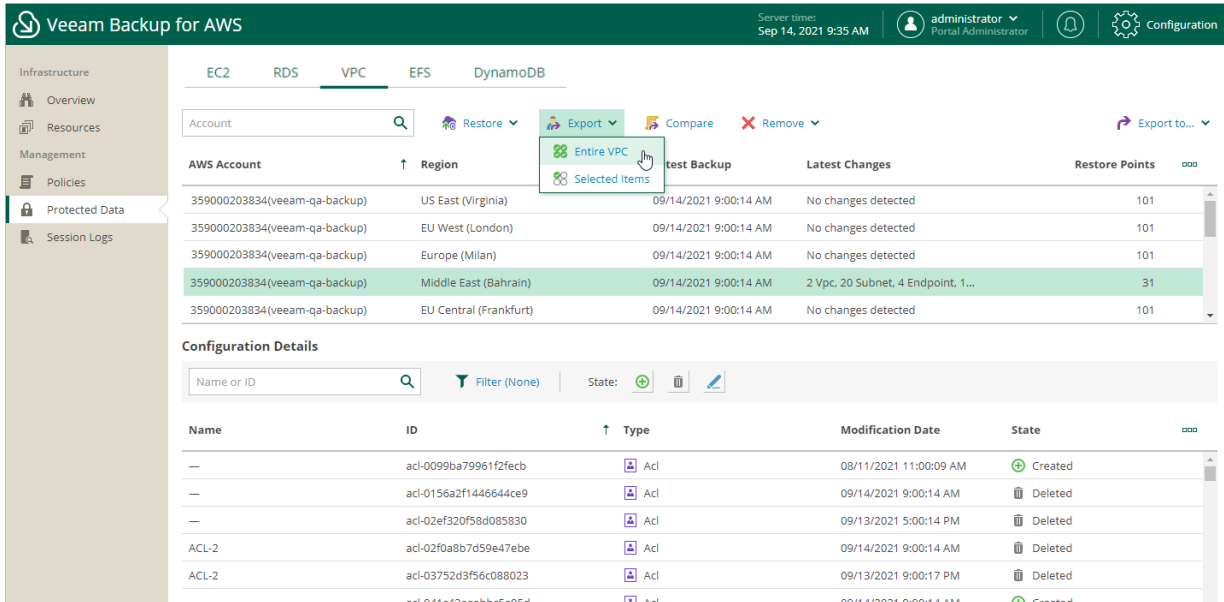
To export the entire VPC configuration to a CloudFormation template, do the following:

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. [Configure 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, the text 'Veeam Backup for AWS', the server time 'Sep 14, 2021 9:35 AM', and the user 'administrator Portal Administrator'. The left sidebar shows the navigation menu with 'Protected Data' selected. The main content area is divided into two sections: a table of VPC configurations and a 'Configuration Details' section.

The VPC configurations table has the following data:

AWS Account	Region	test Backup	Latest Changes	Restore Points
359000203834(veeam-qa-backup)	US East (Virginia)	09/14/2021 9:00:14 AM	No changes detected	101
359000203834(veeam-qa-backup)	EU West (London)	09/14/2021 9:00:14 AM	No changes detected	101
359000203834(veeam-qa-backup)	Europe (Milan)	09/14/2021 9:00:14 AM	No changes detected	101
359000203834(veeam-qa-backup)	Middle East (Bahrain)	09/14/2021 9:00:14 AM	2 Vpc, 20 Subnet, 4 Endpoint, 1...	31
359000203834(veeam-qa-backup)	EU Central (Frankfurt)	09/14/2021 9:00:14 AM	No changes detected	101

The 'Configuration Details' section shows a table of ACLs:

Name	ID	Type	Modification Date	State
—	acl-0099ba79961f2fecb	Acl	08/11/2021 11:00:09 AM	Created
—	acl-0156a2f1446644ce9	Acl	09/14/2021 9:00:14 AM	Deleted
—	acl-02ef320f58d085830	Acl	09/13/2021 5:00:14 PM	Deleted
ACL-2	acl-02f0a8b7d59e47ebe	Acl	09/14/2021 9:00:14 AM	Deleted
ACL-2	acl-03752d3f56c088023	Acl	09/13/2021 9:00:17 PM	Deleted
—	acl-041c42eaabb5a95d	Acl	09/14/2021 9:00:14 AM	Created

The 'Export' dropdown menu is open, showing options for 'Entire VPC' and 'Selected Items'.

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 'Available restore points' window is open, displaying a table of restore points. The most recent restore point is selected.

Date	Changed Objects
09/14/2021 8:00:17 PM	No changes detected
09/14/2021 7:00:11 PM	No changes detected
09/14/2021 6:00:16 PM	No changes detected
09/14/2021 5:00:10 PM	No changes detected
09/14/2021 4:00:10 PM	No changes detected
09/14/2021 3:00:15 PM	No changes detected
09/14/2021 2:00:17 PM	No changes detected
09/14/2021 1:00:17 PM	No changes detected
09/14/2021 12:00:09 PM	No changes detected
09/14/2021 11:00:20 AM	No changes detected
09/14/2021 10:00:09 AM	3 Vpc, 4 Subnet, 3 RouteTable, 6 SecurityGroup, 4 Acl, 2 InternetGateway, 1 Ela...

Step 3. Specify IAM Identity

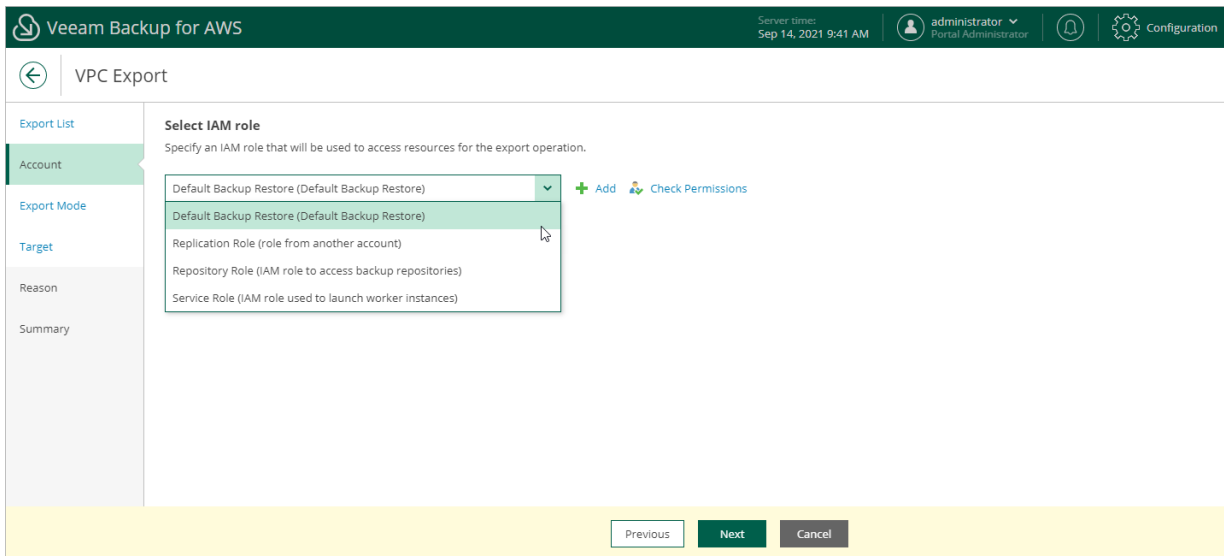
At the **Account** step of the wizard, specify an IAM role whose permissions Veeam Backup for AWS will use to perform the export operation. For more information on permissions required for the IAM role, see [VPC Configuration Restore IAM Permissions](#).

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 add an IAM role, 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 required permissions 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](#).



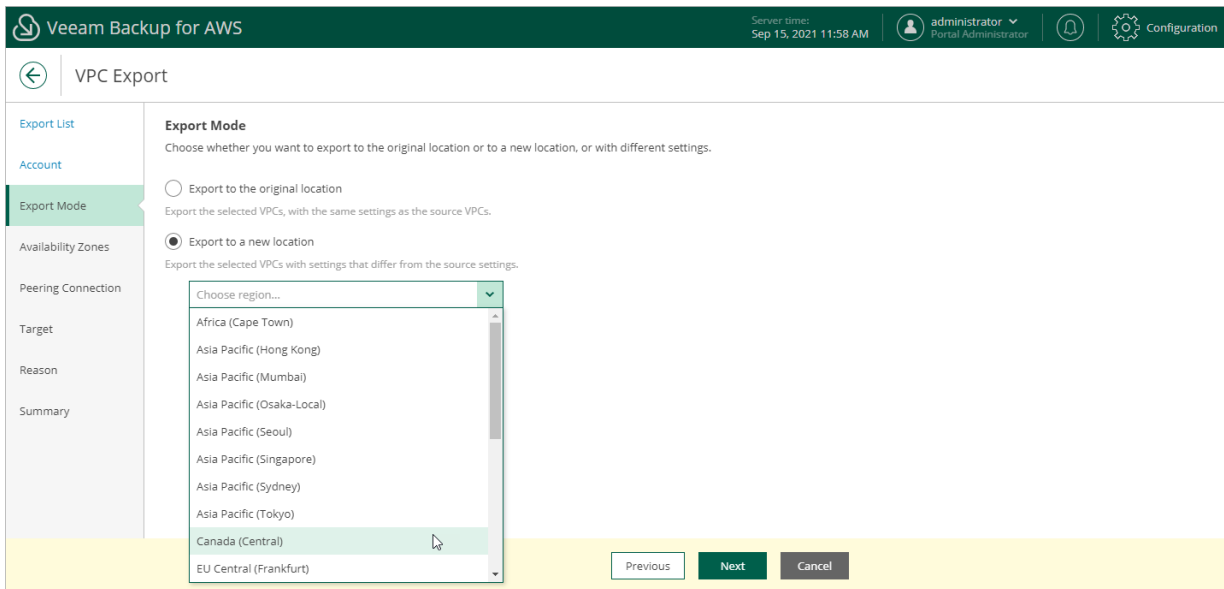
Step 4. Choose Export Mode

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.



Step 5. Configure Availability Zone Mapping

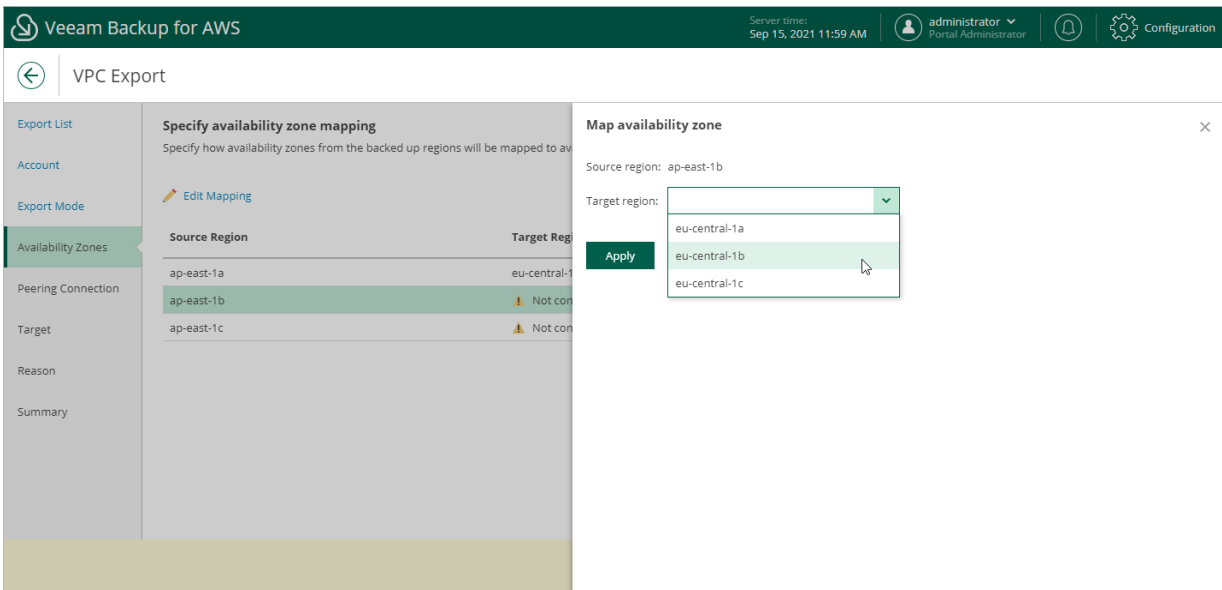
[This step applies only if you have selected the **Export to a new location** option at the **Export Mode** step of the wizard]

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. Configure 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 shows the 'VPC Export' wizard in Veeam Backup for AWS. The current step is 'Review peering connection settings'. The interface includes a sidebar with navigation options: Export List, Account, Export Mode, Availability Zones, Peering Connection (selected), Target, Reason, and Summary. The main content area displays a table with columns for Name, ID, Requested VPC, and Accepted VPC. At the bottom, there are 'Previous', 'Next', and 'Cancel' buttons.

Name	ID	Requested VPC	Accepted VPC
pc_dept01	pcx-0bd5b063924daed33	vpc-00c73e7477f7def0d	vpc-014e41a8eafab1cb9

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 'Target' step is active, showing the 'Specify target location' section. The bucket is set to 'aborbucket'. The 'Create new folder' option is selected with the folder name 'sb_dept01'. A navigation bar at the bottom contains 'Previous', 'Next', and 'Cancel' buttons.

Veeam Backup for AWS Server times: Sep 15, 2021 12:03 PM administrator Portal Administrator Configuration

VPC Export

Export List

Account

Export Mode

Availability Zones

Peering Connection

Target

Reason

Summary

Specify target location
To perform the export operation, specify an S3 bucket where the created CloudFormation template will be stored.

Bucket: `aborbucket`

Use root folder

Create new folder:

i For more information on how to import CloudFormation templates using S3, see [AWS Documentation](#).

Previous Next Cancel

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 Veeam Backup for AWS interface. The top navigation bar includes the Veeam logo, the text "Veeam Backup for AWS", the server time "Sep 15, 2021 12:12 PM", the user "administrator Portal Administrator", and a "Configuration" link. The main header shows a back arrow and "VPC Export". On the left, a sidebar lists steps: Export List, Account, Export Mode, Availability Zones, Peering Connection, Target, Reason (highlighted), and Summary. The main content 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 vpc-0edb0680332a98262". 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, server time (Sep 15, 2021 12:13 PM), user (administrator), and configuration icons. 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 the following information:

Export destination	
Export destination:	As a new VPC
Location name:	EU Central (Frankfurt)

IAM role	
IAM role name:	Default Backup Restore (Default Backup Restore)

Reason	
Reason:	Export of VPC configuration vpc-0edeb0680332a98262

At the bottom of the wizard, there are three buttons: 'Previous' (disabled), 'Finish' (active), and 'Cancel' (disabled).

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 (Sep 14, 2021 9:36 AM), and user information (administrator, Portal Administrator). The left sidebar shows the navigation menu with 'Protected Data' selected. The main content area is divided into two sections: a table of VPC configurations and a 'Configuration Details' section.

The VPC configurations table has the following data:

AWS Account	Region	test Backup	Latest Changes	Restore Points
359000203834(veeam-qa-backup)	US East (Virginia)	09/14/2021 9:00:14 AM	No changes detected	101
359000203834(veeam-qa-backup)	EU West (London)	09/14/2021 9:00:14 AM	No changes detected	101
359000203834(veeam-qa-backup)	Europe (Milan)	09/14/2021 9:00:14 AM	No changes detected	101
359000203834(veeam-qa-backup)	Middle East (Bahrain)	09/14/2021 9:00:14 AM	2 Vpc, 20 Subnet, 4 Endpoint, 1...	31
359000203834(veeam-qa-backup)	EU Central (Frankfurt)	09/14/2021 9:00:14 AM	No changes detected	101

The 'Configuration Details' section shows a table of ACLs:

Name	ID	Type	Modification Date	State
—	acl-0099ba79961f2fecb	Acl	08/11/2021 11:00:09 AM	Created
—	acl-0156a2f1446644ce9	Acl	09/14/2021 9:00:14 AM	Deleted
—	acl-02ef320f58d085830	Acl	09/13/2021 5:00:14 PM	Deleted
ACL-2	acl-02f0a8b7d59e47ebe	Acl	09/14/2021 9:00:14 AM	Deleted
ACL-2	acl-03752d3f56c088023	Acl	09/13/2021 9:00:17 PM	Deleted
—	acl-041c42eaabb5a95d	Acl	09/14/2021 9:00:14 AM	Created

The 'Export' menu is open, showing options for 'Entire VPC' and 'Selected Items'. The 'Selected Items' option is highlighted by the mouse cursor.

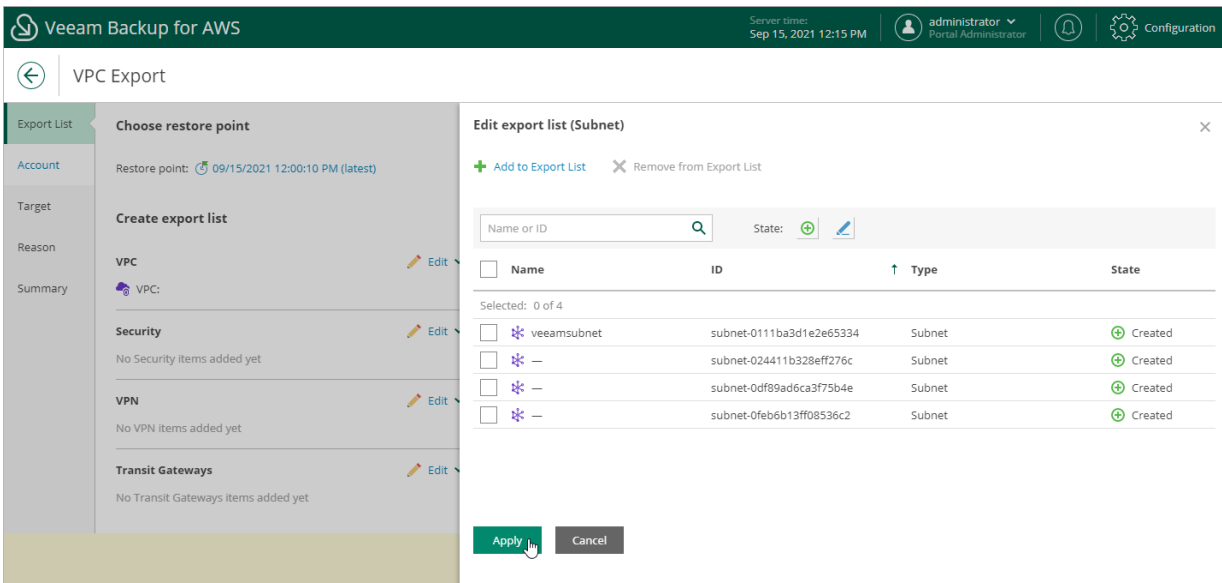
Step 2. Select Restore Point

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 main window is titled 'VPC Export' and has a sidebar with 'Export List', 'Account', 'Target', 'Reason', and 'Summary'. The 'Export List' section is active, showing 'Choose restore point' (Restore point: 09/15/2021 12:00:10 PM (latest)) and 'Create export list' with categories: VPC, Security, VPN, and Transit Gateways. An 'Edit export list (Subnet)' modal window is open, displaying a search bar and a table of subnets. The table has columns: Name, ID, Type, and State. The table content is as follows:

<input type="checkbox"/>	Name	ID	Type	State
<input type="checkbox"/>	veeamsubnet	subnet-0111ba3d1e2e65334	Subnet	Created
<input type="checkbox"/>	—	subnet-024411b328eff276c	Subnet	Created
<input type="checkbox"/>	—	subnet-0df89ad6ca3f75b4e	Subnet	Created
<input type="checkbox"/>	—	subnet-0feb6b13ff08536c2	Subnet	Created

At the bottom of the modal, there are 'Apply' and 'Cancel' buttons. The 'Apply' button is highlighted with a mouse cursor.

Step 3. Specify IAM Identity

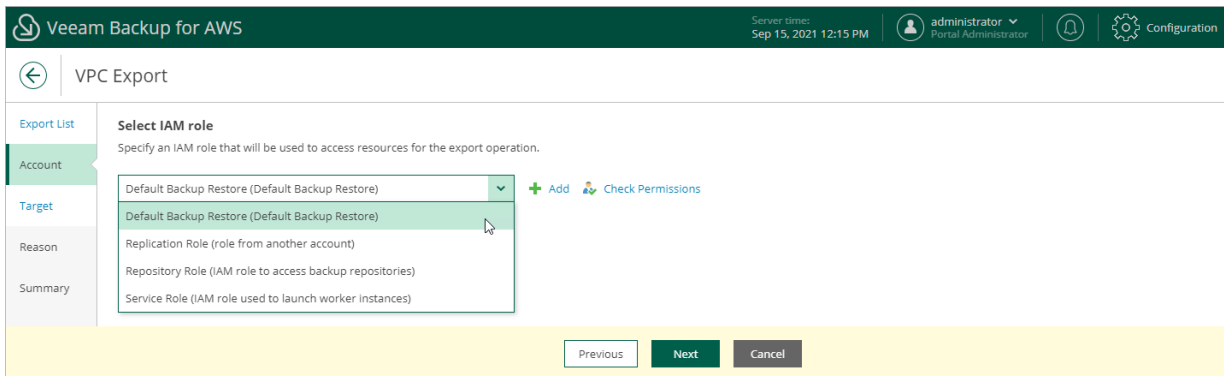
At the **Account** step of the wizard, specify an IAM role whose permissions Veeam Backup for AWS will use to perform the export operation. For more information on permissions required for the IAM role, see [VPC Configuration Restore IAM Permissions](#).

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 add an IAM role, 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 required permissions 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 an Amazon S3 bucket. The bucket name 'abor-stock' is entered. The user has selected the option to 'Create new folder' with the folder name 'sb_dept01'. A 'Next' button is highlighted in green, indicating the user can proceed to the next step.

Veeam Backup for AWS

Server time: Sep 15, 2021 12:17 PM

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: [abor-stock](#)

Use root folder

Create new folder:

i 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 'VPC Export' wizard in the Veeam Backup for AWS console. The interface is divided into a top navigation bar, a left sidebar, and a main content area. The top bar includes the Veeam logo, the text 'Veeam Backup for AWS', the server time 'Sep 15, 2021 12:18 PM', the user 'administrator Portal Administrator', and a 'Configuration' link. The left sidebar contains a list of steps: 'Export List', 'Account', 'Target', 'Reason' (which is highlighted in green), and 'Summary'. The main content area is titled 'VPC Export' and contains the 'Export reason' section. This section has a sub-header 'Export reason' and a description 'Specify a reason for performing the export operation.' Below this is a text input field with the label 'Export reason:' and the text 'Exporting subnet configurations'. At the bottom of the main content area, there are three buttons: 'Previous', 'Next' (which is highlighted in green), and 'Cancel'.

Step 6. 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. The top navigation bar includes the Veeam logo, the text "Veeam Backup for AWS", the server time "Sep 21, 2021 10:55 AM", the user "administrator Portal Administrator", and a "Configuration" link. The main content area is titled "VPC Export" and shows a sidebar with steps: "Export List", "Account", "Target", "Reason", and "Summary" (which is highlighted in green). The main panel displays "Review configured settings" with the following details:

- Export destination:** Original location
- IAM role name:** Default Backup Restore (Default Backup Restore)
- Reason:** Exporting subnet configurations

At the bottom of the wizard, there are three buttons: "Previous", "Finish" (highlighted in green), and "Cancel".

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 EFS 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 a disaster strikes, 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

Consider the following:

- 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](#).
- Deployment of worker instances used for restore operations in [production accounts](#) is available only in the Veeam Backup for AWS Web UI. If you plan to use this functionality, open the Veeam Backup for AWS [appliance Web UI](#) and follow the instructions provided in section [EC2 Restore Using Web UI](#).

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**.

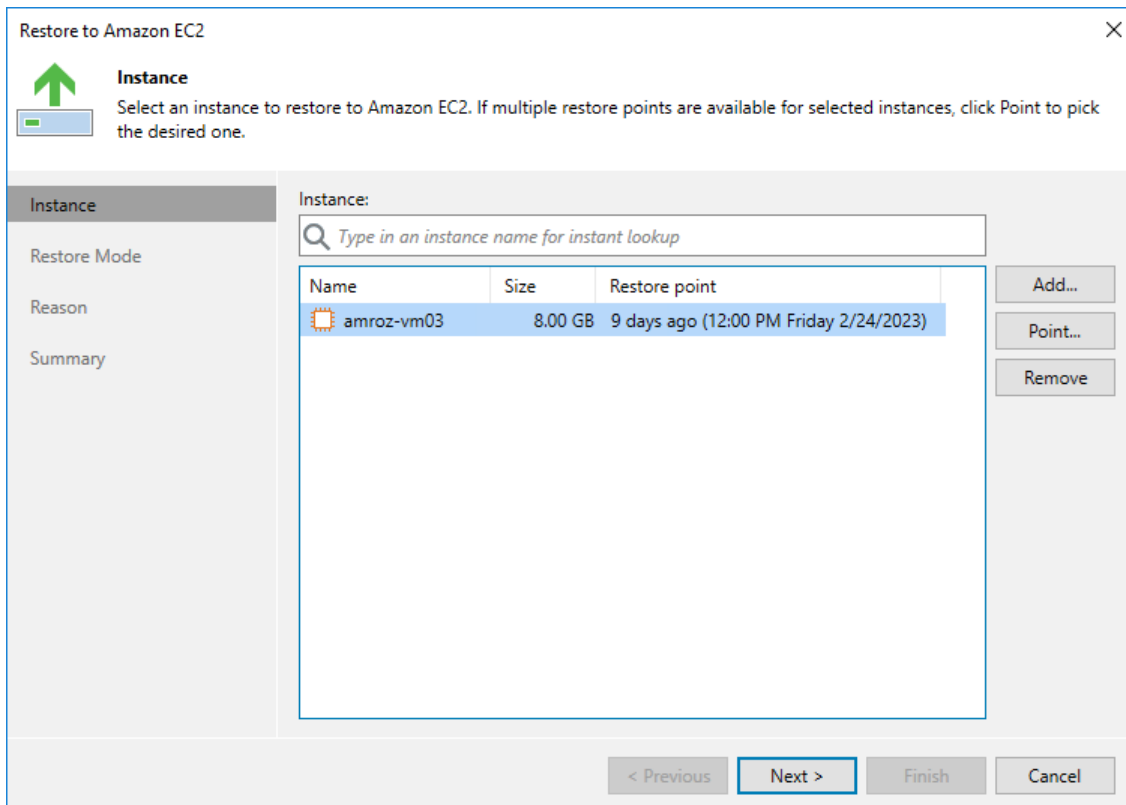
To help you choose a restore point, Veeam Backup & Replication provides the following information on each available restore point:

- **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 accelerated** – the least expensive mode. If you choose this mode, the retrieved data will be available within 25 minutes for data stored in the S3 Glacier Flexible Retrieval storage class and within 8 hours for data stored in the S3 Glacier Deep Archive storage class. With this mode 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 required permissions to perform data retrieval operations. For details, see [Managing Backup Repositories](#).

- **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.

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.

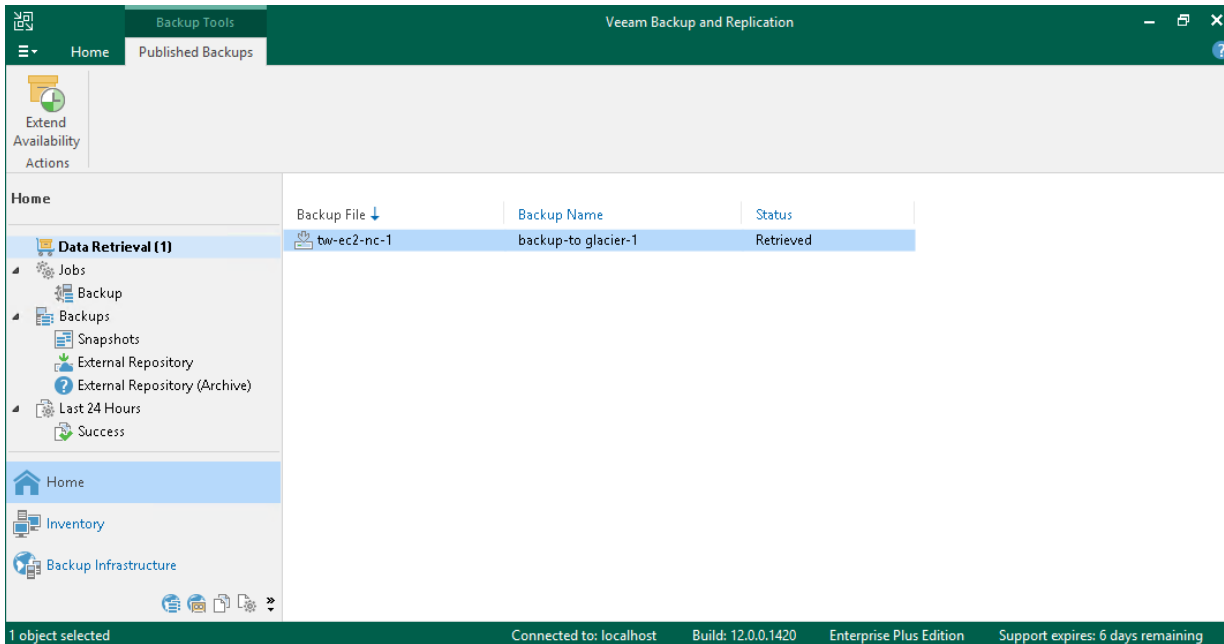
TIP

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.



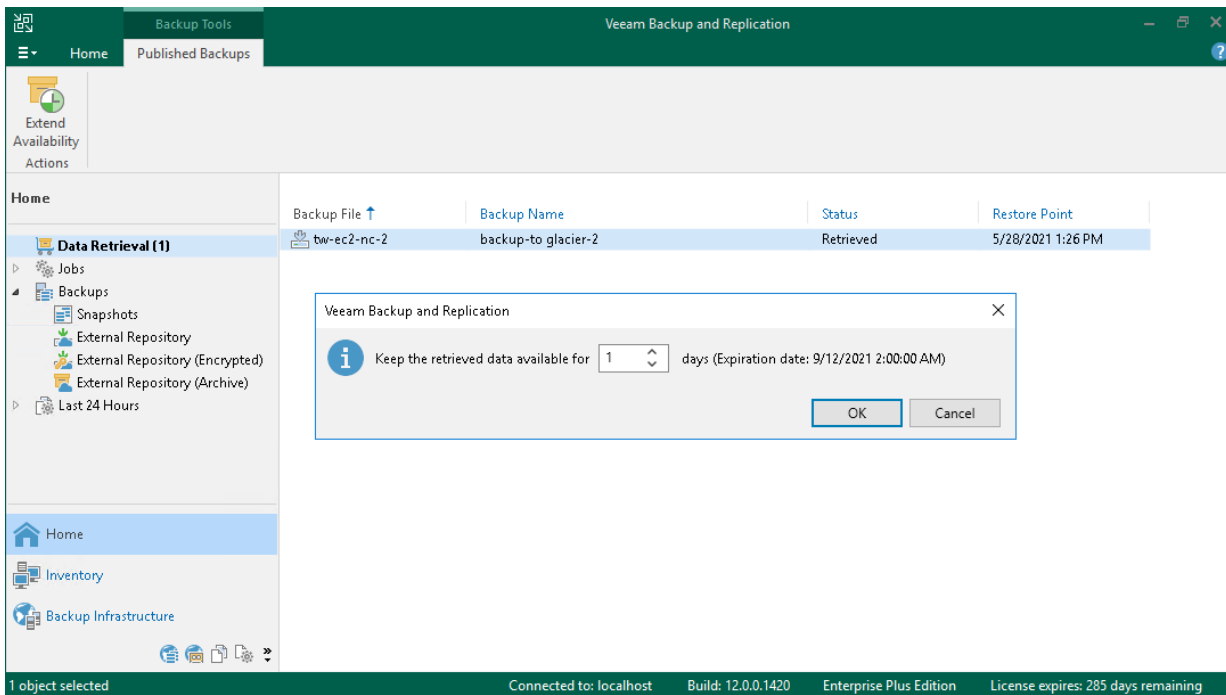
Extending Data Availability

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 account 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 one-time 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 restore operations, Veeam Backup & Replication uses permissions of either the *Default Backup Restore* IAM role, or 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 *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](#).

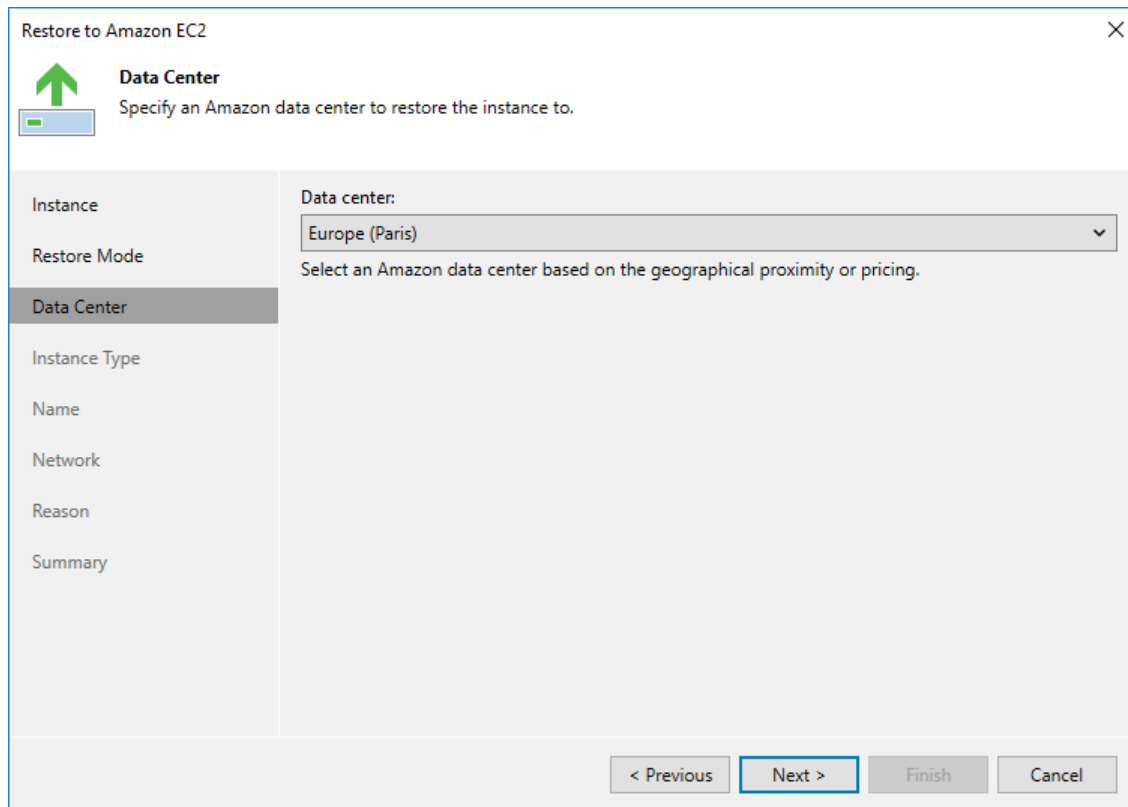
The screenshot shows a wizard window titled "Restore to Amazon EC2" with a close button (X) in the top right corner. The main content area is titled "Restore Mode" and includes a green upward-pointing arrow icon. Below the title, there is a description: "Specify whether selected instances should be restored back to the original location, or to a new location or with different settings." On the left side, there is a vertical navigation pane with the following items: "Instance", "Restore Mode" (which is highlighted), "Data Center", "Instance Type", "Name", "Network", "Reason", and "Summary". The main content area contains two radio button options:
1. "Restore to the original location" (unselected): "Quickly initiate restore of the selected instance to its original location, with the original name and settings. This option minimizes the chance of user input error."
2. "Restore to a new location, or with different settings" (selected): "Customize the restored instance location, and change its settings. The wizard will automatically populate all controls with the original instance settings as the defaults." Below this second option is a blue link that says "Pick account to use". At the bottom of the window, there are four buttons: "< Previous", "Next >" (which is highlighted with a blue border), "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 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' followed by 'Specify an Amazon data center to restore the instance to.' Below this is a dropdown menu labeled 'Data center:' with 'Europe (Paris)' selected. Underneath the dropdown is the instruction 'Select an Amazon data center based on the geographical proximity or pricing.' On the left side of the window is a vertical navigation pane with the following items: 'Instance', 'Restore Mode', 'Data Center' (which is highlighted with a dark background), 'Instance Type', 'Name', 'Network', 'Reason', and 'Summary'. At the bottom of the window 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.

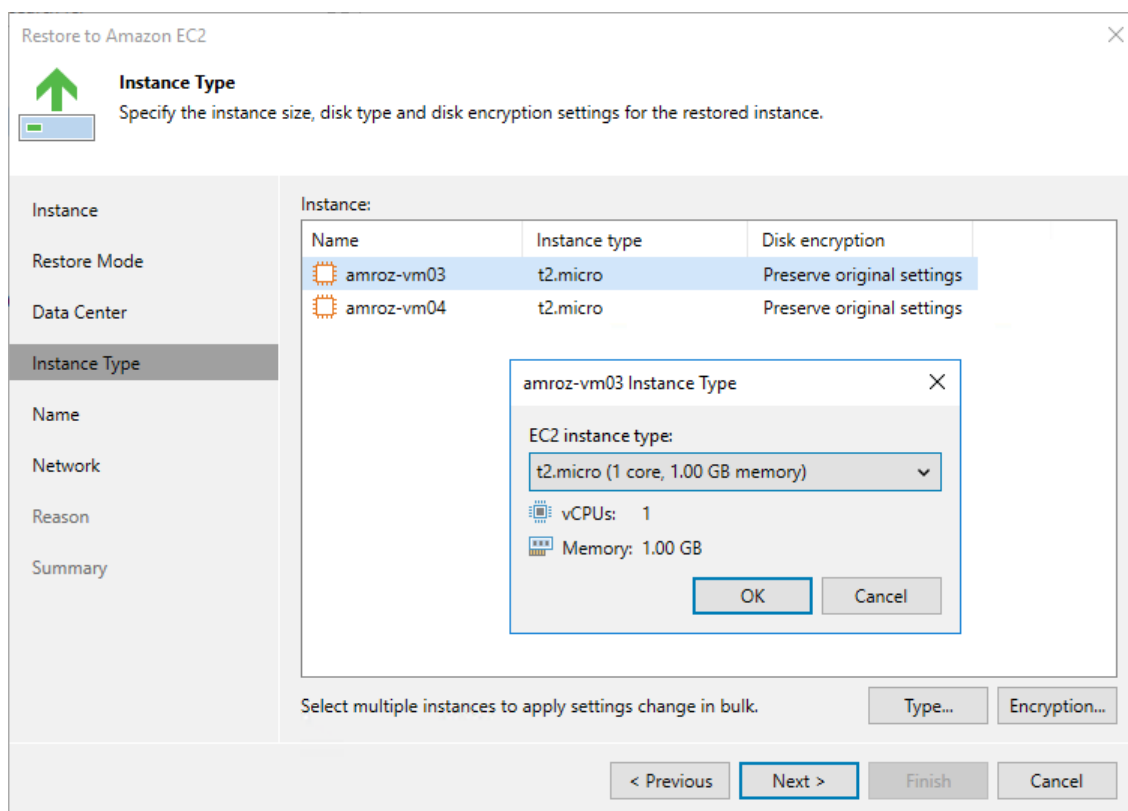
- Select the **Use the following encryption key** option if you want to encrypt the restored EBS volumes of the processed EC2 instance with an AWS KMS key. Then, choose the necessary KMS key from the list.

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.



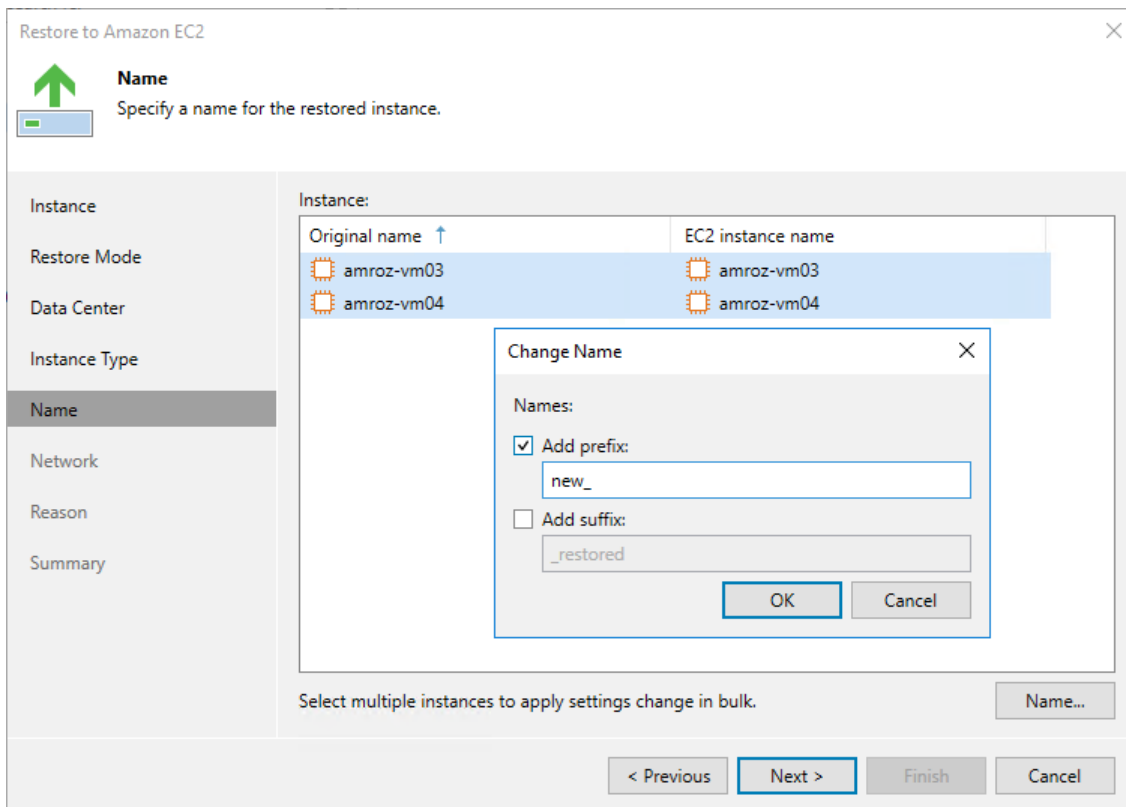
Step 6. Specify Instance Name

[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**.

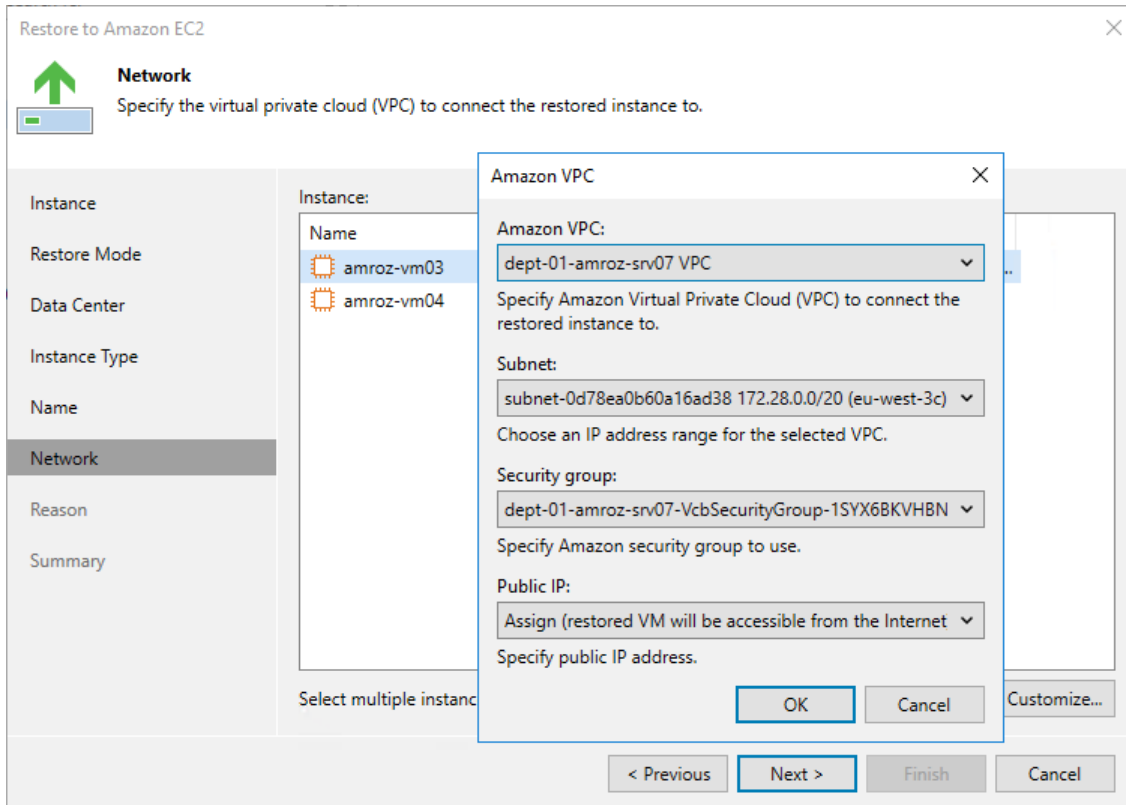


Step 7. Configure Network Settings

[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. The window is divided into a left sidebar and a main content area. The sidebar contains a list of steps: Instance, Restore Mode, Data Center, Instance Type, Name, Network, Reason (highlighted), and Summary. Above the sidebar, there is a green upward-pointing arrow icon and the heading "Reason". Below the heading, 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 content area features a large text input field with the label "Restore reason:" and the text "Restore failed EC2 instances" entered. 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 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.

The screenshot shows the 'Restore to Amazon EC2' wizard at the 'Summary' step. The window title is 'Restore to Amazon EC2'. On the left, there is a navigation pane with options: Instance, Restore Mode, Data Center, Instance Type, Name, Network, Reason, and Summary (which is selected). The main area displays the following summary information:

Summary:
IAM role: Policy role
Data center: Europe (Paris)

Items:

- Original instance name: amroz-vm03
EC2 instance name: amroz-vm03
Restore point: 2/23/2023 12:00:23 PM
EC2 instance type: t2.micro
VPC: dept-01-amroz-srv07 VPC (172.28.0.0/16)
Subnet: subnet-0d78ea0b60a16ad38 172.28.0.0/20 (eu-west-3c)
Security group: dept-01-amroz-srv07-VcbSecurityGroup-1SYX6BKVHBNX
Do not assign public IP address: False
KMS key: Preserve original settings
- Original instance name: amroz-vm04
EC2 instance name: amroz-vm04
Restore point: 2/22/2023 8:00:24 AM
EC2 instance type: t2.micro
VPC: dept-01-amroz-srv07 VPC (172.28.0.0/16)
Subnet: subnet-0d78ea0b60a16ad38 172.28.0.0/20 (eu-west-3c)

At the bottom of the summary area, there is a checkbox labeled 'Power on target instance after restoring' which is currently unchecked. At the bottom of the wizard, there are four buttons: '< Previous', 'Next >', 'Finish' (highlighted in blue), and '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 one-time 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 [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](#).

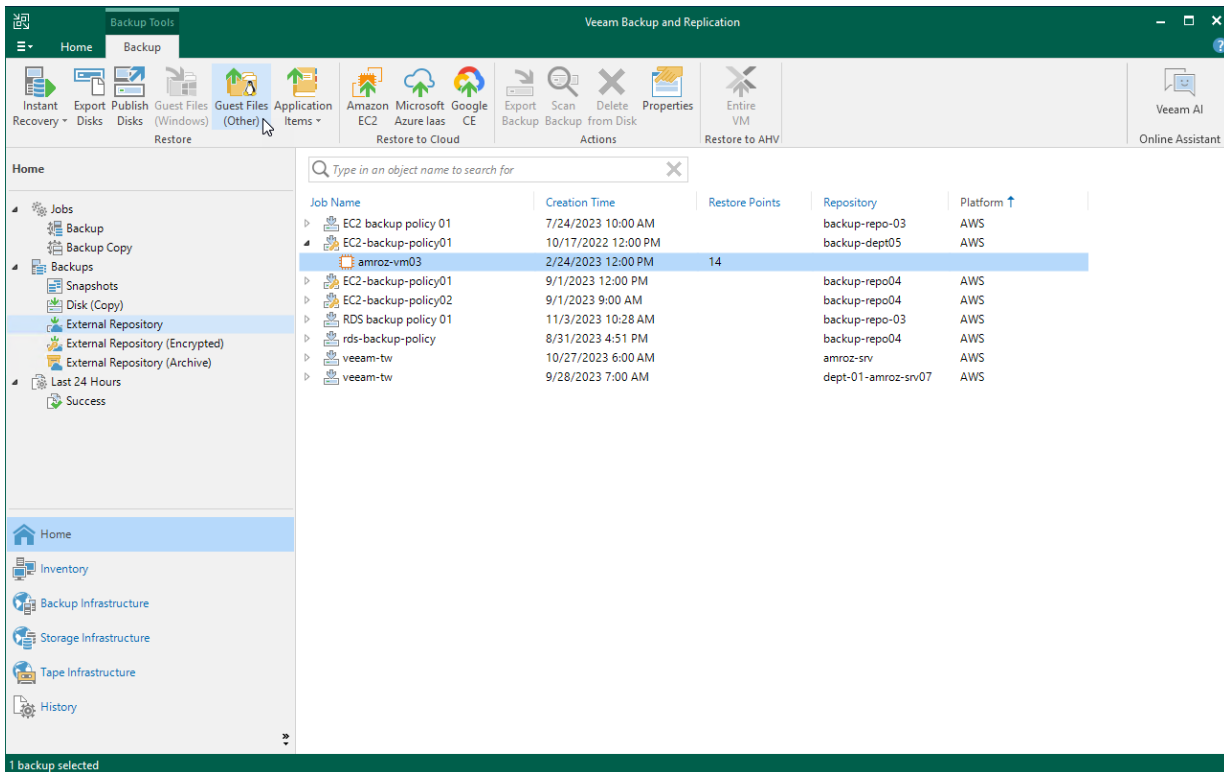
To restore guest OS files and folders, do the following:

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 one-time 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](#).


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. [Launch the Instance Restore wizard](#).
2. [Select a restore point](#).
3. [Specify data retrieval settings for archived backups](#).
4. [Specify restore settings](#).
5. [Choose a restore mode](#).
6. [Enable encryption for EBS volumes](#).
7. [Specify EC2 instance settings](#).
8. [Configure network settings](#).
9. [Specify a restore reason](#).
10. [Finish working with the wizard](#).

Before You Begin

Before you restore EC2 instances, consider the following limitations:

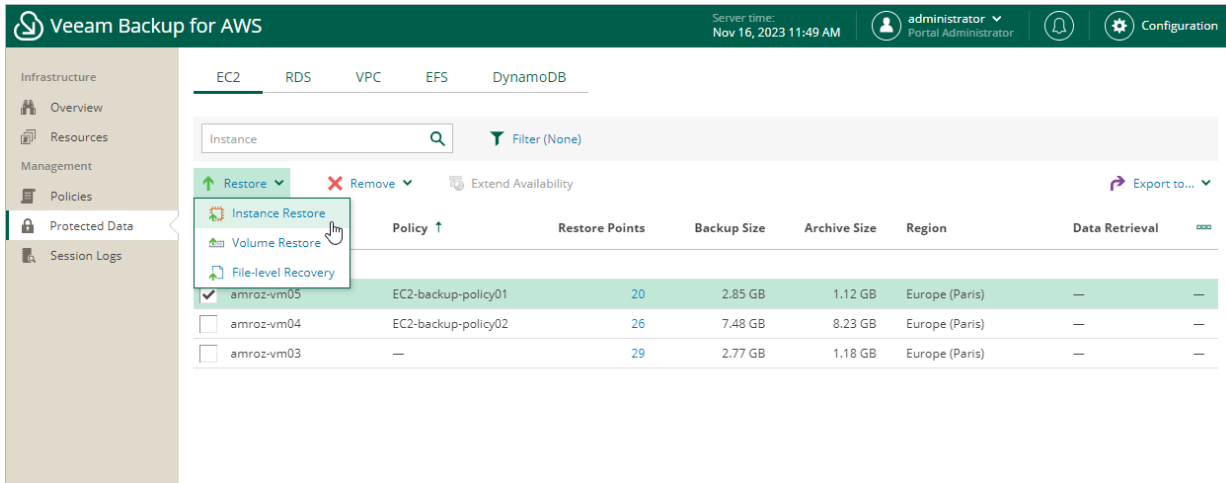
- To restore an EC2 instance from a backup that is stored in an archive backup repository, you must retrieve the archived data first. You can either retrieve the archived data manually before you begin the restore operation, or launch the data retrieval process right from the **Restore** wizard. To learn 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.
- Veeam Backup for AWS does not support restore of IPv6 addresses, tags of Elastic IP addresses, prefixes assigned to Amazon EC2 network interfaces, and the source/destination checking settings configured for 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, Veeam Backup for AWS will raise a warning. If you decide to proceed with the restore operation, 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 allocated to the restored instance.
 - 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 and all network interfaces attached to it will be automatically deleted from AWS, unless [termination protection](#) is enabled for the instance. In the latter case, Veeam Backup for AWS will not be able to restore the EC2 instance and will raise an error notifying that you must disable termination protection on the source instance.
- 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.
- **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).

- o **IAM Role** – an IAM role used to create the restore point (for cloud-native snapshots and snapshot replicas).

Choose instances to restore

Instance	Type	Restore Point
amroz-vm05	Snapshot	10/20/2023

Choose restore point

Date	Type	State	Storage Class	Restore Point Region
10/19/2023 10:00...	Snapshot	—	—	Europe (Paris)
10/13/2023 10:00...	Snapshot	—	—	Europe (Paris)
10/17/2023 10:00...	Snapshot	—	—	Europe (Paris)
10/18/2023 10:00...	Snapshot	—	—	Europe (Paris)
08/31/2023 2:42:...	Replica	—	—	Europe (Milan)
10/20/2023 10:00...	Snapshot	—	—	Europe (Paris)
09/29/2023 10:00...	Archive	✓ Healthy	S3 Glacier Flexible ...	Europe (Paris)
10/10/2023 10:00...	Backup	✓ Healthy	S3 Standard	Europe (Paris)
10/06/2023 10:00...	Backup	✓ Healthy	S3 Standard	Europe (Paris)
09/08/2023 10:00...	Backup	✓ Healthy	S3 Standard	Europe (Paris)
10/15/2023 10:00...	Backup	✓ Healthy	S3 Standard	Europe (Paris)
09/29/2023 10:00...	Backup	✓ Healthy	S3 Standard	Europe (Paris)
09/01/2023 10:00...	Backup	✓ Healthy	S3 Standard	Europe (Paris)

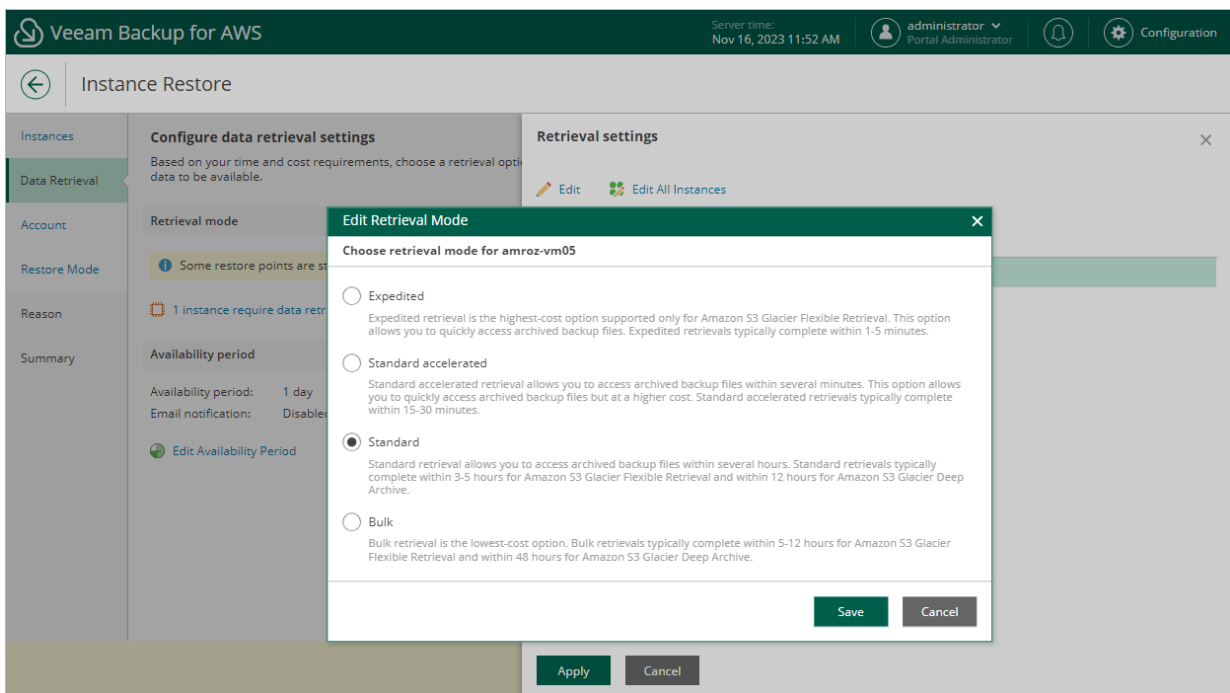
Apply Cancel

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, '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 has a left sidebar with navigation options: 'Instances', 'Data Retrieval', 'Account', 'Restore Mode', 'Reason', and 'Summary'. The 'Data Retrieval' section is active, 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 add an IAM role, click **Add** and complete the [Add IAM Role](#) wizard.

IMPORTANT

It is recommended that you check whether the selected IAM role has all the required permissions 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 launches worker instances used to perform restore operations in the [backup account](#). However, you can instruct Veeam Backup for AWS to launch 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 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 *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 add an IAM role, 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, specifically the 'Account' step. The interface includes a top navigation bar with the Veeam logo, server time (Nov 16, 2023 12:02 PM), and user information (administrator, Portal Administrator). The left sidebar contains navigation options: Instances, Data Retrieval, Account (selected), Restore Mode, Reason, and Summary. The main content area is titled 'Choose IAM role and specify worker deployment settings'. It features two radio buttons: 'IAM role' (selected) and 'Temporary access keys'. Under 'IAM role', there is a dropdown menu showing 'Default Backup Restore (Default Backup Restore)' and buttons for '+ Add' and 'Check Permissions'. Under 'Temporary access keys', there are input fields for 'Access key:' and 'Secret key:', along with 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.' Below this, the 'Worker deployment' section has a toggle switch for 'Deploy workers in production account:' set to 'On'. It also includes a dropdown menu for 'IAM role:' showing 'Worker role (IAM role used to launch worker in production accounts)' and buttons for '+ Add' and 'Check Permissions'. A note at the bottom of this section states: 'To be able to restore instances with volumes encrypted using default AWS managed keys, it is required to deploy worker instances in the production account.' At the bottom of the wizard, there are three buttons: 'Previous', 'Next', and 'Cancel'.

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 [termination protection](#) is 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 'Europe (Paris)'. A link 'Not set...' is visible next to the 'Select AWS Outpost:' label. The interface includes a sidebar with navigation options like 'Instances', 'Data Retrieval', 'Account', 'Restore Mode', 'Encryption', 'Settings', 'Network', 'Reason', and 'Summary'. 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 number (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 'Instance Restore' wizard in Veeam Backup for AWS. The 'Encryption' step is active, showing options to 'Use original encryption scheme' (unselected) or 'Restore as encrypted instance' (selected). An 'Encryption key' dropdown menu is set to 'am-key'. A help message states: 'To learn how to work with AWS encryption keys, see this Veeam KB article.' The bottom navigation bar includes 'Previous', 'Next', and 'Cancel' buttons.

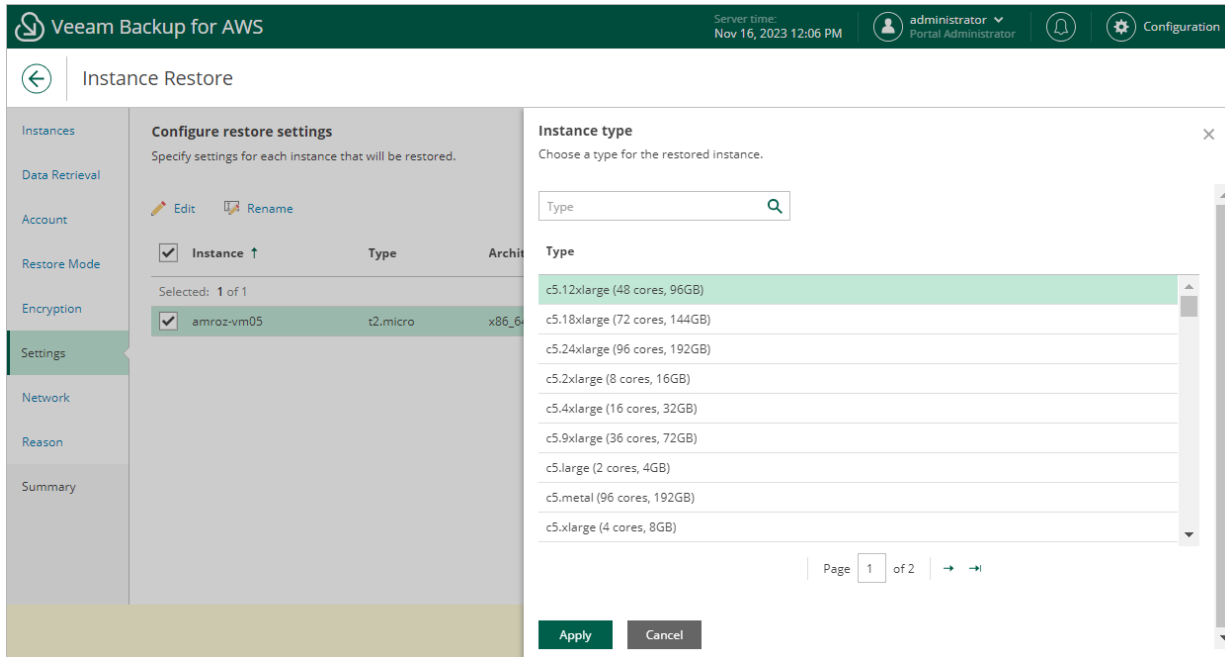
Step 7. Specify Instance Settings

[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, or the original Amazon machine image (AMI) that was used to launch one of the source instances has not been found]

At the **Settings** step of the wizard, do the following for each EC2 instance added to the restore session:

- To specify a new name for the restored EC2 instance, select the source instance from the list and click **Rename**. In the **Instance name** window, specify the name and click **Apply**.
- To change the instance type for the restored EC2 instance, select the source instance from the list and click **Edit**. In the **Instance type** window, select the necessary instance type and click **Apply**. For the list of all existing instance types, see [AWS Documentation](#).
- [This step applies only if the original AMI that was used to launch the source instance has not been found] To specify an AMI that will be used to launch the restored EC2 instance, select the source instance from the list and click **Change AMI**. In the **Instance settings** window, select an AMI that will be used to perform the restore operation.

By default, Veeam Backup for AWS automatically selects an AMI whose configuration is similar to the configuration of the restored instance and allows you to change the selected AMI if required.



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 to which Amazon VPC a restored EC2 instance must be connected, select a subnet in which the EC2 instance will be launched and security groups that must 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**.

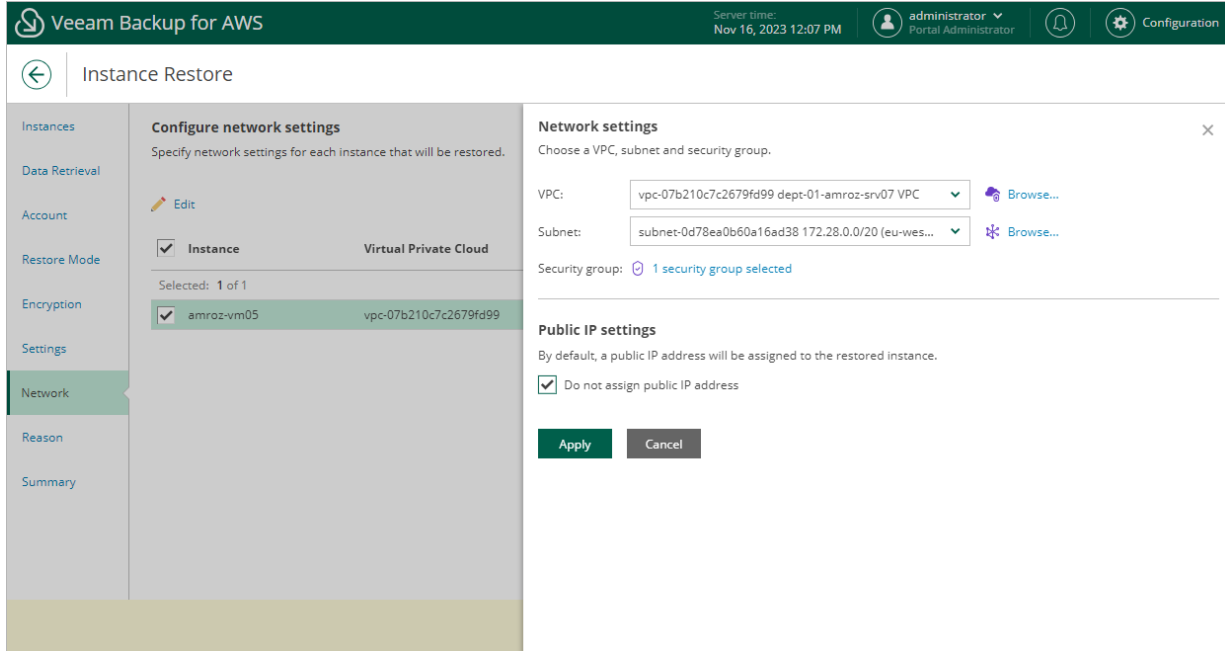
For a VPC, subnet and security group to be displayed in the lists of available network specifications, they must be created in the AWS Region specified at [step 5](#) of the wizard as described in [AWS Documentation](#).

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 EC2 instances. The information you provide will be saved in the session history and you can reference it later.

The screenshot shows the 'Instance Restore' wizard in Veeam Backup for AWS. The interface includes a top navigation bar with the Veeam logo, 'Veeam Backup for AWS', server time 'Nov 16, 2023 12:08 PM', and user information 'administrator Portal Administrator'. A left sidebar lists steps: Instances, Data Retrieval, Account, Restore Mode, Encryption, Settings, Network, Reason (highlighted), and Summary. The main area is titled 'Restore reason' and contains a text input field with the text 'Restoring failed EC2 instance'. At the bottom, 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.

The screenshot shows the 'Instance Restore' wizard in the 'Summary' step. The interface includes a top navigation bar with the Veeam logo, server time (Nov 16, 2023 12:08 PM), and user information (administrator, Portal Administrator). A left sidebar lists various settings categories: Instances, Data Retrieval, Account, Restore Mode, Encryption, Settings, Network, Reason, and Summary (which is highlighted). The main content area is titled 'Review configured settings' and contains the following information:

- General settings**
 - IAM role name: Default Backup Restore
 - Restore mode: New location
 - Location name: Europe (Paris)
 - Encryption: Restore as an encrypted instance
 - KMS key: am-key
- Restore settings**
 - Items: 1 will be restored
- Reason**
 - Reason: Restoring failed EC2 instance

At the bottom, there is a checked checkbox for 'Power on target instances after restoring'. At the very bottom of the wizard, there are three buttons: 'Previous', 'Finish', and 'Cancel'.

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 attached to a protected EC2 instance, do the following:

1. [Launch the Volume Restore wizard](#).
2. [Select a restore point](#).
3. [Specify data retrieval settings for archived backups](#).
4. [Specify restore settings](#).
5. [Choose a restore mode](#).

6. [Enable encryption for EBS volumes.](#)
7. [Specify the restored EBS volume name.](#)
8. [Specify a restore reason.](#)
9. [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:

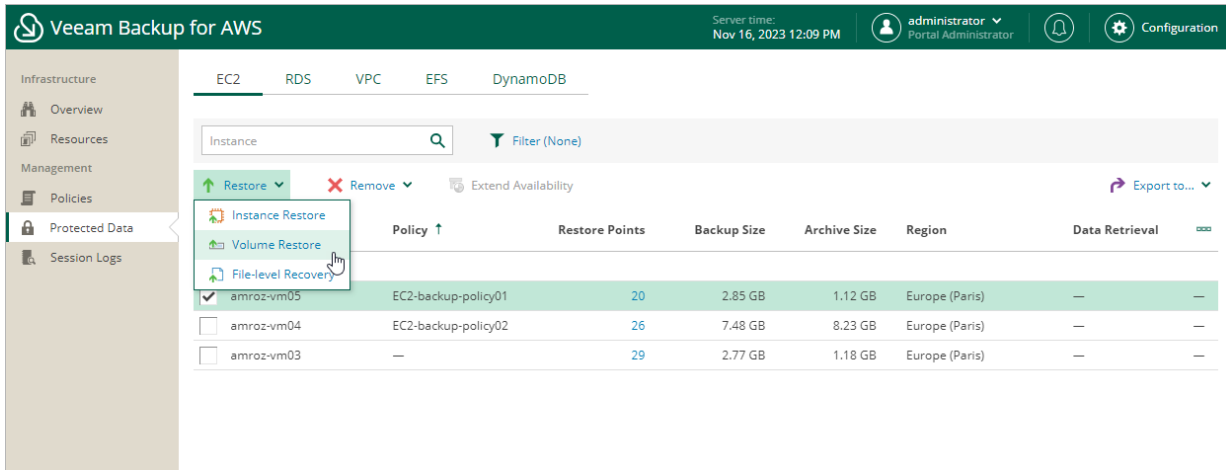
1. An IAM role you plan to specify for the restore operation must have the following permissions: `outposts:ListOutposts`, `outposts:GetOutpostInstanceTypes`. To grant the necessary permissions for the IAM role, use the [AWS Management Console](#).
2. If the Outpost subnet is specified in the [worker 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).
- **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 Volume Restore interface. The top navigation bar shows the Veeam logo, the product name 'Veeam Backup for AWS', the server time 'Nov 16, 2023 12:10 PM', and the user 'administrator Portal Administrator'. The main window is titled 'Volume Restore' and is divided into two main sections: 'Choose instances which volumes you want to restore.' and 'Choose restore point'.

The 'Choose instances' section contains a search bar and a table with the following data:

Instance	Type	Restore Point
amroz-vm05	Snapshot	10/20/2023

The 'Choose restore point' section contains a table with the following data:

Date	Type	State	Storage Class	Restore Point Region
10/08/2023 10:00...	Backup	Healthy	S3 Standard	Europe (Paris)
09/08/2023 10:00...	Backup	Healthy	S3 Standard	Europe (Paris)
10/15/2023 10:00...	Backup	Healthy	S3 Standard	Europe (Paris)
09/29/2023 10:00...	Backup	Healthy	S3 Standard	Europe (Paris)
09/01/2023 10:00...	Backup	Healthy	S3 Standard	Europe (Paris)
10/17/2023 10:00...	Backup	Healthy	S3 Standard	Europe (Paris)
10/14/2023 10:00...	Backup	Healthy	S3 Standard	Europe (Paris)
10/16/2023 10:00...	Backup	Healthy	S3 Standard	Europe (Paris)
09/29/2023 10:00...	Archive	Healthy	S3 Glacier Flexible ...	Europe (Paris)
10/20/2023 10:00...	Snapshot	—	—	Europe (Paris)
10/17/2023 10:00...	Snapshot	—	—	Europe (Paris)
10/19/2023 10:00...	Snapshot	—	—	Europe (Paris)
10/13/2023 10:00...	Snapshot	—	—	Europe (Paris)
08/31/2023 2:42:...	Replica	—	—	Europe (Milan)
10/18/2023 10:00...	Snapshot	—	—	Europe (Paris)

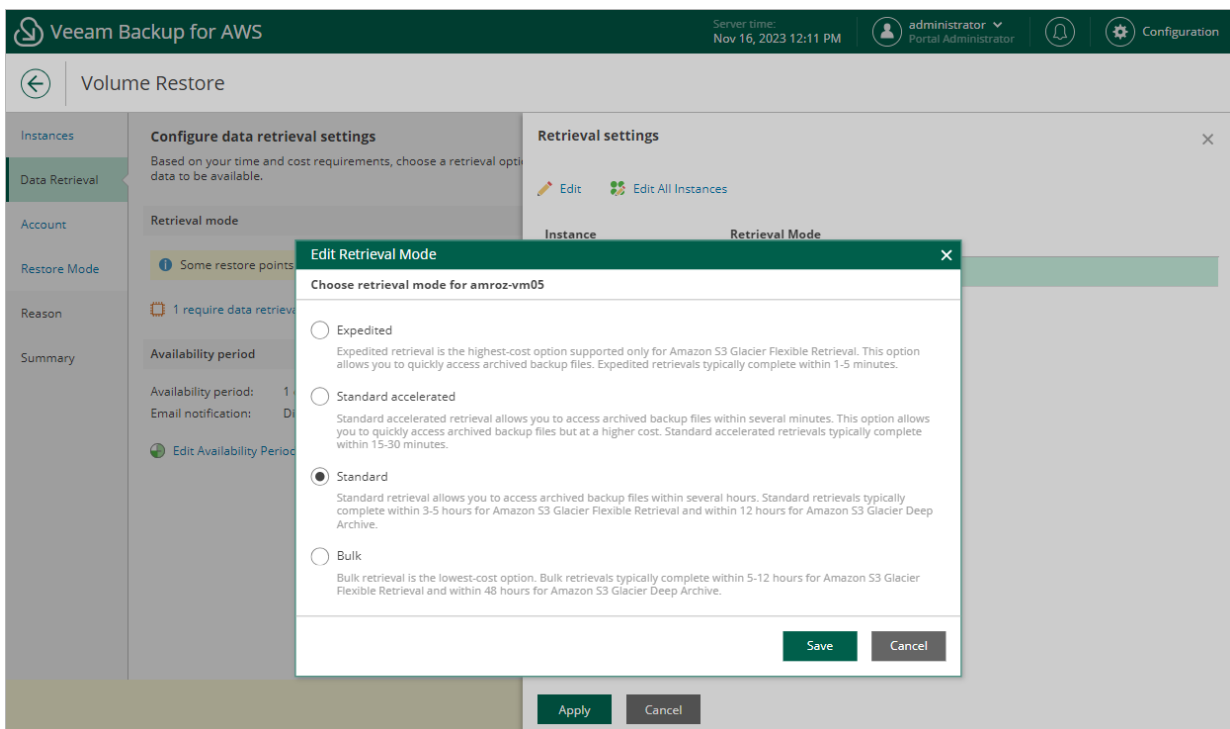
At the bottom of the 'Choose restore point' panel, there are 'Apply' and 'Cancel' buttons.

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, '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 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: 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 add an IAM role, click **Add** and complete the [Add IAM Role](#) wizard.

IMPORTANT

It is recommended that you check whether the selected IAM role has all the required permissions 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 launches worker instances used to perform restore operations in the [backup account](#). However, you can instruct Veeam Backup for AWS to launch 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 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 *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 add an IAM role, 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 sidebar with 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', and an information box stating that keys are not saved or stored. The 'Worker deployment' section has a toggle for 'Deploy workers in production account' set to 'On'. An information box explains that for encrypted volumes, a worker role must be specified. A dropdown menu for 'Worker role (IAM role used to launch worker in production accounts)' is shown with an '+ Add' button. 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 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 'Restore Mode' step is active, showing two options: 'Restore to original location' (unselected) and 'Restore to new location, or with different settings' (selected). The selected option includes configuration fields for 'Region' (set to 'Europe (Milan)') and 'Availability zone' (set to 'eu-south-1a'). A link 'Not set...' is visible for 'Select AWS Outpost:'. The interface includes a sidebar with navigation options (Instances, Data Retrieval, Account, Restore Mode, Encryption, Settings, Reason, Summary) and a top navigation bar with user information and server time.

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 number (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.

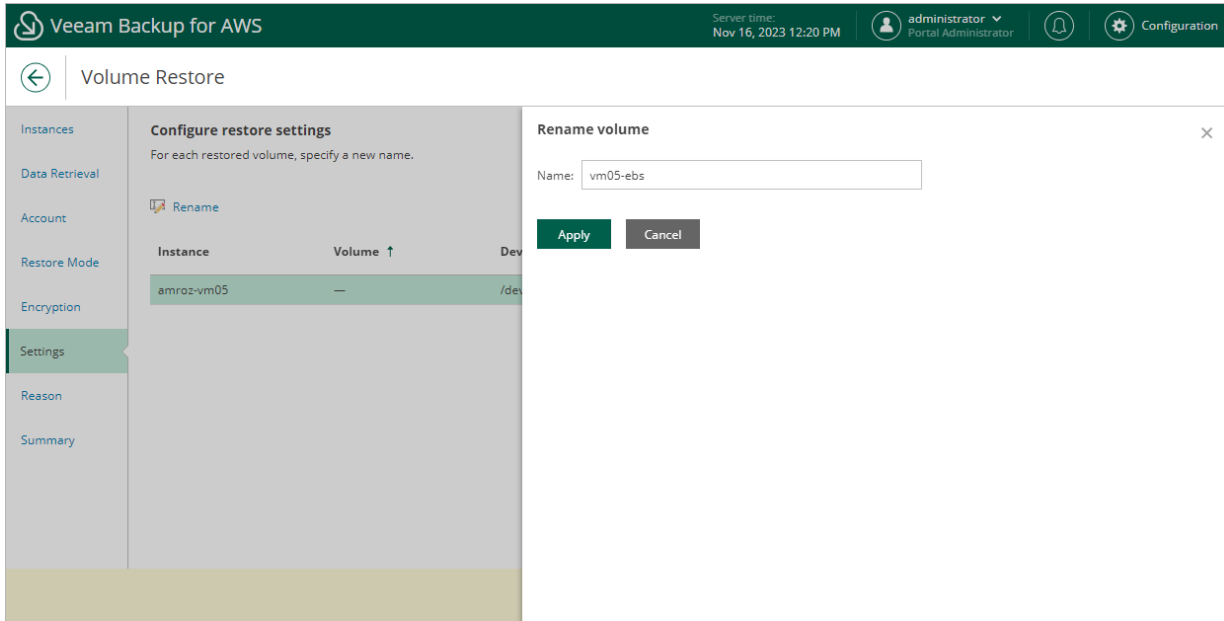
The screenshot shows the 'Volume Restore' wizard in Veeam Backup for AWS. The 'Encryption' step is active, showing options to 'Configure encryption settings'. The 'Restore as encrypted volume' option is selected. An 'Encryption key' dropdown menu is set to 'am-key'. A help message is displayed: 'To learn how to work with AWS encryption keys, see this Veeam KB article.' The navigation bar at the bottom includes 'Previous', 'Next', and 'Cancel' buttons.

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**.



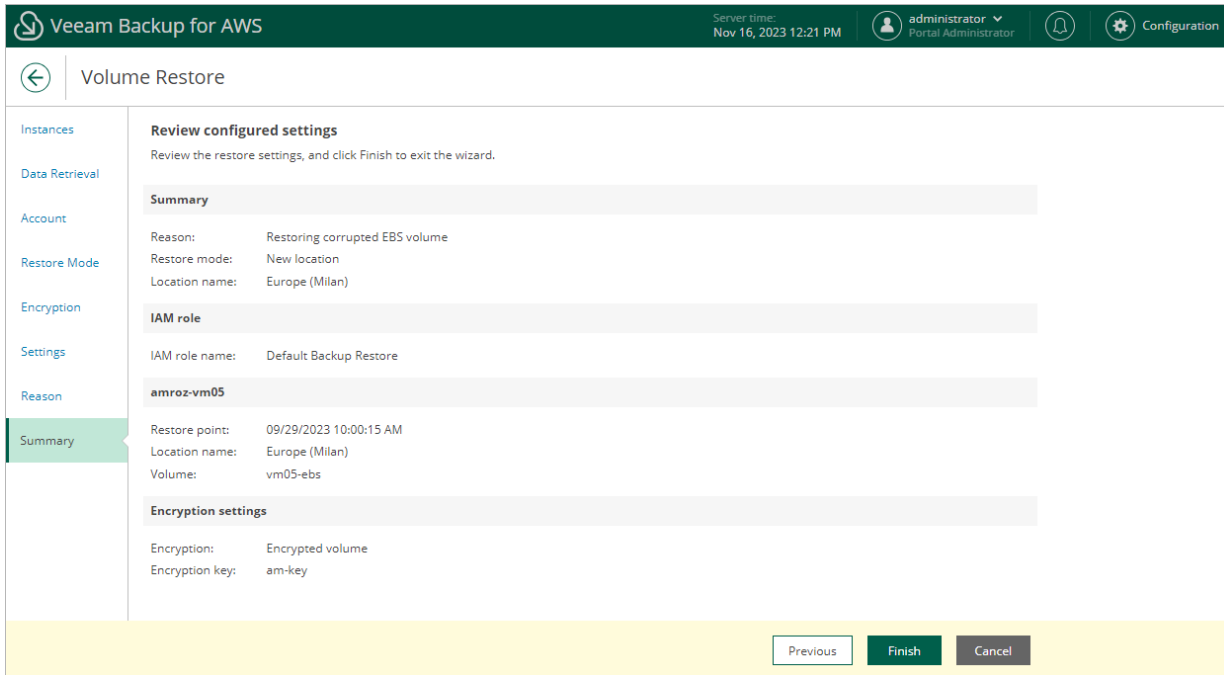
Step 8. Specify Restore Reason

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 'Volume Restore' wizard in Veeam Backup for AWS. The interface includes a top navigation bar with the Veeam logo, 'Veeam Backup for AWS', server time 'Nov 16, 2023 12:21 PM', and user information 'administrator Portal Administrator'. A left sidebar contains navigation links: 'Instances', 'Data Retrieval', 'Account', 'Restore Mode', 'Encryption', 'Settings', 'Reason' (highlighted), and 'Summary'. The main content area is titled 'Restore reason' and contains a text input field with the text 'Restoring corrupted EBS volume'. 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**.



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. [Launch the EC2 File-level Recovery wizard.](#)
2. [Select a restore point.](#)
3. [Specify restore settings.](#)
4. [Specify a restore reason.](#)
5. [Finish working with the wizard – start a recovery session.](#)
6. [Choose files and folders to recover.](#)
7. [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 launches 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:
 - `com.amazonaws.<region>.ec2messages`
 - `com.amazonaws.<region>.ssm`
 - `com.amazonaws.<region>.sqs`
 - `com.amazonaws.<region>.kinesis-streams`
 - `com.amazonaws.<region>.sts`

To learn how to create interface VPC endpoints, see [AWS Documentation](#).

Step 1. Launch EC2 File-level Recovery Wizard

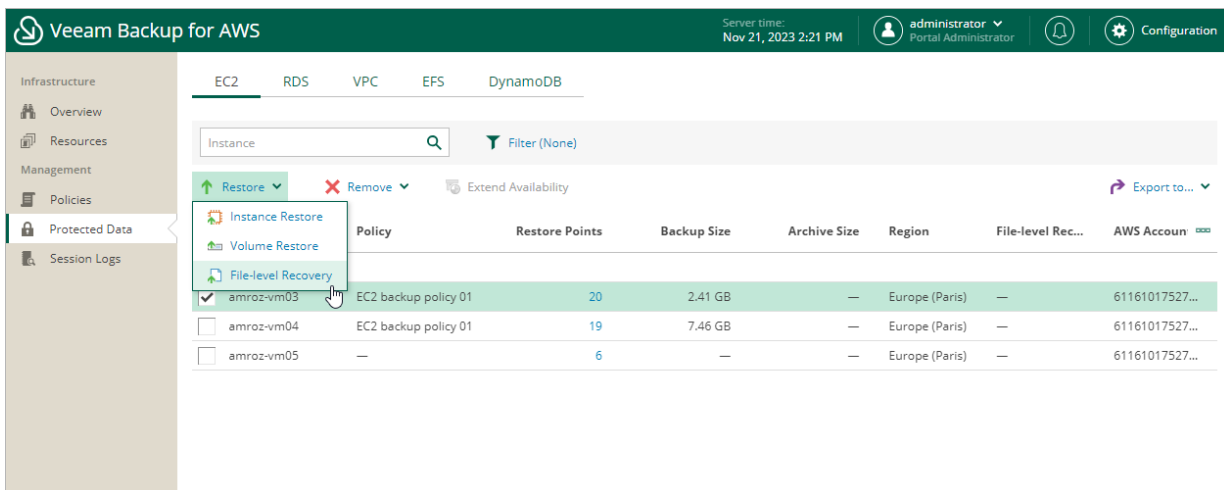
To launch the **EC2 File-level Recovery** wizard, do the following:

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).
- **IAM Role** – an IAM role used to create the restore point (for cloud-native snapshots and snapshot replicas).

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 displays the Veeam Backup for AWS interface. At the top, the header shows 'Veeam Backup for AWS' and 'Server time: Nov 21, 2023 2:24 PM'. The user is logged in as 'administrator Portal Administrator'. The main window is titled 'EC2 File-level Recovery' and is divided into two panes.

The left pane, 'Choose instances to restore', contains a search bar and a table with the following data:

Instance	Type	Restore Point
amroz-vm03	Snapshot	10/19/2023

The right pane, 'Choose restore point', contains a table with the following data:

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	Healthy	—	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)

At the bottom of the right pane, there are 'Apply' and 'Cancel' buttons.

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 launches worker instances used to perform restore operations in the [backup account](#). However, you can instruct Veeam Backup for AWS to launch 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 launch 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 add an IAM role, click **Add** and complete the [Add IAM Role](#) wizard.

IMPORTANT

It is recommended that you check whether the selected IAM role has all the required permissions 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. The top bar displays the server time as Nov 21, 2023 2:25 PM and the user as administrator (Portal Administrator). The main window is titled "EC2 File-level Recovery" and is divided into several sections: "Instances", "Configure restore settings", and "Permission check".

The "Permission check" section shows a message: "Your account does not meet the required permissions." Below this message are three buttons: "Grant", "Recheck", and "Export Missing Permissions". A table below lists the permissions and their status:

Type	Status	Missing Permissions
EC2MESSAGES permissions	Passed	—
SQS permissions	Passed	—
SSM permissions	Passed	—

The "Grant Permissions" dialog box is open, showing a form to provide temporary credentials. The form includes an information 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](#)." The form has two input fields: "Access key:" with the value "AKIA4ZWOU4WMVRAGEVN" and "Secret key:" with a masked value. There are "Apply" and "Cancel" buttons at the bottom of the dialog.

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 'EC2 File-level Recovery' wizard in the Veeam Backup for AWS console. The interface includes a top navigation bar with the Veeam logo, 'Veeam Backup for AWS', server time 'Nov 21, 2023 2:26 PM', and user information 'administrator Portal Administrator'. A left sidebar contains navigation links for 'Instances', 'Restore Settings', 'Reason' (highlighted), and 'Summary'. The main content area is titled 'Restore reason' and contains a text input field with the text 'restoring corrupted files'. At the bottom of the wizard, 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 launch 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:

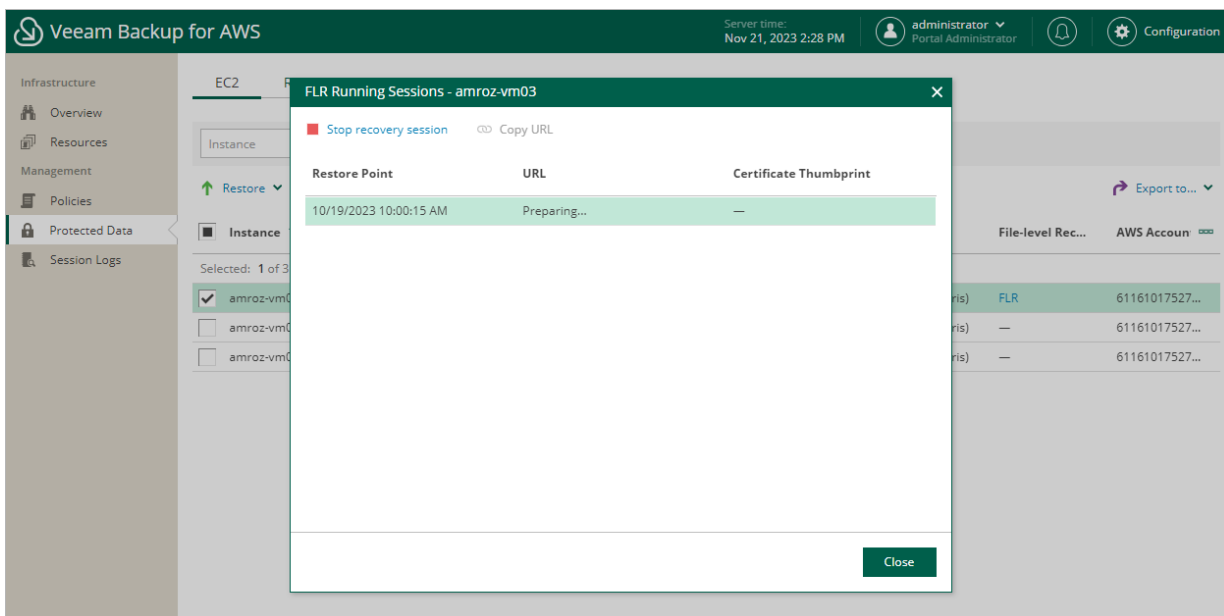
- Click the link to open the file-level recovery browser on your local machine while the recovery session is running.
- Copy the link, close the **FLR Running Sessions** window and open the file-level recovery browser on another machine.

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.

The screenshot displays the Veeam Backup for AWS interface. At the top, there are two tabs: "Browse" and "Restore List (3)". The "Restore List" tab is active, showing a list of items for recovery. The list has columns for Name, Location, Size, Last Modified, Restore Point, Restore Date, and Restore Status. Two items, "dev" and "home", are selected. Below the list is a "Session Log" section with a table for tracking restore operations.

Restore List: amroz-vm03

Restore Status: All [Green Check] [Yellow Warning] [Red X]

Download [Red Stop] Remove [X]

<input type="checkbox"/>	Name ↑	Location	Size	Last Modified	Restore Point	Restore Date	Restore Status
<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	—	—

Selected: 2 of 3

Session Log

Status: All [Green Check] [Yellow Warning] [Red X]

Action	Status	Start Time	End Time	Duration
Select a single item to view sessions details				

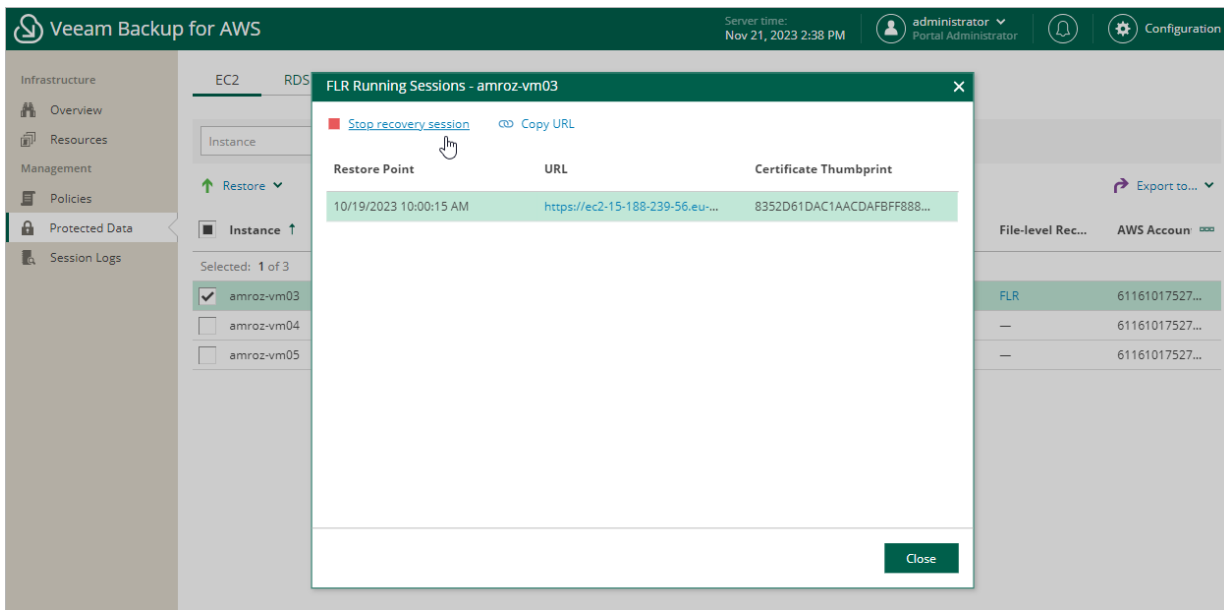
Step 7. Stop Recovery Session

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 DB instance running the PostgreSQL database engine.

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**.

NOTE

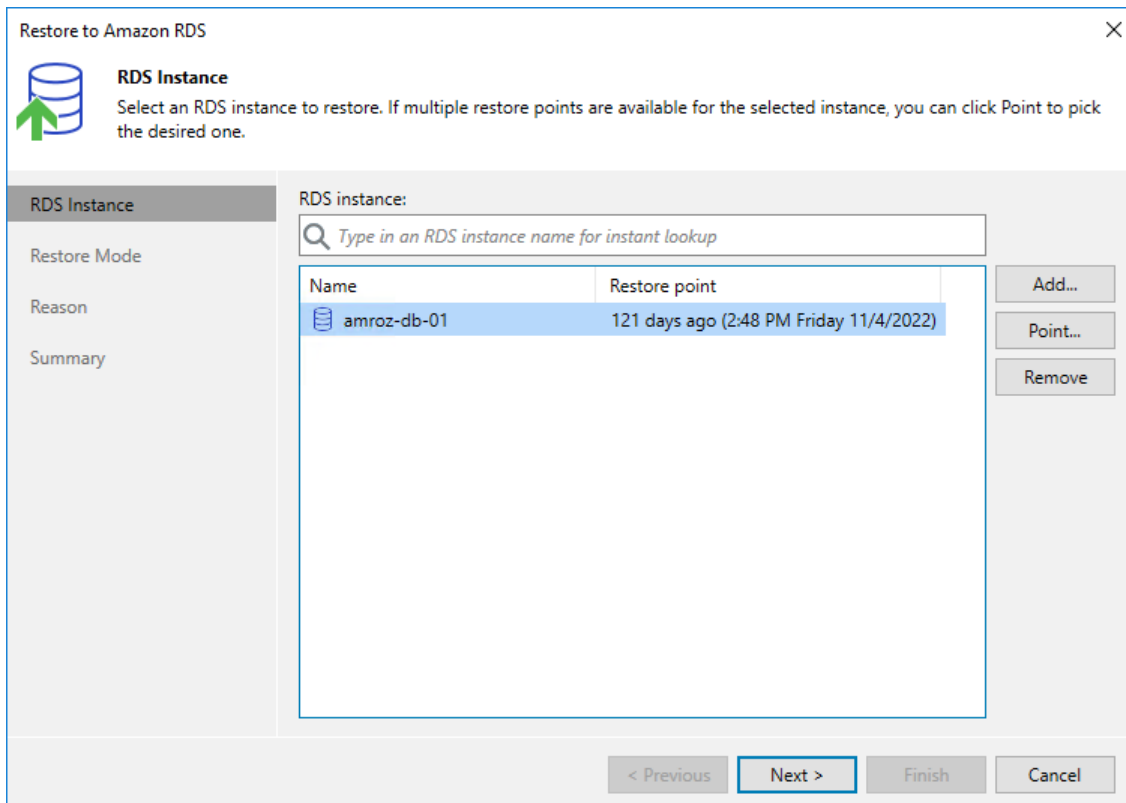
If you want to restore a DB instance from an Amazon DB snapshot created in AWS, expand the *<Appliance name>* node and select the necessary snapshot of an *AWS Snapshot* type in the **Restore Points** window, and then click **OK**.

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 in the following cases:

- If the restore point that you have selected at [step 2](#) of the wizard is of the *AWS Snapshot* type.
- If the IAM role that will be used to perform the restore operation belongs to an AWS account that differs from the AWS account where the source resources belong.

2. Click **Pick account 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 one-time 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 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 *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](#).

Restore to Amazon RDS

Restore Mode
Specify whether selected RDS instances should be restored back to the original location, or to a new location or with different settings.

RDS Instance

Restore Mode

Data Center

Instance Type

Instance Configuration

Identifier

Network

Reason

Summary

Restore to the original location
Quickly initiate restore of the selected RDS 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 RDS instance location, and change its settings. The wizard will automatically populate all controls with the original RDS instance settings as the defaults.

[Pick account to use](#)

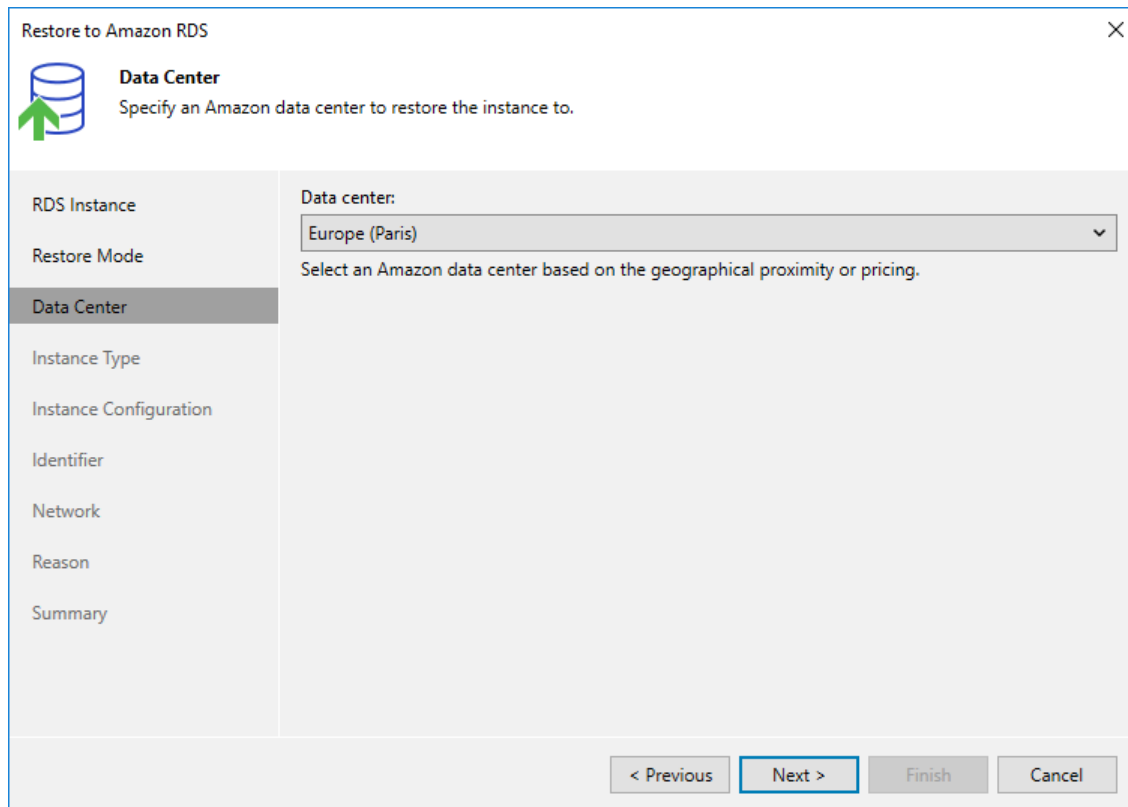
< Previous Next > Finish Cancel

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 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 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:' dropdown menu with 'Europe (Paris)' selected. 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 >', '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](#).

The screenshot shows the 'Restore to Amazon RDS' wizard at the 'Instance Type' step. The main window has a sidebar with navigation options: RDS Instance, Restore Mode, Data Center, Instance Type (selected), Instance Configuration, Identifier, Network, Reason, and Summary. The main content area shows a table of RDS instances:

Name	Instance type	Encryption
amroz-db-01	db.t3.micro	Preserve original settings

An 'Instance Type' dialog box is open over the table, showing the following configuration:

- RDS instance type: db.t3.micro (2 cores, 1.00 GB memory)
- Disk type: General Purpose SSD (GP2)
- Provisioned IOPS SSD (IO1) 6000
- Magnetic

Buttons for 'OK' and 'Cancel' are at the bottom of the dialog. Below the table, there are buttons for 'Type...' and 'Encryption...'. At the bottom of the wizard, there are navigation buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'.

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.

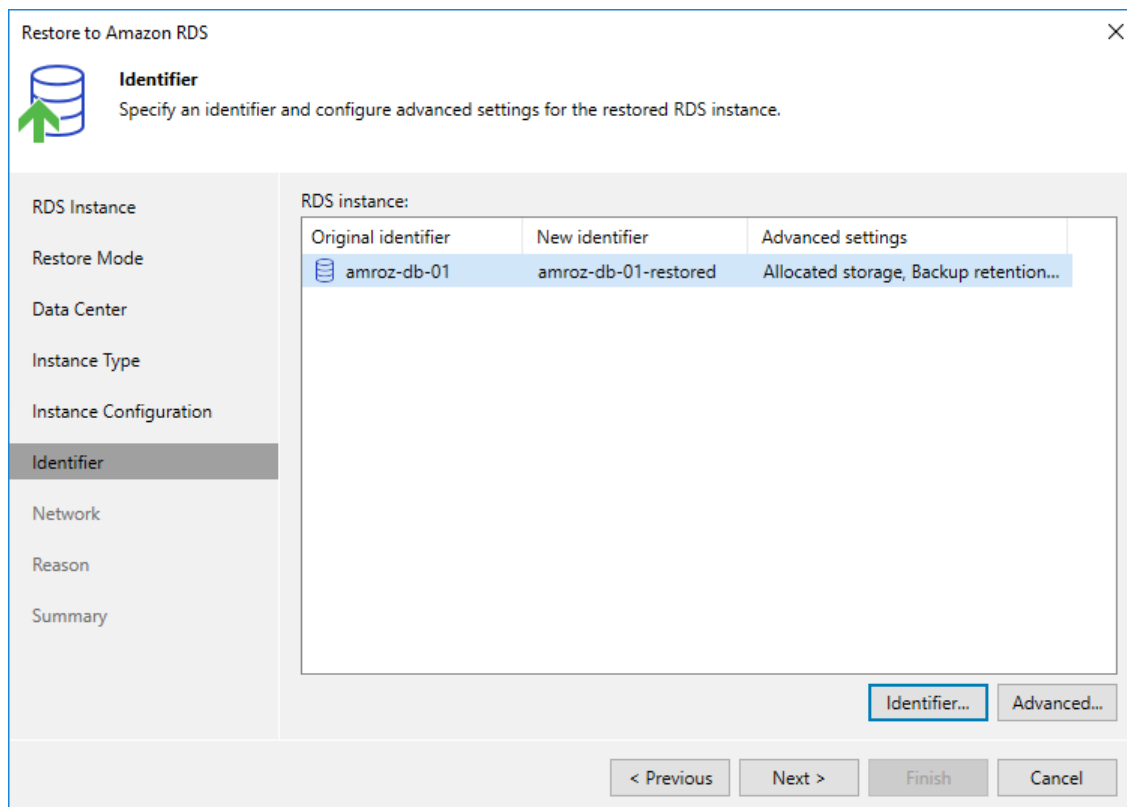
Consider the following limitations:

- 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, in the main window, are 'Customize...' and 'Availability...' buttons. At the very bottom of the wizard 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. [Launch the Restore Amazon RDS Cluster wizard.](#)
2. [Select a restore point.](#)
3. [Choose a restore mode.](#)
4. [Select an AWS Region.](#)
5. [Choose capacity type and enable encryption.](#)
6. [Specify cluster and instance parameter groups.](#)
7. [Specify cluster and database identifiers.](#)
8. [Configure network and availability settings.](#)
9. [Specify a restore reason.](#)
10. [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](#).

Public Access

The security group associated with the restored Aurora DB cluster 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.

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 where the source cluster belongs.
- Restore of Aurora DB clusters to the original location is not supported using restore points of the *AWS Snapshot* type – you can restore these resources only to a new location.
- 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.

Note that restore of Aurora multi-master clusters using restore points of the *AWS Snapshot* type is supported only to the source region within the same AWS account.

- 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 the following DB engine versions: MySQL 5.6.10a, MySQL 2.07.1, PostgreSQL 10.12 and PostgreSQL 10.14.
- Aurora Serverless v1 is supported for a limited list of AWS Regions and specific DB engine versions. For more information, see [AWS Documentation](#).

Step 1. Launch Restore Amazon RDS Cluster Wizard

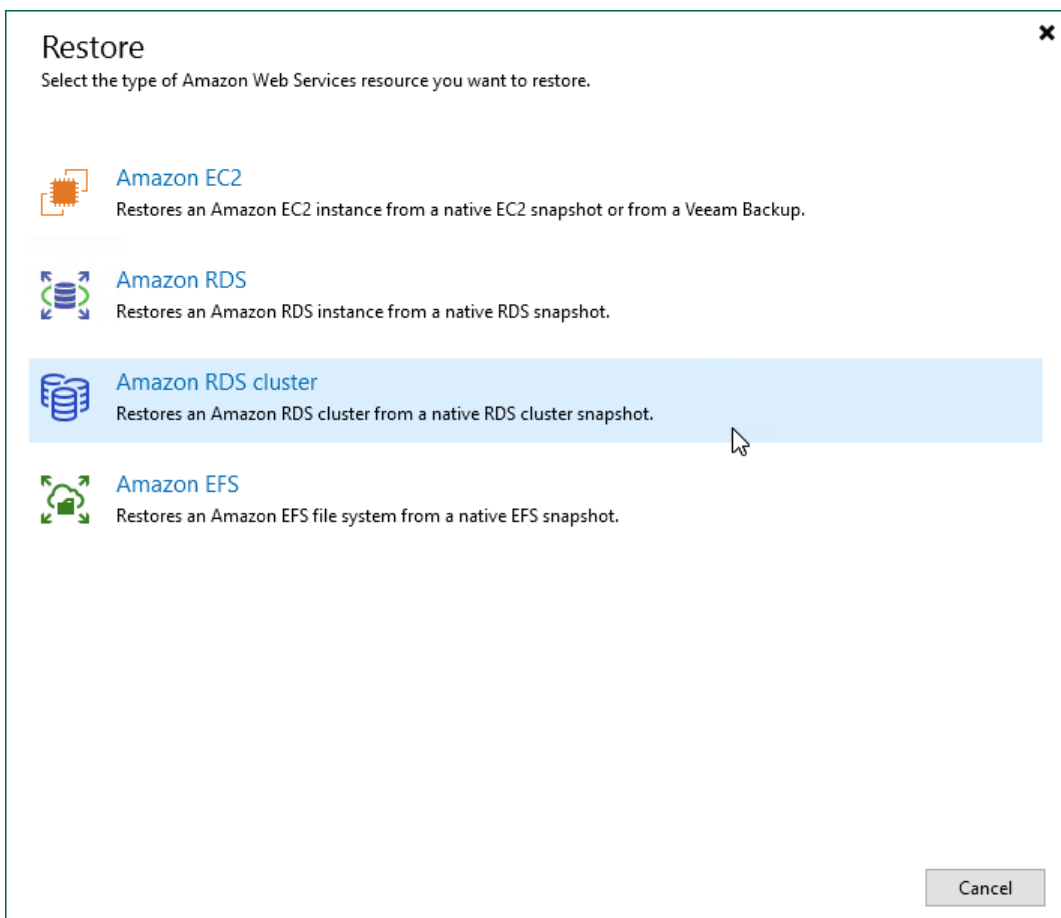
To launch the **Restore to Amazon RDS cluster** wizard, do the following:

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**.

NOTE

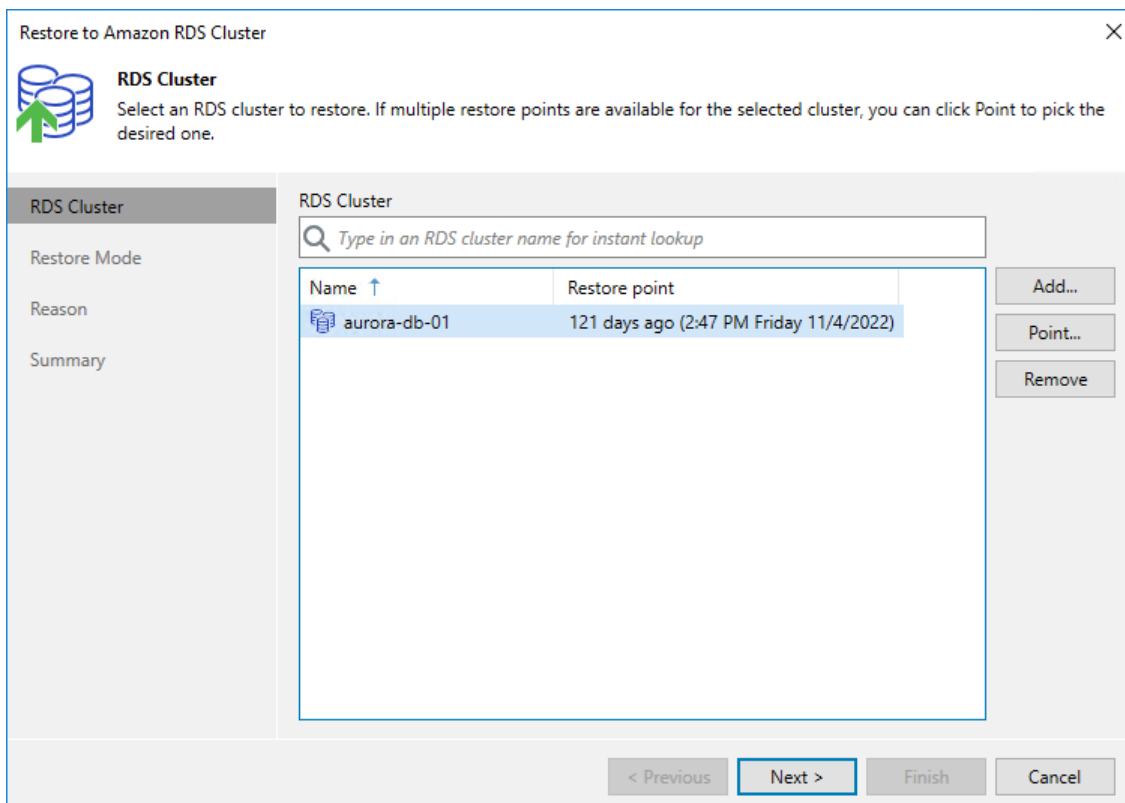
If you want to restore an Aurora DB cluster from an Amazon DB snapshot created in AWS, expand the *<Appliance name>* node and select the necessary snapshot of an *AWS Snapshot* type in the **Restore Points** window, and then click **OK**.

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](#).

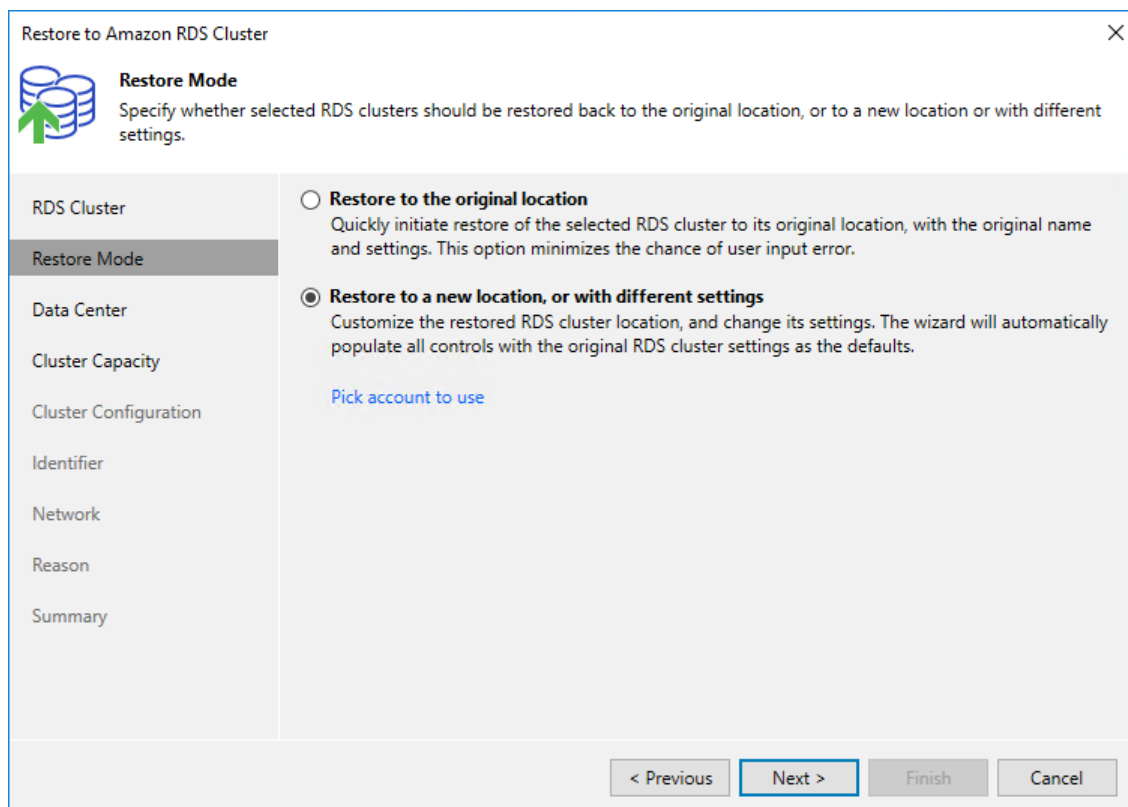
1. Choose whether you want to restore the Aurora DB cluster to the original or to a new location.
2. Click **Pick account to use** to select an IAM identity whose permissions will be used to perform the restore operation:
 - To specify an IAM role for the restore operation, select the **IAM role** option 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 one-time 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 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 *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 RDS Cluster" with a close button (X) in the top right corner. On the left is a navigation pane with the following items: "RDS Cluster", "Restore Mode" (highlighted), "Data Center", "Cluster Capacity", "Cluster Configuration", "Identifier", "Network", "Reason", and "Summary". An icon of three database cylinders with a green arrow pointing up is positioned to the left of the "Restore Mode" header. Below the icon, the text reads: "Restore Mode Specify whether selected RDS clusters should be restored back to the original location, or to a new location or with different settings." The main content area contains two radio button options:

- Restore to the original location**
Quickly initiate restore of the selected RDS cluster 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 RDS cluster location, and change its settings. The wizard will automatically populate all controls with the original RDS cluster settings as the defaults.
[Pick account to use](#)

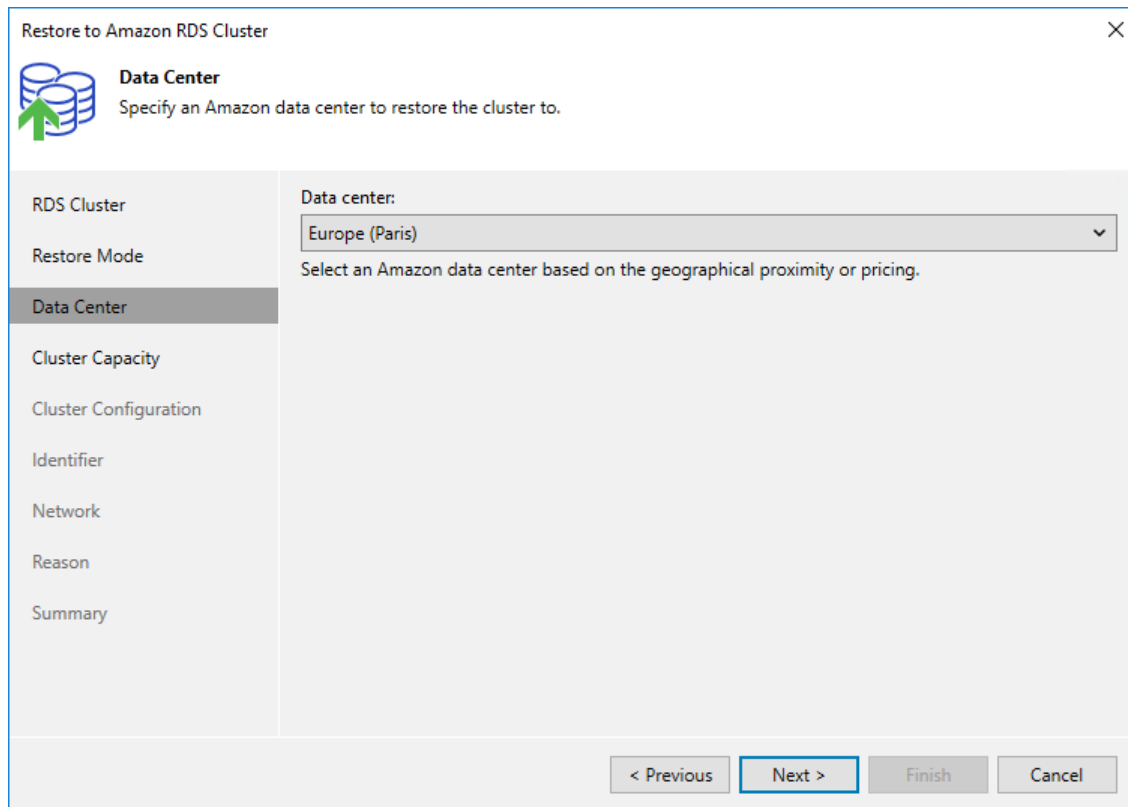
At the bottom of the window are four buttons: "< Previous", "Next >" (highlighted with a blue border), "Finish", and "Cancel".

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.
- 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](#).

NOTE

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](#).

Restore to Amazon RDS Cluster

Cluster Capacity
Specify the capacity and disk encryption settings for the restored RDS cluster.

RDS Cluster: **aurora-db-01**

Name	Compute	Engine version	Encryption
aurora-db-01	Provisioned	13.7	Preserve original settings

Capacity Settings

Select capacity options for the restored cluster. You can either provision the fixed amount of compute resources or let AWS automatically scale capacity based on the database load.

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

OK Cancel Encryption...

< Previous Next > Finish Cancel

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.

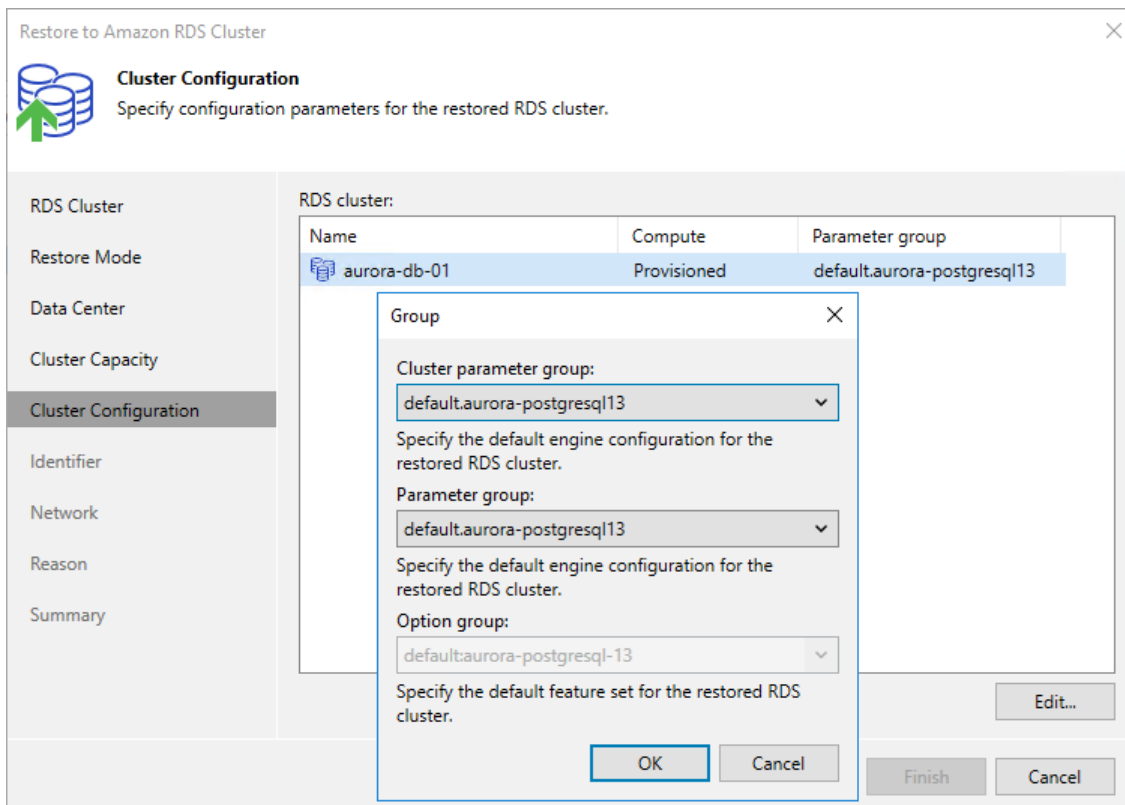
For a DB cluster parameter group to be displayed in the list of available groups, the group must be created in AWS as described in [AWS Documentation](#).

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](#).

Restore to Amazon RDS Cluster

Identifier
Specify an identifier and configure advanced settings for the restored RDS cluster.

RDS Cluster

Restore Mode

Data Center

Cluster Capacity

Cluster Configuration

Identifier

Network

Reason

Summary

RDS cluster:

Name	Compute	Engine version	Cluster identifier
aurora-db-01	Provisioned	13.7	aurora-db-01-restored

Identifier

Cluster identifier:
aurora-db-01-restored
Specify database cluster identifier.

Instance identifier:
aurora-db-01-instance-1-restored
Specify database instance identifier.

OK Cancel

Identifier... Advanced...

< Previous Next > Finish Cancel

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.

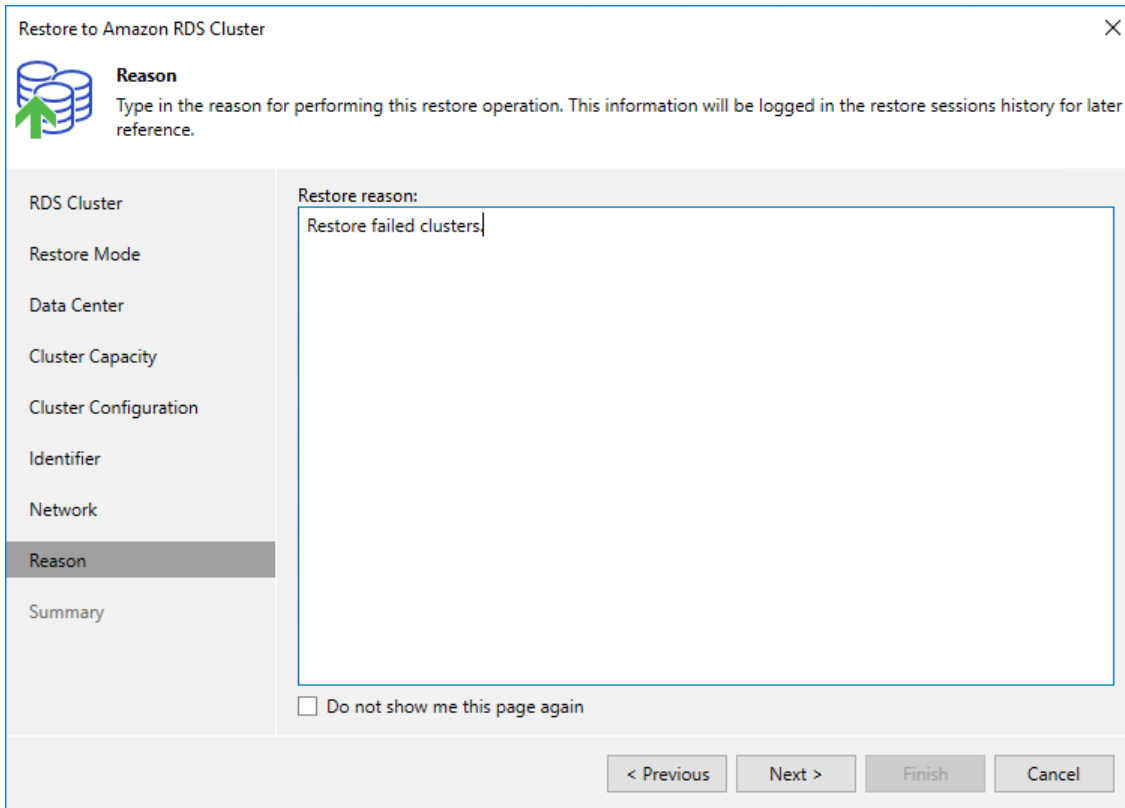
The screenshot shows the 'Restore to Amazon RDS Cluster' wizard at the 'Network' step. A dialog box titled 'Amazon VPC' is open, allowing configuration of network settings. The dialog includes the following fields:

- 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 wizard window shows a sidebar with 'Network' selected and a 'Next >' button highlighted.

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 green icon of a database with an upward arrow and the heading "Reason". Below the icon, the text reads: "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 label "Restore reason:" above a large text input field. The input field contains 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**.



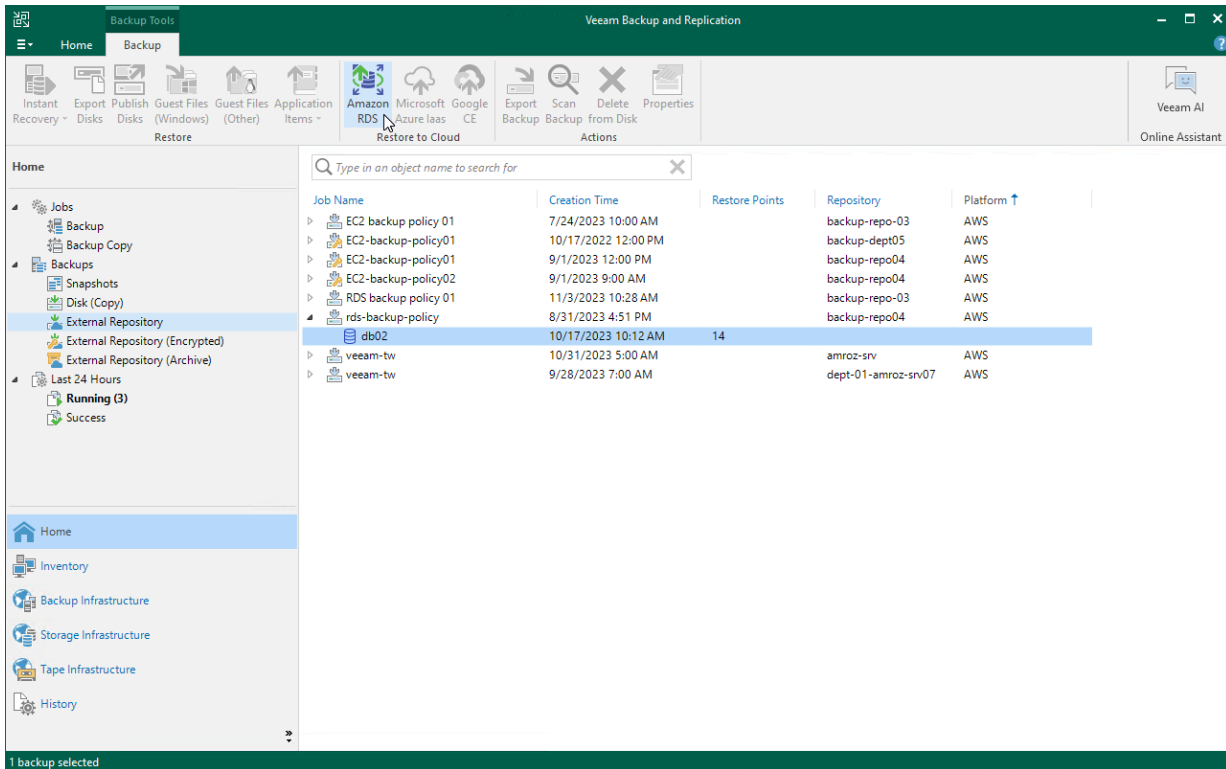
Restoring RDS Databases

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, snapshot replica or an AWS snapshot. 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.

NOTE

Restore of RDS resources with gp3 storage volumes is not supported. For more information on General Purpose gp3 storage volumes, see [AWS Documentation](#).

How to Perform RDS Restore

To restore a protected RDS resource, do the following:

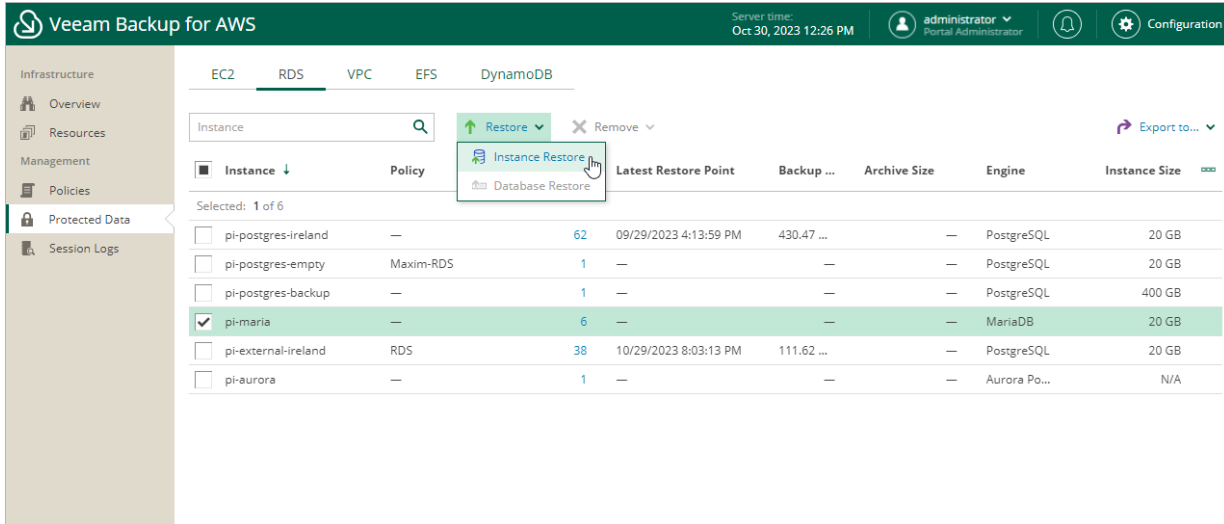
1. [Launch the RDS Restore wizard](#).
2. [Select a restore point](#).
3. [Specify an IAM identity for restore](#).
4. [Choose a restore mode](#).
5. [Enable encryption](#).
6. [Configure RDS instance settings](#).
7. [Configure network settings](#).
8. [Specify a restore reason](#).
9. [Finish working with the wizard](#).

Step 1. Launch RDS Restore Wizard

To launch the **RDS Restore** wizard, do the following:

1. Navigate to **Protected Data > 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**.



The screenshot shows the Veeam Backup for AWS interface. The top navigation bar includes the Veeam logo, server time (Oct 30, 2023 12:26 PM), user (administrator), and Configuration settings. The left sidebar shows navigation options: Infrastructure, Overview, Resources, Management, Policies, Protected Data (selected), and Session Logs. The main content area is titled 'RDS' and displays a table of instances. A dropdown menu is open over the 'Restore' button, showing 'Instance Restore' and 'Database Restore' options. The table below shows a list of RDS instances with columns for Instance, Policy, Latest Restore Point, Backup, Archive Size, Engine, and Instance Size.

Instance	Policy	Latest Restore Point	Backup ...	Archive Size	Engine	Instance Size
<input type="checkbox"/> pi-postgres-ireland	—	62	09/29/2023 4:13:59 PM	430.47 ...	—	PostgreSQL 20 GB
<input type="checkbox"/> pi-postgres-empty	Maxim-RDS	1	—	—	—	PostgreSQL 20 GB
<input type="checkbox"/> pi-postgres-backup	—	1	—	—	—	PostgreSQL 400 GB
<input checked="" type="checkbox"/> pi-maria	—	6	—	—	—	MariaDB 20 GB
<input type="checkbox"/> pi-external-ireland	RDS	38	10/29/2023 8:03:13 PM	111.62 ...	—	PostgreSQL 20 GB
<input type="checkbox"/> pi-aurora	—	1	—	—	—	Aurora Po... N/A

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.
- **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.
 - *AWS Snapshot* – an Amazon DB snapshot created in AWS.

IMPORTANT

If you select a restore point of the **AWS Snapshot** type, you will not be able to restore an RDS resource to the original location.

- **Restore Point Region** – the AWS Region where the restore point is stored (for cloud-native snapshots and snapshot replicas).

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 'Oct 30, 2023 12:37 PM', and user information 'administrator Portal Administrator'. The main content area is titled 'RDS Restore' and is divided into two panels. The left panel, 'Specify instances to restore', contains a search bar and a table with columns 'Instance' and 'Engine'. The right panel, 'Choose restore point', displays a table of restore points with columns 'Date', 'Size', 'Type', and 'Restore Point Region'. The table lists 16 restore points, all 20 GB in size and of type 'Snapshot', located in the 'Europe (Ireland)' region. The most recent restore point, dated 10/27/2023 at 8:04:09 PM, is highlighted in green. At the bottom of the right panel are 'Apply' and 'Cancel' buttons.

Date ↓	Size	Type	Restore Point Region
10/29/2023 8:03:13 PM	20 GB	Snapshot	Europe (Ireland)
10/28/2023 8:02:51 PM	20 GB	Snapshot	Europe (Ireland)
10/27/2023 8:04:09 PM	20 GB	Snapshot	Europe (Ireland)
10/26/2023 8:03:19 PM	20 GB	Snapshot	Europe (Ireland)
10/25/2023 8:05:00 PM	20 GB	Snapshot	Europe (Ireland)
10/24/2023 8:03:33 PM	20 GB	Snapshot	Europe (Ireland)
10/23/2023 8:03:03 PM	20 GB	Snapshot	Europe (Ireland)
10/23/2023 12:14:26 PM	20 GB	Snapshot	Europe (Ireland)
10/23/2023 11:54:33 AM	20 GB	Snapshot	Europe (Ireland)
10/23/2023 11:22:34 AM	20 GB	Snapshot	Europe (Ireland)
10/22/2023 8:57:39 PM	20 GB	Snapshot	Europe (Ireland)
10/22/2023 8:31:18 PM	20 GB	Snapshot	Europe (Ireland)
10/20/2023 6:28:42 PM	20 GB	Snapshot	Europe (Ireland)
10/20/2023 2:50:52 PM	20 GB	Snapshot	Europe (Ireland)
10/20/2023 12:58:29 PM	20 GB	Snapshot	Europe (Ireland)

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.

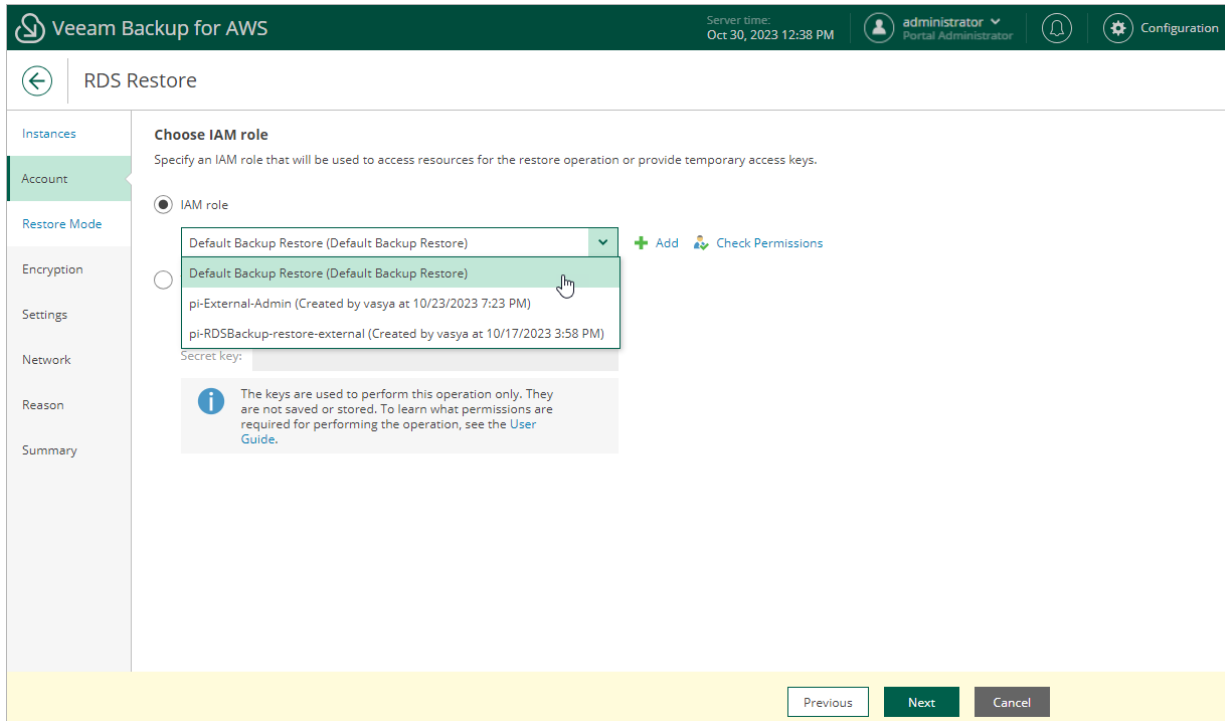
Specifying IAM Role

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 add an IAM role, click **Add** and complete the [Add IAM Role](#) wizard.

IMPORTANT

It is recommended that you check whether the selected IAM role has all the required permissions 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 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. At the top, the header includes the Veeam logo, the text 'Veeam Backup for AWS', the server time 'Oct 30, 2023 12:40 PM', and the user 'administrator Portal Administrator'. A navigation menu on the left lists 'Instances', 'Account', 'Restore Mode', 'Encryption', 'Settings', 'Network', 'Reason', and 'Summary'. The main content area is titled 'RDS Restore' and contains the 'Choose IAM role' section. This section instructs the user to '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 'Default Backup Restore (Default Backup Restore)' and buttons for '+ Add' and 'Check Permissions'. Under 'Temporary access keys', there are two input fields: 'Access key:' with the value 'DFRT6TGGDJKLP' and 'Secret key:' with a masked value '.....'. An information icon (i) is present next to the secret key field. Below the input fields, a note states: '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 screen, 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 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:

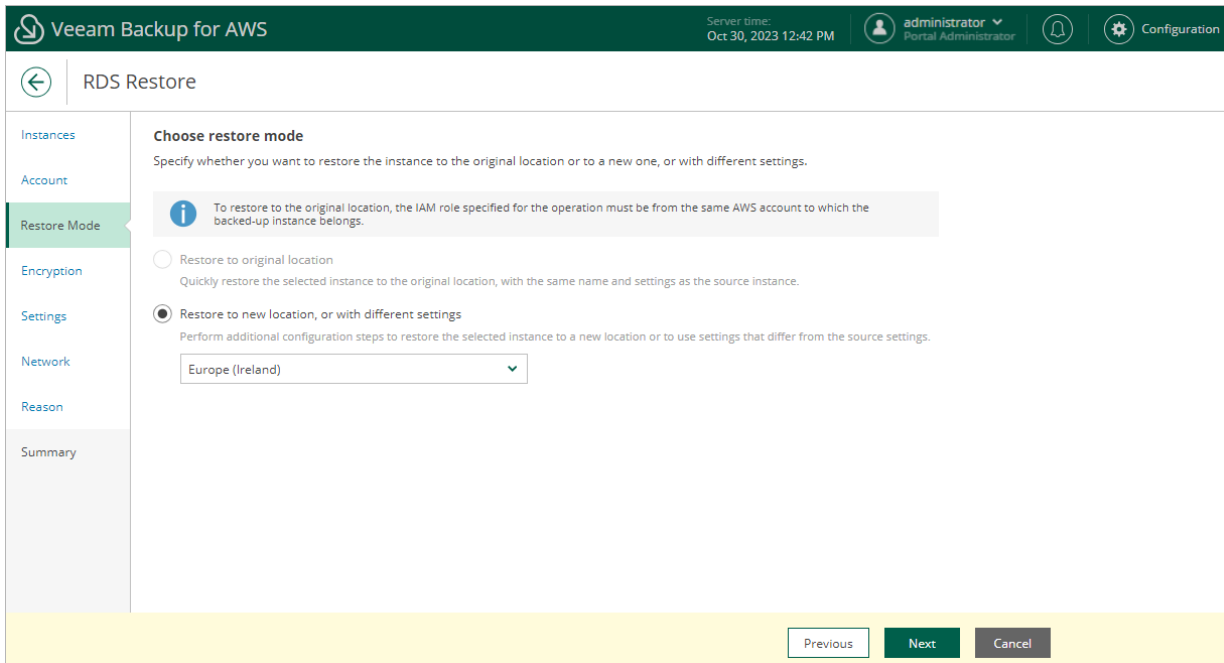
- Restore of RDS resources to the original location is not supported if the IAM role specified for the restore operation belongs to an AWS account that differs from the AWS account where the source resources belong.
- Restore of RDS resources to the original location is not supported using restore points of the **AWS Snapshot** type – you can restore these resources only to a new location.
- Restore of RDS resources to the original location is not supported if **termination protection** is enabled for the source resource.
- 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 another AWS account. To specify an AWS account in which the clusters will be restored, select an IAM role that belongs to the necessary account at [step 3](#) of the wizard.

Note that restore of Aurora multi-master clusters using restore points of the **AWS Snapshot** type is supported only to the source region within the same AWS account.

- 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. 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](#).



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](#).

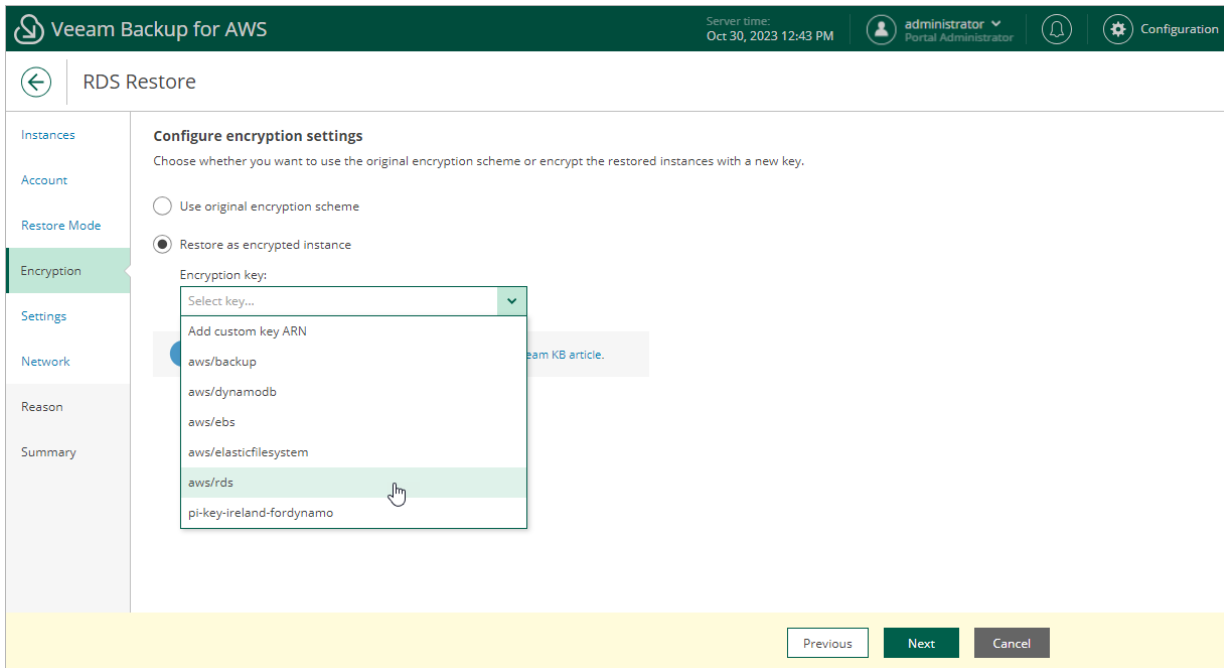
- If you want to encrypt the RDS resources, select the **Restore as encrypted instance** option and choose the necessary KMS key from the **Encryption key** list.

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 number (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 'RDS Restore' configuration window in Veeam Backup for AWS. The interface includes a sidebar on the left with navigation options: Instances, Account, Restore Mode, Encryption, Settings (selected), Network, Reason, and Summary. The main content area is titled 'Configure restore settings' and contains a table of existing instances:

Name	Engine	Instance Class
pi-aurora	Aurora P...	—
pi-maria	MariaDB	—
pi-external-irela...	PostgreS...	db.t3.micro

Below the table, there is a 'Rescan' button and configuration fields for the restored instance:

- Instance identifier:** pi-maria-restored
- Instance specifications:**
 - Instance class: db.m6g.4xlarge (16 cores, 64GB)
 - Storage type: Provisioned IOPS (SSD)
 - Provisioned IOPS: 1000
- Instance options:**
 - Parameter group: default.mariadb10.6
 - Option group: default:mariadb-10-6

An information message states: 'Changing the storage type can impact instance performance.' At the bottom of the configuration area are 'Apply' and 'Cancel' buttons.

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 with single-master replication, and multiple DB instances (writers) in Aurora multi-master clusters. For more information on provisioned DB clusters, see [AWS Documentation](#).
- **Aurora Serverless DB cluster** – 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](#).

Before you choose a capacity type for the restored cluster, consider the following limitations:

- Aurora Serverless v1 is supported only for a limited list of AWS Regions and specific DB engine versions. For more information, see [AWS Documentation](#).
- 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 the following DB engine versions: MySQL 5.6.10a, MySQL 2.07.1, PostgreSQL 10.12 and PostgreSQL 10.14.

Configuring Settings for Provisioned Cluster

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. From the **Capacity type** drop-down list, select *Provisioned*.

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.

- b. [This step applies only to Aurora MySQL DB clusters with single-master replication and Aurora PostgreSQL DB clusters] 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](#).

- c. [This step applies only to Aurora MySQL DB clusters with single-master replication] 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](#).

- d. 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.

- e. 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 two 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](#).

- f. 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](#).

NOTE

Veeam Backup for AWS supports Aurora Serverless v2. 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](#).

- 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](#).

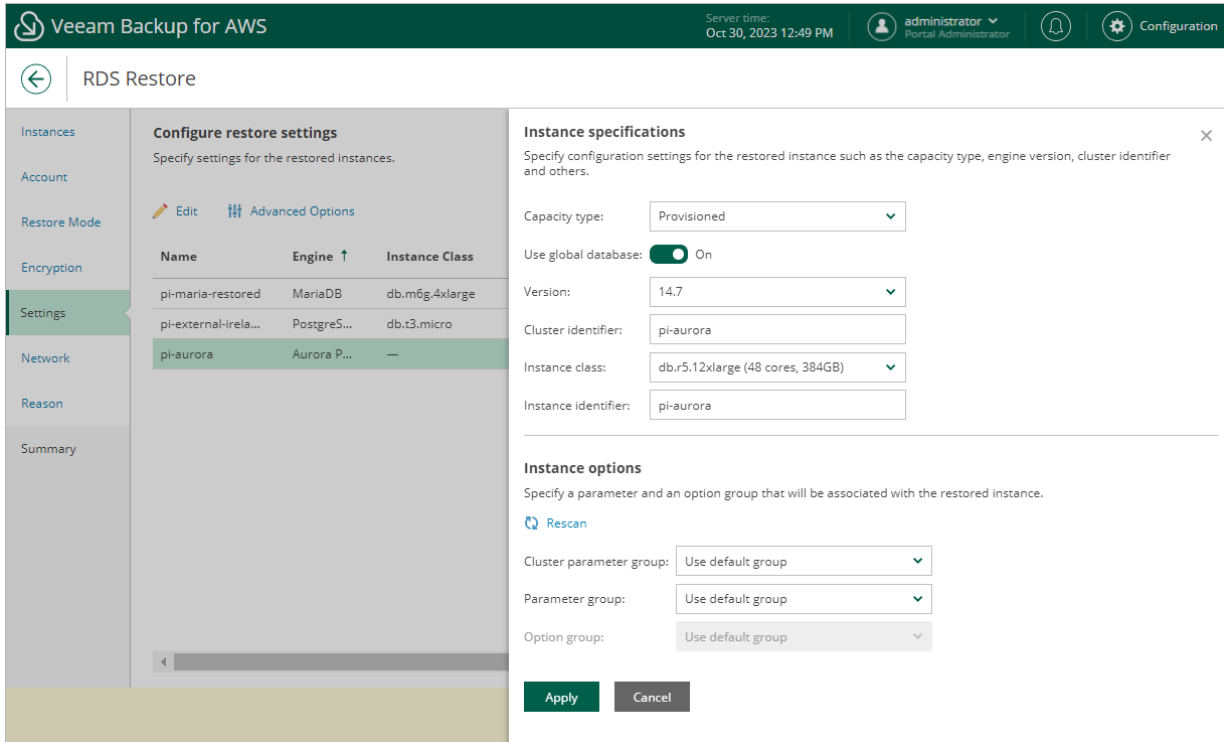
- b. From the **Parameter group** drop-down list, select the DB parameter group containing database engine configuration values that will be applied to the primary DB instance in the restored cluster.

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**.



Configuring Settings for Serverless Cluster

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. From the **Capacity type** drop-down list, select *Serverless*.
 - b. 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.

For more information on Amazon Aurora database engine versions, see [AWS Documentation](#).

- c. 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 two 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](#).

- d. Use the **Minimum capacity unit** and **Maximum capacity unit** 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](#).

- In the **Instance options** section, specify a DB cluster parameter group containing database engine configuration values that will be applied to the restored cluster.

For a DB cluster parameter group to be displayed in the **Cluster parameter group** 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 with the default DB parameter group that will be automatically created by AWS during the restore operation.

- Click **Apply**.

The screenshot displays the Veeam Backup for AWS interface for RDS Restore. The top navigation bar shows the Veeam logo, the text 'Veeam Backup for AWS', the server time 'May 27, 2022 11:51 AM', and user information 'administrator Portal Administrator'. The main interface is titled 'RDS Restore' and is divided into several sections:

- Configure restore settings:** A section for specifying settings for the restored instances, with options for 'Edit' and 'Advanced Options'.
- Instance specifications:** A section for specifying configuration settings for the restored instance, including:
 - Capacity type: Serverless
 - Version: 5.7.mysql_aurora.2.07.1
 - Cluster identifier: bev-aurora-mysql-paris-sl
 - Minimum capacity unit: ACU 4 (8 GIB RAM)
 - Maximum capacity unit: ACU 16 (32 GIB RAM)
- Instance options:** A section for specifying a cluster parameter group that will be associated with the restored instance. It includes a 'Rescan' button and a dropdown menu for the 'Cluster parameter group' set to 'default.aurora-mysql5.7'.

At the bottom of the configuration area, there are two buttons: 'Apply' and 'Cancel'. A mouse cursor is pointing at the 'Apply' button.

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](#).

TIP116

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 provisioned 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.
 - For a restored Aurora Serverless 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, and one or more security groups that will control access to the Aurora DB cluster.
2. [This step applies only to DB instances and Aurora provisioned DB clusters] In the **Security settings** section, specify security settings to control what IP addresses will be able to connect to databases on the restored RDS resource.
 - a. If you want to make the restored 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, do the following:
 - i. Click the link in the **Security** group field.
 - ii. In the **Select Security Group** window, select the necessary groups and click **Add**. Then, click **Save** to close the window.

3. Click Apply.

Veeam Backup for AWS Server time: Oct 30, 2023 12:51 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**

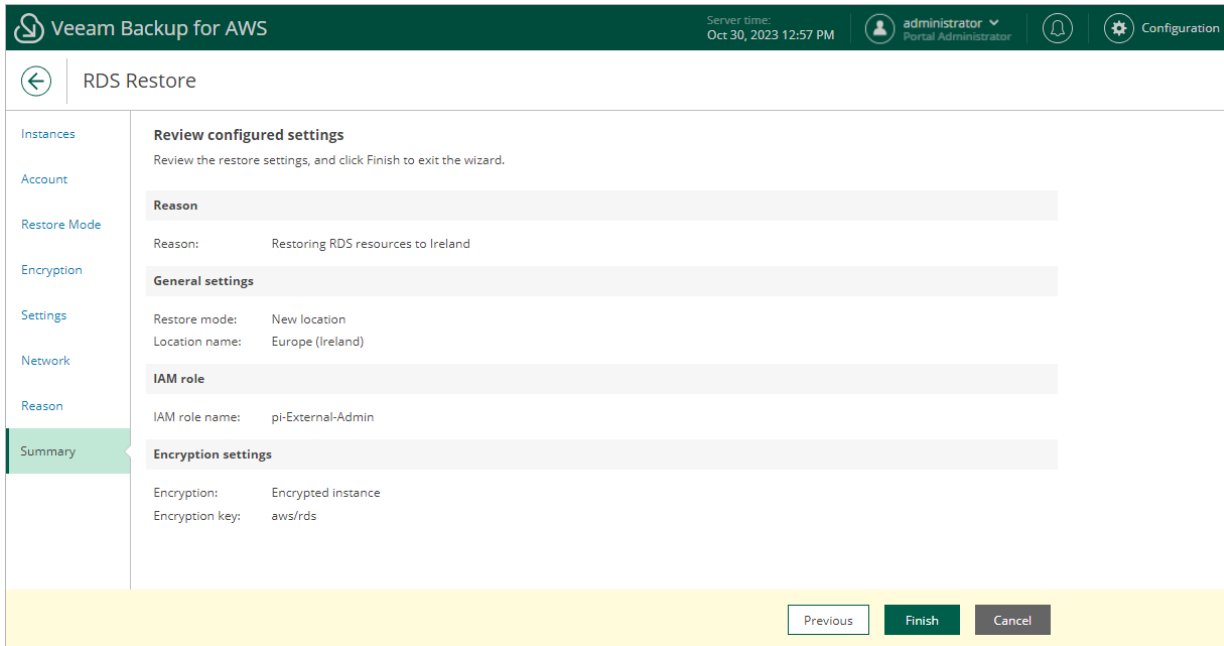
Step 8. Specify Restore Reason

At the **Reason** step of the wizard, you can specify a reason for restoring the RDS instance. 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 "Oct 30, 2023 12:56 PM", the user "administrator Portal Administrator", and a "Configuration" link. The main content area is titled "RDS Restore" and features a left-hand navigation menu with options: Instances, Account, Restore Mode, Encryption, Settings, Network, Reason (highlighted in green), and Summary. The main panel 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

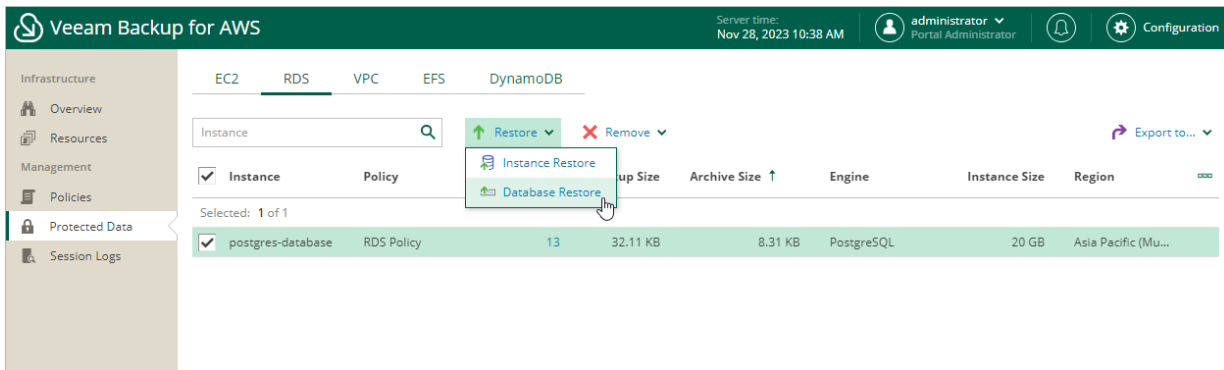
To launch the **RDS Database Restore** wizard, do the following.

1. Navigate to **Protected Data > 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:
 - *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 (for image-level backups).
- **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. The top bar displays the Veeam logo, the product name 'Veeam Backup for AWS', the server time 'Nov 28, 2023 10:39 AM', the user 'administrator', and the role 'Portal Administrator'. The main window title is 'RDS Database Restore: postgres-database'. The left sidebar shows navigation options: 'Databases', 'Account', 'Data Retrieval', 'Instance', 'Reason', and 'Summary'. The 'Databases' step is selected. The main content area is titled 'Choose databases to restore' and contains instructions: 'Choose a restore point and databases that will be used to perform'. Below this, there is a 'Restore point' section with a selected restore point: '11/28/2023 10:22:48 AM'. The 'Databases' section is currently empty, with a search bar and an '+ Add' button. An 'Add database' dialog is open, showing a table with the following data:

Database	Size	Instance Region
postgres	7.5 MB	Asia Pacific (Mumbai)

The 'postgres' database is selected. At the bottom of the dialog 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 add an IAM role, 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 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 *Worker deployment 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 add an IAM role, click **Add** and complete the **Add IAM Role** wizard.

IMPORTANT

It is recommended that you check whether the selected IAM roles have all the required permissions 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 launches 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 launch worker instances. However, you can add [specific worker configurations](#) that will be used to launch 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 launch worker instances for the database restore operation. For Veeam Backup for AWS to be able to launch a worker instance used to perform the restore operation:

- The VPC to which the DB instance is connected must have at least one security group that allows outbound access on port **443**. This port is used by worker instances 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 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: postgres-database". 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 sections: "Choose IAM role" and "Worker deployment". The "Choose IAM role" section has a sub-section "IAM role" with a text input field containing "RDS Backup and Restore Role (Created by adm)". The "Worker deployment" section has a sub-section "Worker deployment" with a text input field containing "Default Backup Restore (Default Backup Resto".

Overlaid on the right side of the interface is a "Permission check" dialog box. At the top, it displays a green checkmark and the message "Your account meets the required permissions." Below this, there are three buttons: "Grant", "Recheck", and "Export Missing Permissions". The dialog contains a table with the following structure:

<input type="checkbox"/>	Type	Status	Missing Permissions
Selected: 0 of 1			
<input type="checkbox"/>	Checking backup policy role ...	Passed	—

At the bottom of the dialog box, 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 details, 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 required permissions 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.

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 28, 2023 10:43 AM', and the user 'administrator Portal Administrator'. The main content area is titled 'RDS Database Restore: postgres-database'. 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'. This section includes a note: 'Some restore points are stored in an archive repository and...'. Below this, 'Retrieval Mode' is set to 'Standard' and 'Availability Period' is set to '3 days'. A notification email is 'Enabled (1 hour before data expiration)'. A green 'Apply' button is visible at the bottom of the settings panel. A modal window titled 'Choose retrieval mode' is open, showing three radio button options: 'Expedited', 'Standard accelerated', and 'Standard' (which is selected). A fourth option, 'Bulk', is also present. Each option has a brief description of its characteristics and completion times. At the bottom of the modal, there are 'Apply' and 'Cancel' buttons.

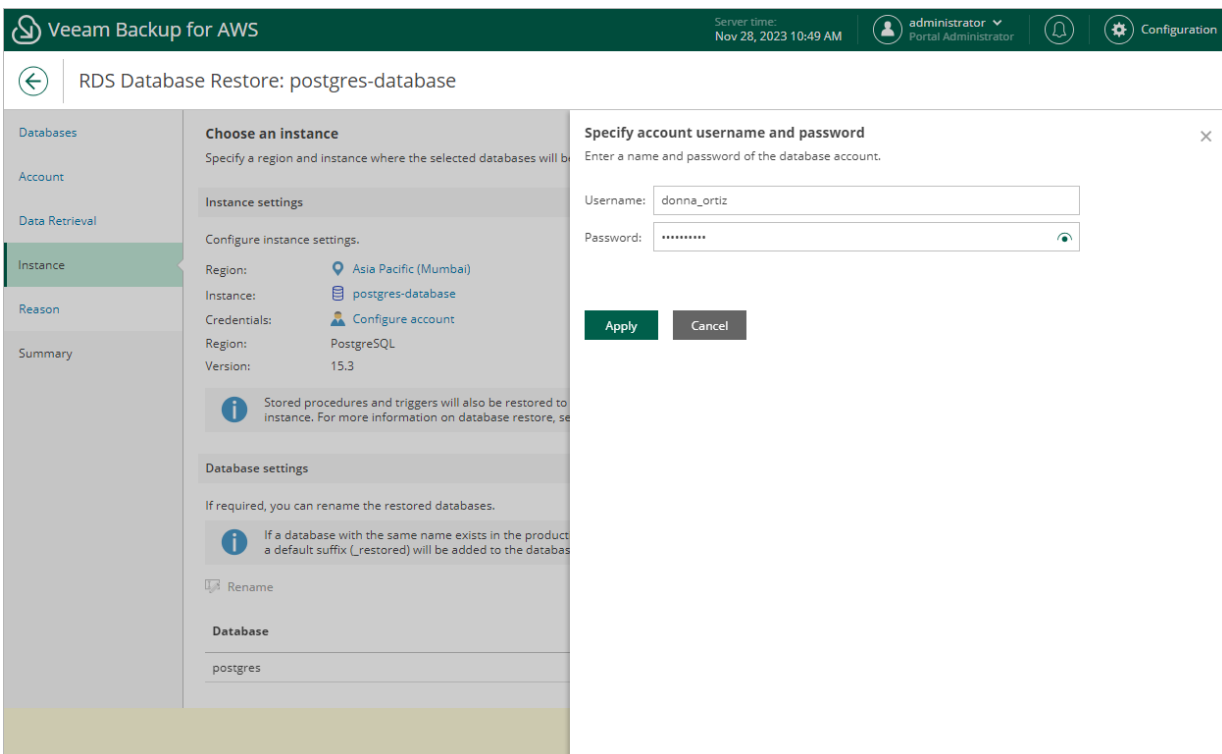
Step 5. Configure Target Instance Settings

At the **Instance** step of the wizard, do the following:

1. In the **Instance settings** section, you can specify the target AWS Region where a DB instance will host the restored databases and choose the target DB instance. By default, Veeam Backup for AWS uses the original location of the source DB instance and the source instance, if it exists.

You must also specify a database account that Veeam Backup for AWS will use to connect to the DB instance. To do that, click a link in the **Credentials** field, and provide a name and password of the account in the **Specify account username and password** window. Note that the specified user must be created on the target DB instance.

2. In the **Database settings** section, you can specify a new name for the restored database. To do that, select the database from the list and click **Rename**. In the **Database name** window, specify the name and click **Apply**.



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 shows the Veeam Backup for AWS interface. The top navigation bar includes the Veeam logo, the text "Veeam Backup for AWS", the server time "Nov 28, 2023 10:50 AM", the user "administrator" (Portal Administrator), and a "Configuration" link. The main content area is titled "RDS Database Restore: postgres-database". On the left, a sidebar lists the wizard steps: Databases, Account, Data Retrieval, Instance, Reason (highlighted), and Summary. The main area is titled "Reason" and contains the instruction "Specify the reason for performing the restore operation." Below this is a text input field with the text "Restoring PostgreSQL database". At the bottom, 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: postgres-database' wizard in Veeam Backup for AWS. The interface includes a top navigation bar with the Veeam logo, server time (Nov 28, 2023 10:51 AM), user information (administrator, Portal Administrator), and a Configuration icon. 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:

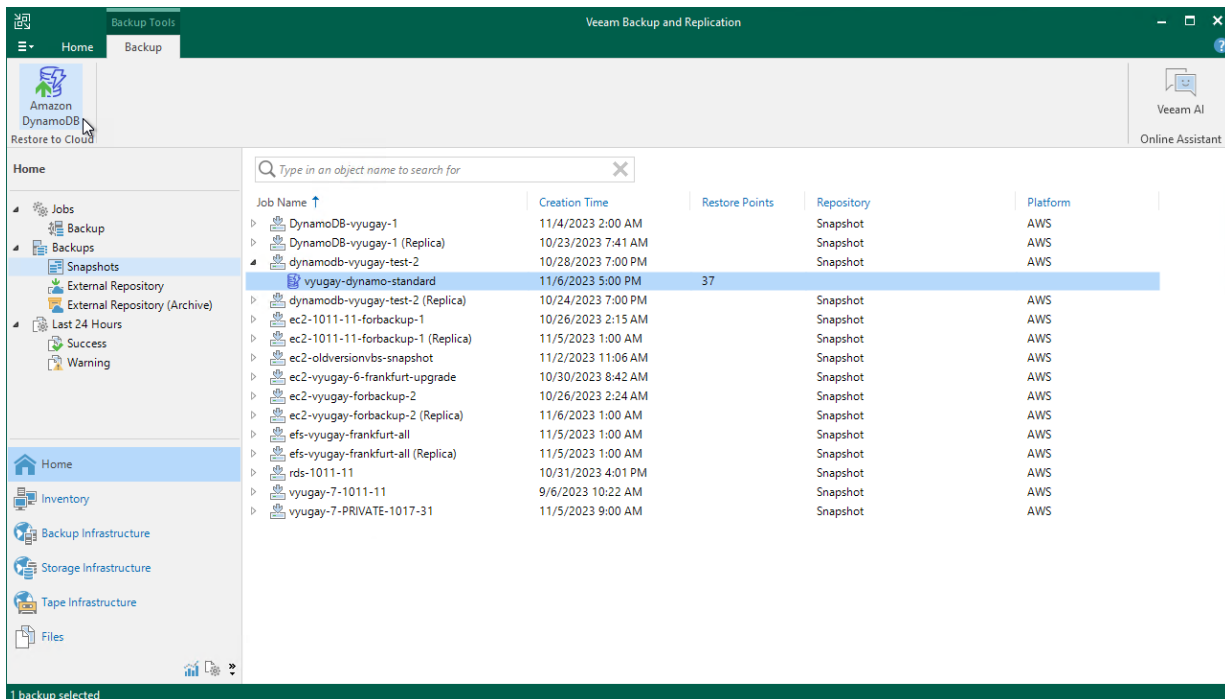
1. In the Veeam Backup & Replication console, open the **Home** view.
2. Navigate to **Backups > Snapshots**.
3. Expand the backup policy that protects the DynamoDB tables that you want to recover, select the necessary table and click **Amazon DynamoDB** on the ribbon.

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:

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

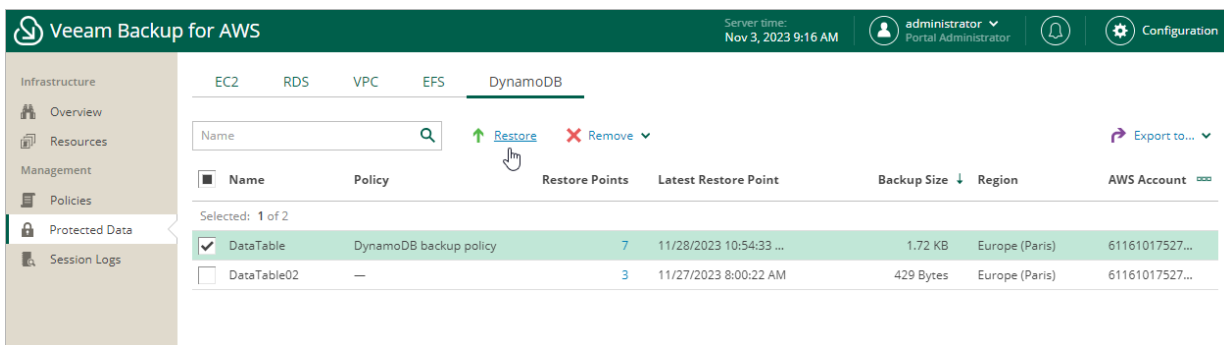
To launch the **DynamoDB Restore** wizard, do the following:

1. Navigate to **Protected Data > 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 for DataTable** 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 bar displays the server time as Nov 3, 2023 9:16 AM and the user as administrator (Portal Administrator). The navigation menu on the left includes Infrastructure, Overview, Resources, Management, Policies, Protected Data, and Session Logs. The main content area is titled 'DynamoDB' and features a search bar, a 'Restore' button, and a 'Remove' button. Below these are two tables of DynamoDB tables and their restore points.

Name	Policy	Restore Points	Latest Restore Point	Backup Size	Region	AWS Account	
Selected: 1 of 2							
<input checked="" type="checkbox"/>	DataTable	DynamoDB backup policy	7	11/28/2023 10:54:33 ...	1.72 KB	Europe (Paris)	61161017527...
<input type="checkbox"/>	DataTable02	—	3	11/27/2023 8:00:22 AM	429 Bytes	Europe (Paris)	61161017527...

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**.

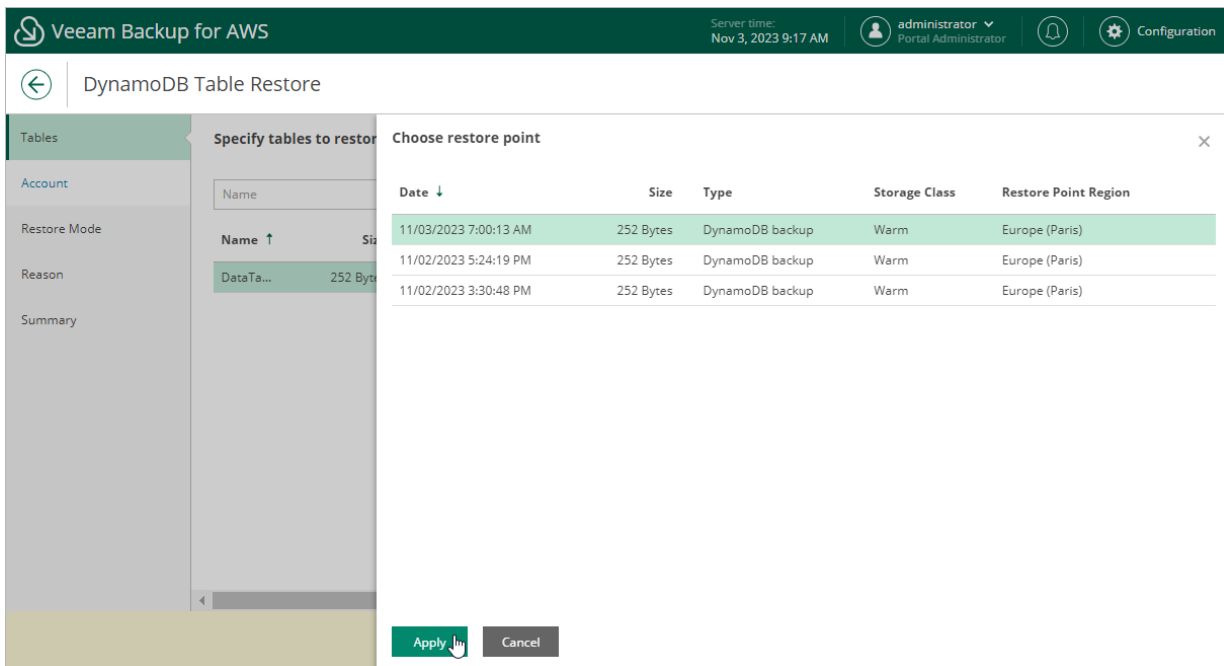
To help you choose a restore point, Veeam Backup for AWS provides the following information on each available restore point:

- **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 another AWS Region. 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 AWS Region in which the selected restore points stored in the cold storage tier are located 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. The top bar displays the Veeam logo, the text "Veeam Backup for AWS", the server time "Nov 3, 2023 9:17 AM", and the user "administrator Portal Administrator". The main window is titled "DynamoDB Table Restore". On the left, there is a sidebar with options: "Tables", "Account", "Restore Mode", "Reason", and "Summary". The "Specify tables to restore" section is active, showing a search box and a table of restore points. The "Choose restore point" dialog is open, displaying a table with the following data:

Date ↓	Size	Type	Storage Class	Restore Point Region
11/03/2023 7:00:13 AM	252 Bytes	DynamoDB backup	Warm	Europe (Paris)
11/02/2023 5:24:19 PM	252 Bytes	DynamoDB backup	Warm	Europe (Paris)
11/02/2023 3:30:48 PM	252 Bytes	DynamoDB backup	Warm	Europe (Paris)

At the bottom of the dialog, there are "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 add an IAM role, click **Add** and complete the [Add IAM Role](#) wizard.

IMPORTANT

It is recommended that you check whether the selected IAM role has all the required permissions 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 (Nov 3, 2023 9:18 AM), user profile (administrator), and configuration icon. A left sidebar contains navigation links for Tables, Account (selected), Restore Mode, Reason, and Summary. 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 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 'Access key:' and 'Secret key:' input fields. A blue information icon is followed by 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 '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 displays the 'DynamoDB Table Restore' configuration interface in Veeam Backup for AWS. The top navigation bar includes the Veeam logo, 'Veeam Backup for AWS', server time 'Nov 3, 2023 9:19 AM', and user information 'administrator Portal Administrator'. A left sidebar contains navigation links for 'Tables', 'Account', 'Restore Mode', 'Reason', and 'Summary'. The main content area is titled 'Choose 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.' Two radio buttons are present: 'IAM role' (unselected) and 'Temporary access keys' (selected). Under 'IAM role', there is a dropdown menu showing 'Default Backup Restore (Default Backup Restore)' and buttons for '+ Add' and 'Check Permissions'. Under 'Temporary access keys', there are two input fields: 'Access key:' with the value 'AKIAY4ZWOU4WMVRAGEVN' and 'Secret key:' with masked characters. An information icon and text 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, a yellow bar contains three buttons: 'Previous' (highlighted), 'Next', and '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 server time is 'Nov 3, 2023 9:20 AM' and the user is 'administrator Portal Administrator'. The wizard has a sidebar with options: Tables, Account, Restore Mode (selected), Encryption, Settings, Capacity, Reason, and Summary. The main 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 a warning 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 buttons 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 'Europe (Milan)'. 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 number (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 'Encryption' step is active, showing two options: 'Use original encryption scheme' (unselected) and 'Change server-side encryption' (selected). Under 'Change server-side encryption', there is a dropdown menu for 'Encryption key' with 'am-key' selected. An information icon and a link to a Veeam KB article are also visible. The bottom of the wizard has 'Previous', 'Next', and 'Cancel' buttons.

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 during the 'DynamoDB Table Restore' process. The top navigation bar includes the Veeam logo, 'Veeam Backup for AWS', server time 'Nov 3, 2023 9:21 AM', and user information 'administrator Portal Administrator'. The main content area is titled 'DynamoDB Table Restore' and features a left-hand navigation menu with options: Tables, Account, Restore Mode, Encryption, Settings (highlighted), Capacity, Reason, and Summary. The 'Configure table settings' panel is active, showing a table with one entry: 'DataTable-2' with a 'Standard' table class. Above this table are 'Edit' and 'Rename' icons. The 'General settings' panel is open, displaying 'Table class' as 'Standard', 'Deletion protection' and 'PITR backup' both enabled with toggle switches, and 'Apply' and 'Cancel' buttons at the bottom.

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 'DynamoDB Table Restore' wizard in the Veeam Backup for AWS interface. The 'Capacity' step is active, showing 'Capacity mode settings' for a table. The 'Capacity mode' is set to 'Provisioned'. A message indicates that auto-scaling is disabled. Below, 'Read capacity' and 'Write capacity' are both set to 30 provisioned capacity units. 'Apply' and 'Cancel' buttons are at the bottom.

Veeam Backup for AWS Server time: Nov 3, 2023 9:22 AM administrator Portal Administrator Configuration

DynamoDB Table Restore

Capacity mode settings Specify which capacity mode to use for the restored table.

Capacity mode: Provisioned

Auto-scaling is disabled. Auto-scaling can be enabled using the AWS management console.

Read capacity Specify the number of read units per second that you require for your application.

Provisioned capacity units: 30

Write capacity Specify the number of write units per second that you require for your application.

Provisioned capacity units: 30

Apply Cancel

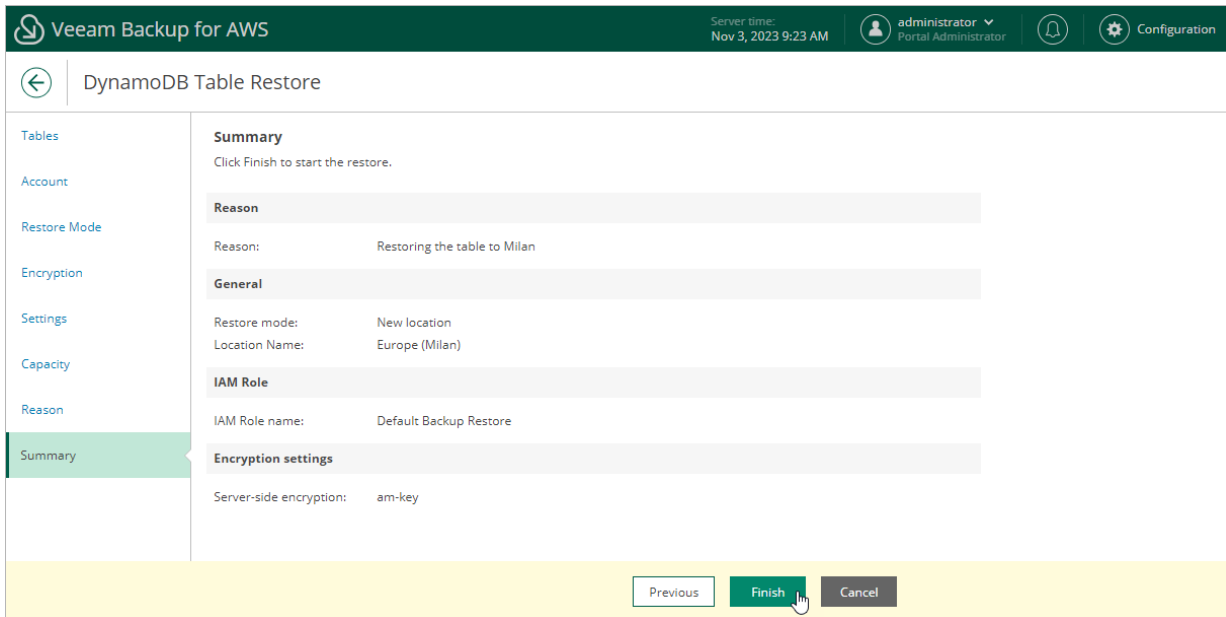
Step 8. Specify Restore Reason

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 'DynamoDB Table Restore' wizard in the Veeam Backup for AWS console. The interface includes a top navigation bar with the Veeam logo, server time (Nov 3, 2023 9:23 AM), user information (administrator, Portal Administrator), and a Configuration icon. A left sidebar lists the wizard steps: Tables, Account, Restore Mode, Encryption, Settings, Capacity, Reason (highlighted), and Summary. The main content area is titled 'Reason' and contains the instruction 'Specify the reason for performing the restore operation.' Below this is a text input field labeled 'Restore reason:' with the text 'Restoring the table to Milan' entered. 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**.



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 where 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 account 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 one-time 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](#).

Restore to Amazon EFS

Restore Mode
Specify whether selected EFS file systems should be restored back to the original location, or to a new location or with different settings.

EFS File System

Restore Mode

Data Center

EFS Configuration

Name

Network

Reason

Summary

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](#)

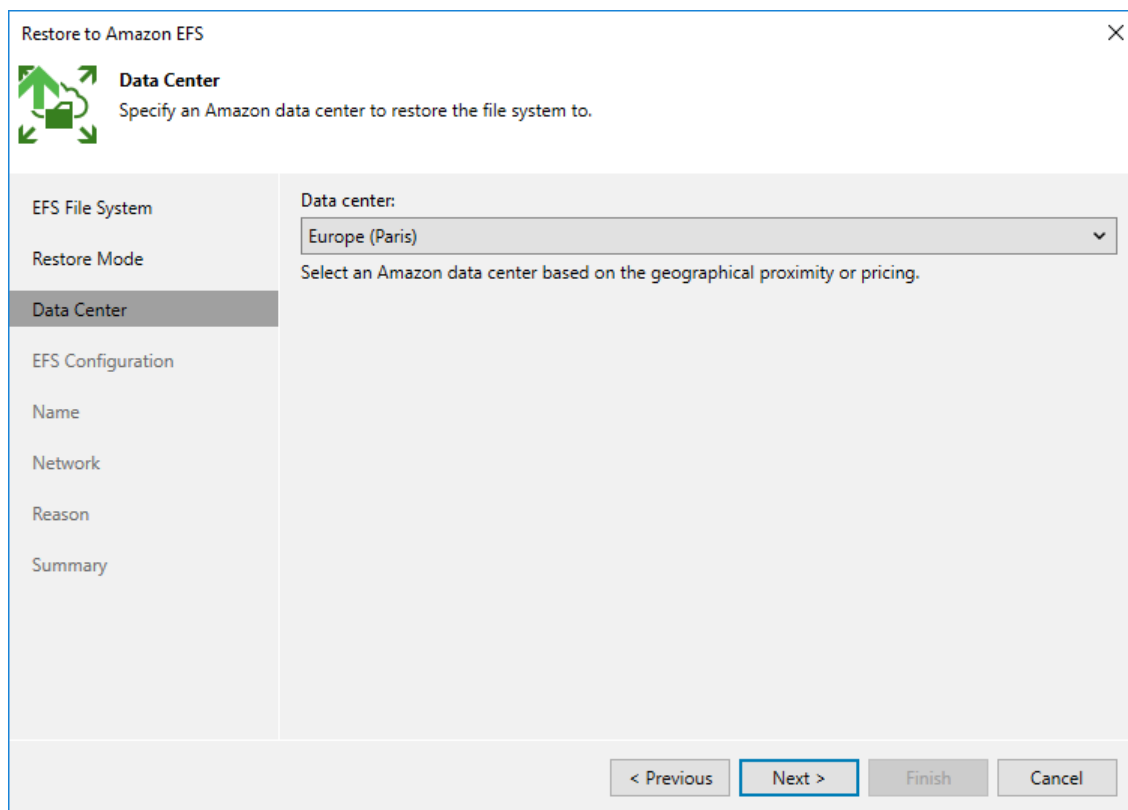
< Previous Next > Finish Cancel

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" with a close button (X) in the top right corner. The main heading is "Data Center" with a green icon of a server rack and arrows, and the instruction "Specify an Amazon data center to restore the file system to." Below this is a list of steps on the left: "EFS File System", "Restore Mode", "Data Center" (highlighted), "EFS Configuration", "Name", "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".

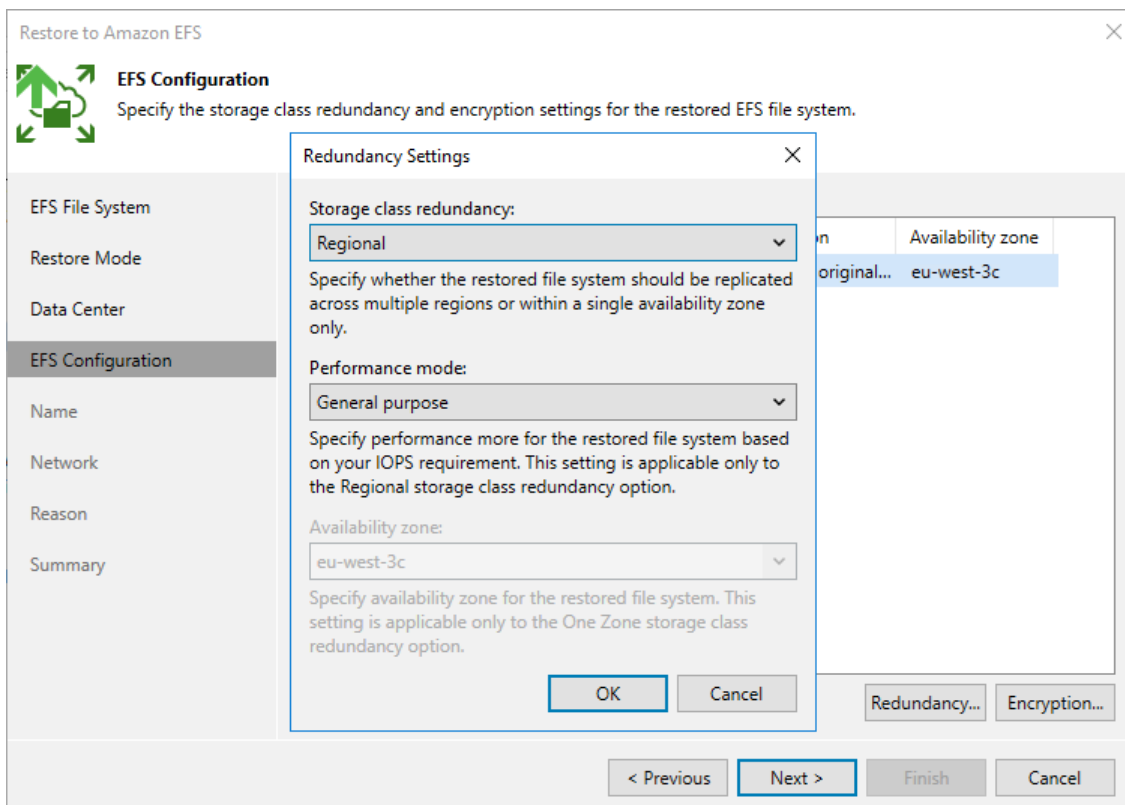
Step 5. Configure Restore Settings

[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](#).



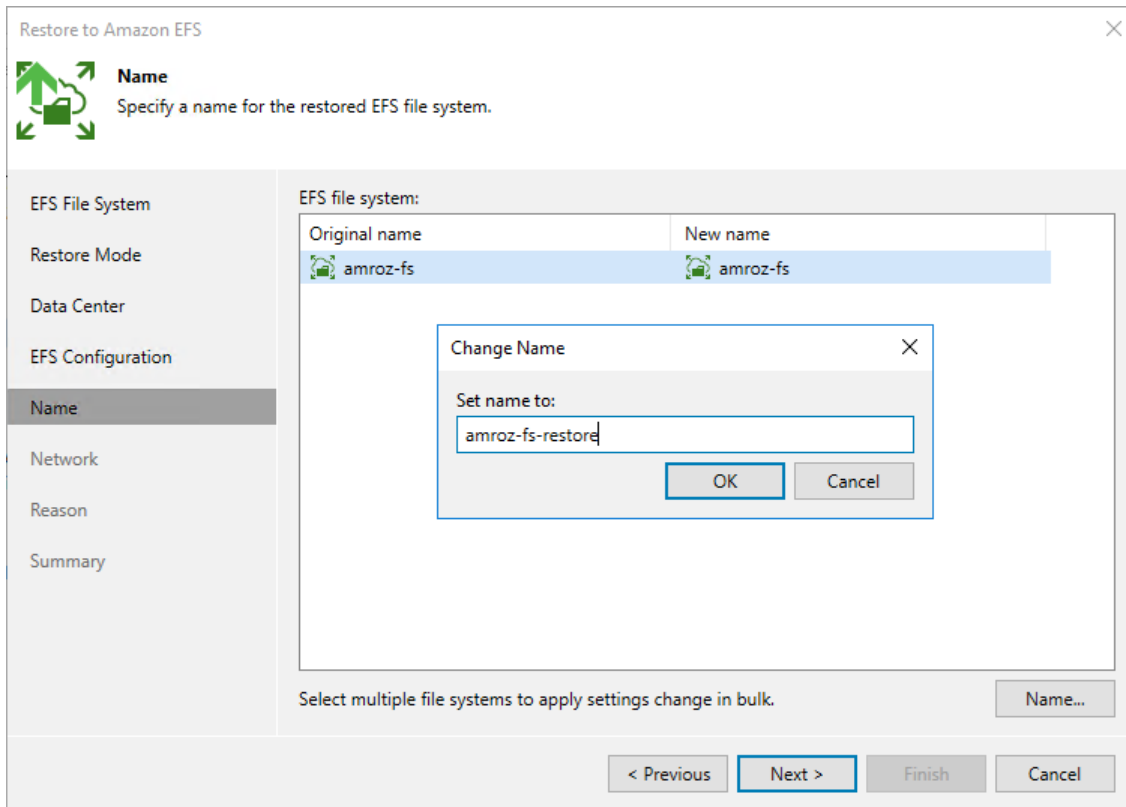
Step 6. Specify File System Name

[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

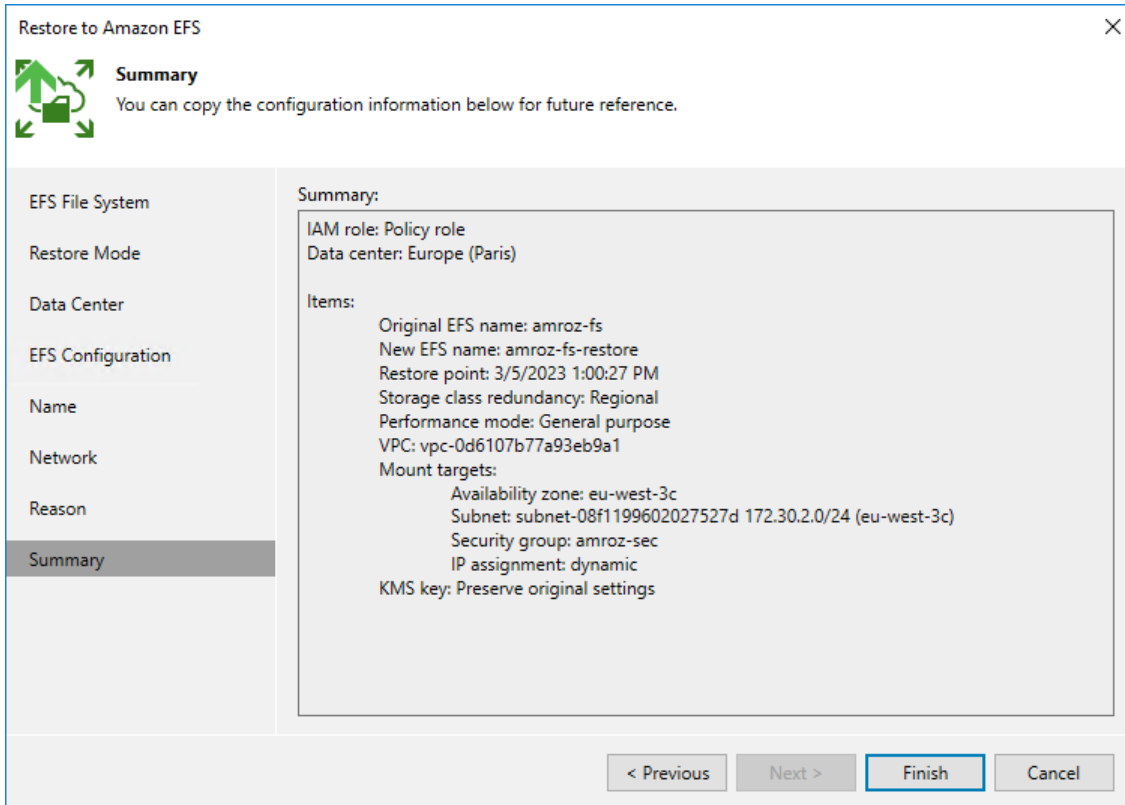
Restore reason:
Restore EFS

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**.



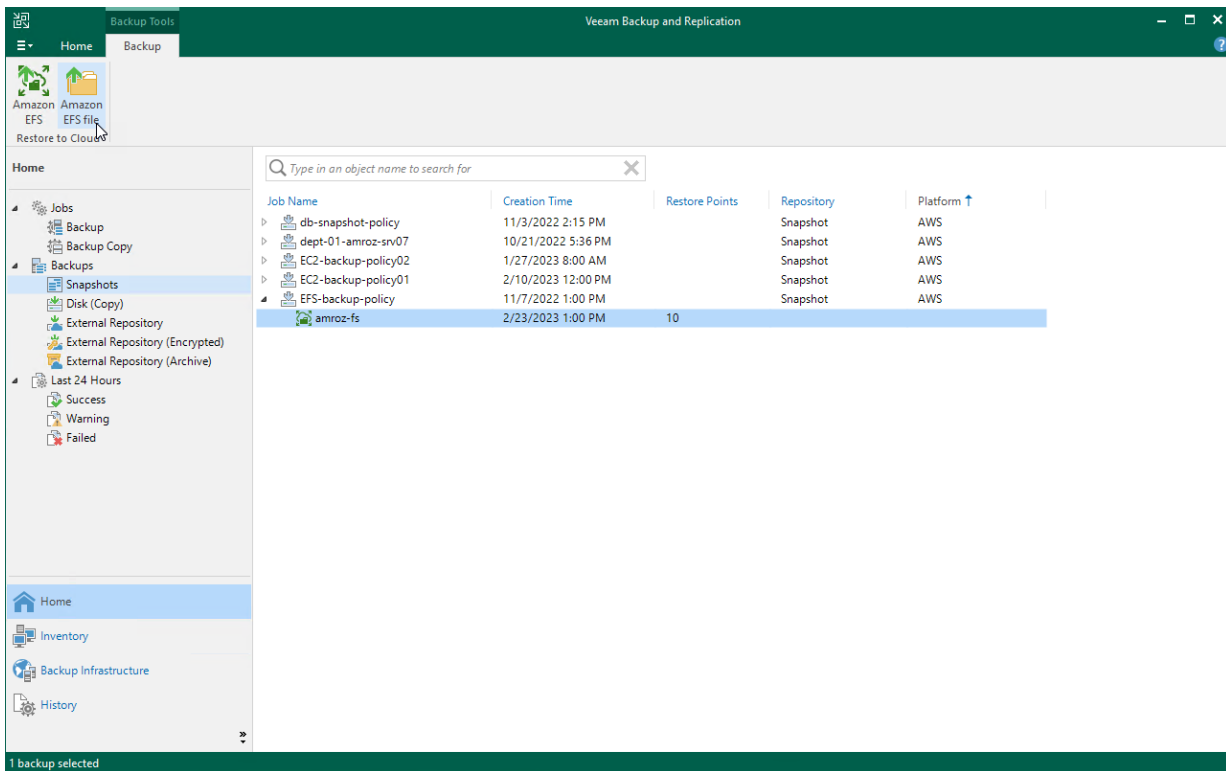
Performing EFS File-Level Restore

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 where 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

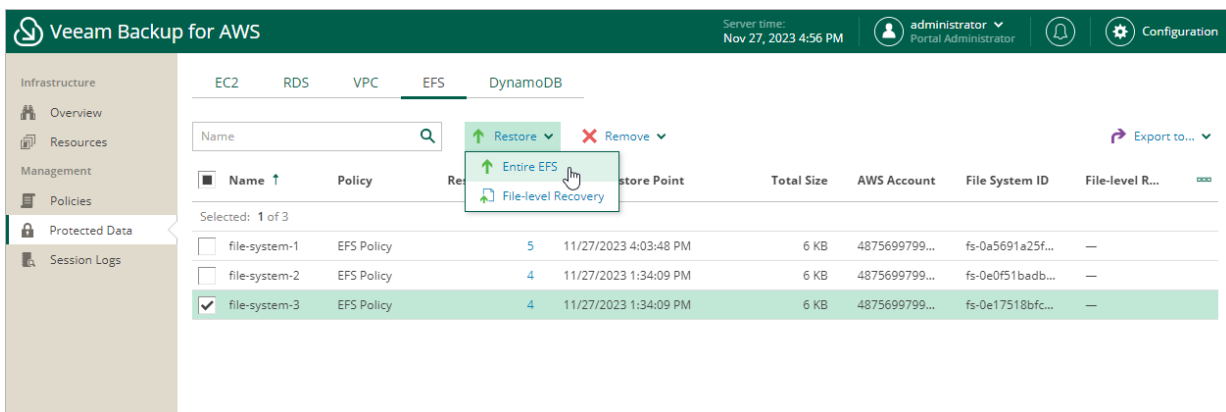
To launch the **EFS Restore** wizard, do the following:

1. Navigate to **Protected Data > 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.



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.

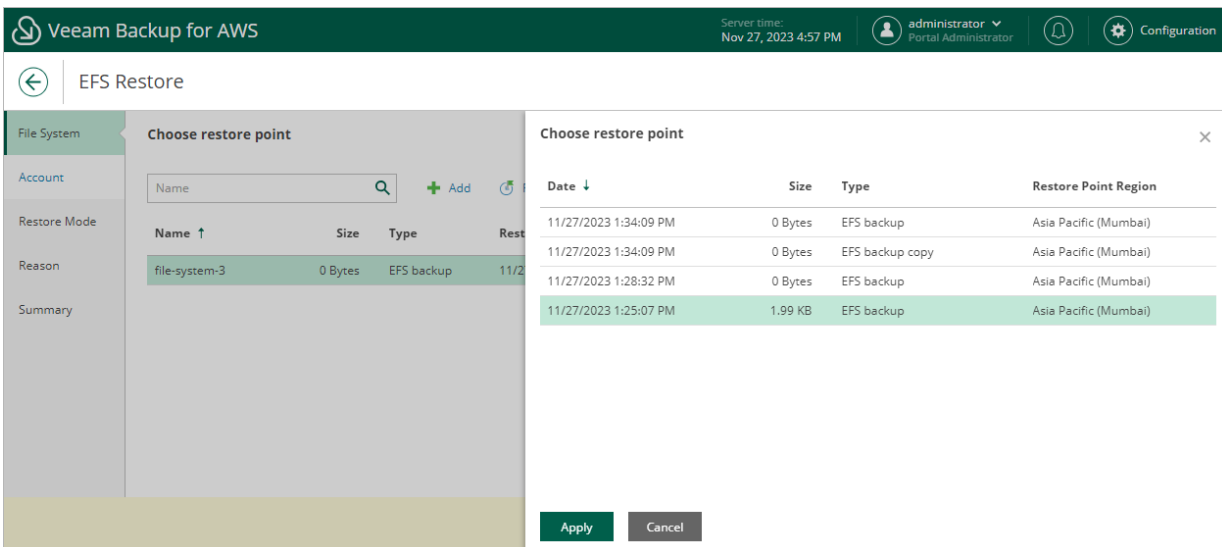
By default, Veeam Backup for AWS uses the most recent valid restore point. However, you can restore a file system to an earlier state.

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 navigation bar includes the Veeam logo, the text 'Veeam Backup for AWS', the server time 'Nov 27, 2023 4:57 PM', and user information 'administrator Portal Administrator'. The main window is titled 'EFS Restore' and shows a 'Choose restore point' dialog box. The dialog box contains a search bar and a table of restore points. The table has the following data:

Date ↓	Size	Type	Restore Point Region
11/27/2023 1:34:09 PM	0 Bytes	EFS backup	Asia Pacific (Mumbai)
11/27/2023 1:34:09 PM	0 Bytes	EFS backup copy	Asia Pacific (Mumbai)
11/27/2023 1:28:32 PM	0 Bytes	EFS backup	Asia Pacific (Mumbai)
11/27/2023 1:25:07 PM	1.99 KB	EFS backup	Asia Pacific (Mumbai)

The last row is highlighted in green. At the bottom of the dialog box, there are '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 [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 add an IAM role, click **Add** and complete the [Add IAM Role](#) wizard.

IMPORTANT

It is recommended that you check whether the selected IAM role has all the required permissions 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 'Restore Mode' is set to 'IAM role'. A list of IAM roles is displayed, including 'Default Backup Restore (Default Backup Restore)' and 'EFS Backup and Restore role (Created by administrator at 11/27/2023 1:25 PM)'. The 'EFS Backup and Restore role' is selected. Below the list, there is a 'Secret key' field and a warning message: '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: Nov 27, 2023 4:58 PM 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: 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 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, server time (Nov 27, 2023 4:59 PM), user information (administrator, Portal Administrator), and a Configuration icon. A left sidebar contains navigation links: File System, 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 file system to the original location or to a new one, or with different settings.' There are two radio button options: '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 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 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 number (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 the Veeam Backup for AWS console. The 'Encryption' step is active, showing two radio button options: 'Use original encryption scheme' (unselected) and 'Restore as encrypted file system' (selected). Below the options is a dropdown menu for 'Encryption key' with 'aws/elasticfilesystem' selected. An information icon and a link to a Veeam KB article are also visible. The bottom of the wizard has 'Previous', 'Next', and 'Cancel' buttons.

Step 6. Configure General Settings

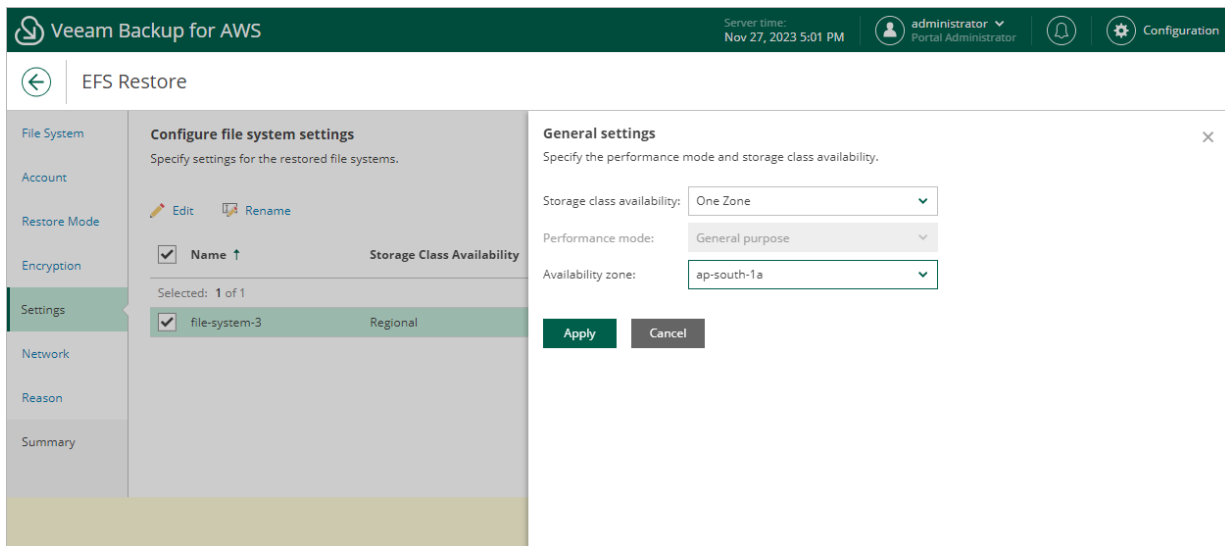
[This step applies only if you have selected the **Restore to new location, or with different settings** option at the **Restore Mode** step of the wizard]

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:

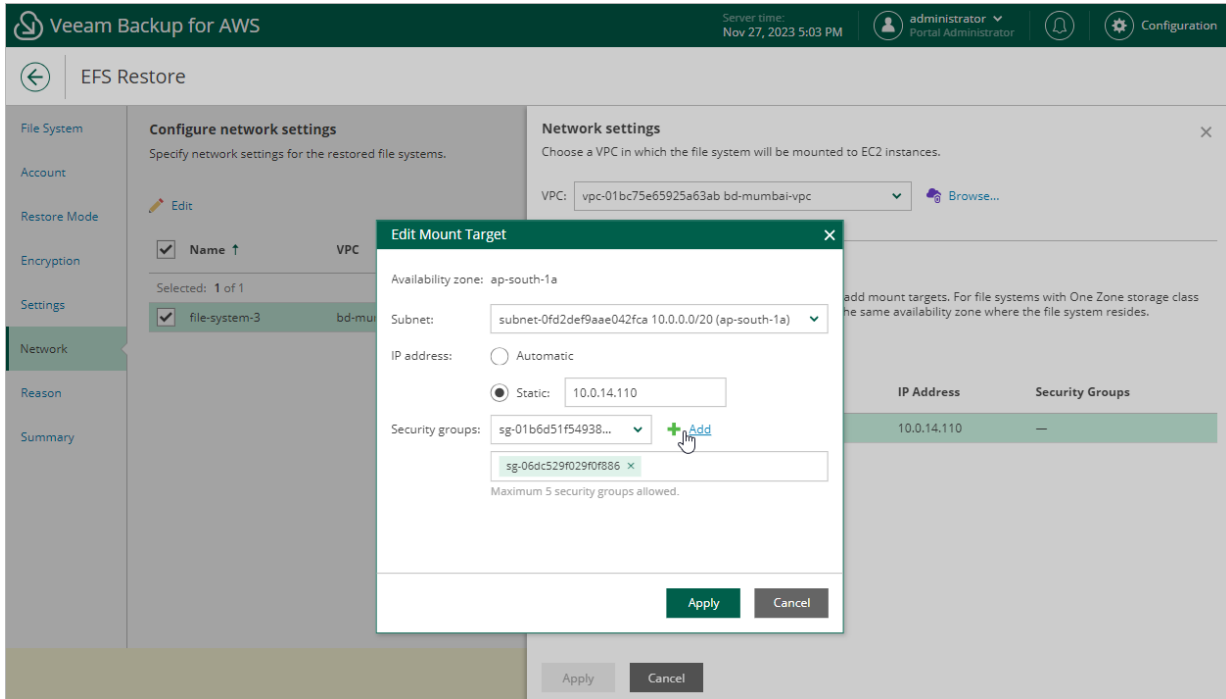
1. In the **Network** section, click **Edit Network Settings**.
2. In the **Network specifications** window, select the necessary Amazon VPC.
For a VPC to be displayed in the **VPC** list, it must be created in the AWS Region specified at [step 4](#) of the wizard as described in [AWS Documentation](#).
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**.



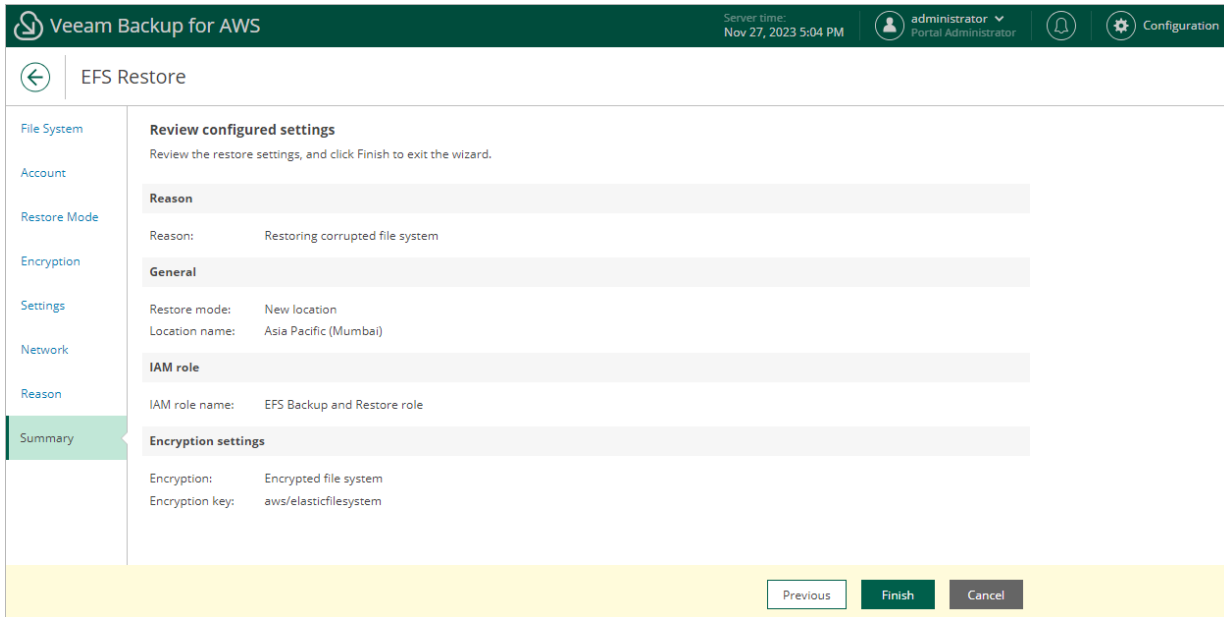
Step 8. Specify Restore Reason

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 "Nov 27, 2023 5:04 PM", 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 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**.



The screenshot shows the 'EFS Restore' wizard in the 'Summary' step. The interface includes a top navigation bar with the Veeam logo, server time (Nov 27, 2023 5:04 PM), user (administrator), and configuration icons. 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:** Restoring corrupted file system
- General:**
 - Restore mode: New location
 - Location name: Asia Pacific (Mumbai)
- IAM role:**
 - IAM role name: EFS Backup and Restore role
- Encryption settings:**
 - Encryption: Encrypted file system
 - Encryption key: aws/elasticfilesystem

At the bottom of the wizard, there are three buttons: 'Previous', 'Finish', and 'Cancel'.

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:

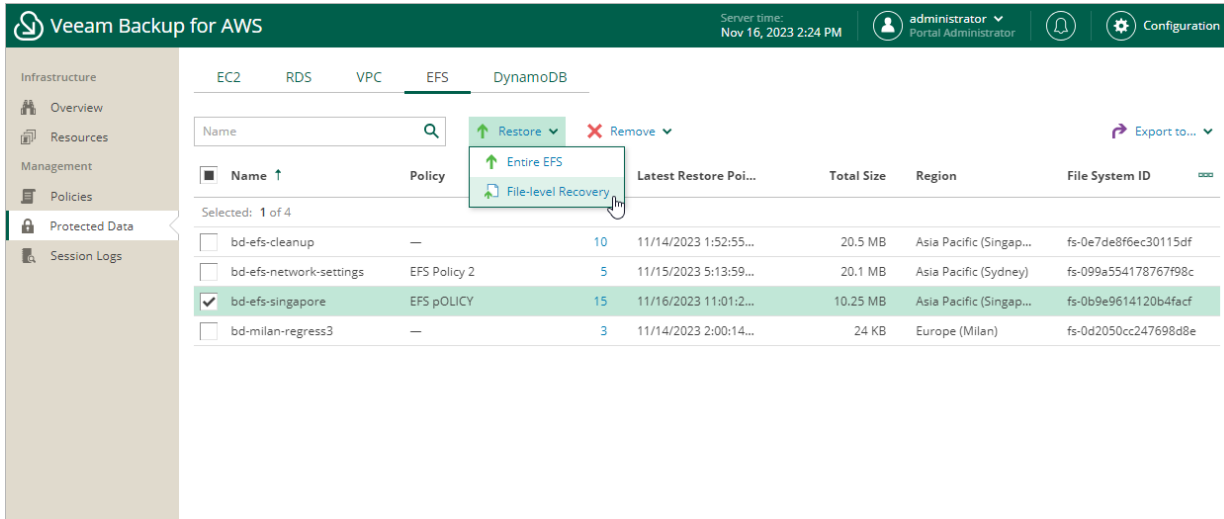
1. [Launch the EFS File-level Recovery wizard.](#)
2. [Choose a restore type.](#)
3. [Configure restore settings.](#)
4. [Specify an IAM identity for restore.](#)
5. [Choose a restore mode.](#)
6. [Specify a restore reason.](#)
7. [Finish working with the wizard.](#)
8. [Open the file-level recovery 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 > 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**.



The screenshot displays the Veeam Backup for AWS interface. The top navigation bar includes the Veeam logo, server time (Nov 16, 2023 2:24 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 for EC2, RDS, VPC, EFS (selected), and DynamoDB. A search bar and 'Export to...' button are visible. Below the search bar, a table lists EFS file systems. The 'bd-efs-singapore' file system is selected, and a context menu is open over it, showing 'Restore' and 'Remove' options. The 'Restore' option is expanded, showing 'Entire EFS' and 'File-level Recovery' sub-options. A mouse cursor is pointing at 'File-level Recovery'.

Name	Policy	Latest Restore Poi...	Total Size	Region	File System ID
bd-efs-cleanup	—	10 11/14/2023 1:52:55...	20.5 MB	Asia Pacific (Singap...	fs-0e7de8f6ec30115df
bd-efs-network-settings	EFS Policy 2	5 11/15/2023 5:13:59...	20.1 MB	Asia Pacific (Sydney)	fs-099a554178767f98c
<input checked="" type="checkbox"/> bd-efs-singapore	EFS pOLICY	15 11/16/2023 11:01:2...	10.25 MB	Asia Pacific (Singap...	fs-0b9e9614120b4facf
bd-milan-regress3	—	3 11/14/2023 2:00:14...	24 KB	Europe (Milan)	fs-0d2050cc247698d8e

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. At the top, the header includes the Veeam logo, the text 'Veeam Backup for AWS', the server time 'Nov 16, 2023 2:24 PM', the user 'administrator Portal Administrator', and a 'Configuration' link. Below the header, the main area is titled 'EFS File-level Recovery'. On the left, a sidebar contains navigation links: 'Restore Type' (highlighted), 'Restore List', 'Account', 'Restore Mode', 'Reason', and 'Summary'. The main content area is titled 'Choose restore type' and contains the instruction: 'Specify whether you want to select files from an index or restore files using the exact file paths.' There are two radio button options: 'Browse files' (unselected) with the subtext 'Browse the data of one or more restore points to find the files that will be restored to the file system.', and 'Specify file paths' (selected) with the subtext 'Specify direct paths to the files that will be restored to the file system.'. At the bottom of the main area, there are two buttons: 'Next' and 'Cancel'.

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 'Nov 16, 2023 2:25 PM', user 'administrator Portal Administrator', and a 'Configuration' link. The main content area is titled 'EFS File-level Recovery'. On the left, there is a sidebar with 'Restore Type' selected. The main area is divided into two panels: 'Choose restore point and items to restore' and 'Choose restore point'. The 'Choose restore point' panel contains a table with the following data:

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. At the top, the header shows 'Veeam Backup for AWS' on the left, 'Server time: Nov 16, 2023 2:26 PM' in the center, and user information 'administrator Portal Administrator' on the right. Below the header, the main area is titled 'EFS File-level Recovery'. On the left, a sidebar contains navigation options: 'Restore Type', 'Restore List' (selected), 'Account', 'Restore Mode', 'Reason', and 'Summary'. The 'Restore List' section is active, showing 'Choose restore point and items to restore'. Under 'Restore point', it says 'Choose a restore point to use for creating the restore list.' and 'Restore point: 11/15/2023 4:35:15 PM'. Under 'Restore list', it says 'Items' and 'No items added yet' with an 'Edit' button. The 'Edit restore list' window is open on the right, showing 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.' Below the instructions, there is an 'Item path' field containing '/ortiz/accounts' and an '+ Add' button. There is also a 'Remove from Restore List' button. Below that, there is a checkbox for 'Item' and a 'Selected: 0 of 1' indicator. At the bottom of the window, there is a list of items with a checkbox for '/ortiz/balance.docx'. At the very bottom of the interface, there 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 add an IAM role, click **Add** and complete the [Add IAM Role](#) wizard.

IMPORTANT

It is recommended that you check whether the selected IAM role has all the required permissions 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 two options: 'Default Backup Restore (Default Backup Restore)' and 'EFS restore role (Created by bd-regress-2 at 11/14/2023 6:45 PM)'. The second option 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 is divided into a left sidebar with navigation options and a main content area. The sidebar includes 'Restore Type', 'Restore List', 'Account', 'Restore Mode' (highlighted), 'Reason', and 'Summary'. The main content area is titled 'Choose restore mode' and contains the following text: 'Specify whether you want to restore files to the original location or to a new one, or with different settings.' There are two radio button options: 'Restore to original location' (unselected) and 'Restore to new location, or with different settings' (selected). Below the selected option, it says 'Perform additional configuration steps to restore the selected files to a new file system.' 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'.

Veeam Backup for AWS

Server time: Nov 16, 2023 2:28 PM

administrator Portal Administrator

Configuration

EFS File-level Recovery

Restore Type

Restore List

Account

Restore Mode

Reason

Summary

Choose restore mode

Specify whether you want to restore files to the original location or to a new one, or with different settings.

Restore to original location
Quickly restore the selected files to the original location.

Restore to new location, or with different settings
Perform additional configuration steps to restore the selected files to a new file system.

Region: Asia Pacific (Singapore)

File system: bd-efs-singapore

Previous Next Cancel

Step 6. Specify Restore Reason

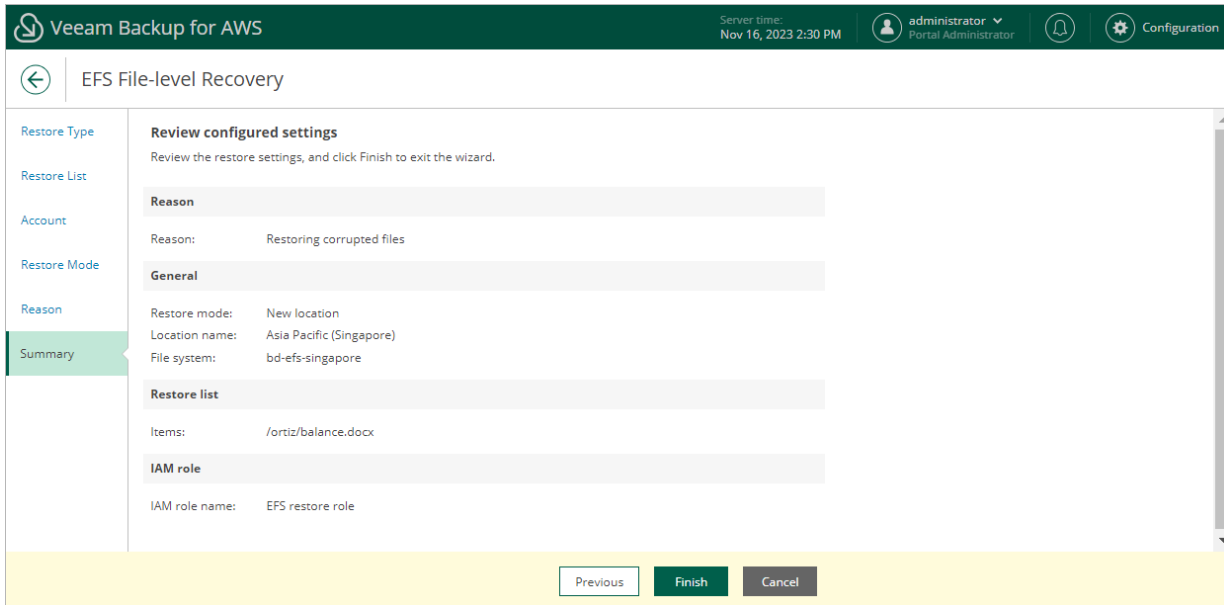
At the **Reason** step of the wizard, you can specify a reason for restoring the files and folders. The information you provide will be saved in the session history and you can reference it later.

The screenshot shows the 'EFS File-level Recovery' wizard in Veeam Backup for AWS. The interface includes a top navigation bar with the Veeam logo, 'Veeam Backup for AWS', server time 'Nov 16, 2023 2:29 PM', user 'administrator Portal Administrator', and a 'Configuration' link. A left sidebar contains navigation options: 'Restore Type', 'Restore List', 'Account', 'Restore Mode', 'Reason' (highlighted), and 'Summary'. The main area is titled 'Restore reason' and contains the instruction 'Specify a reason for performing the restore operation.' Below this is a text input field with the text 'Restoring corrupted files'. At the bottom, 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**.

[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 7-9.



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 > 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:

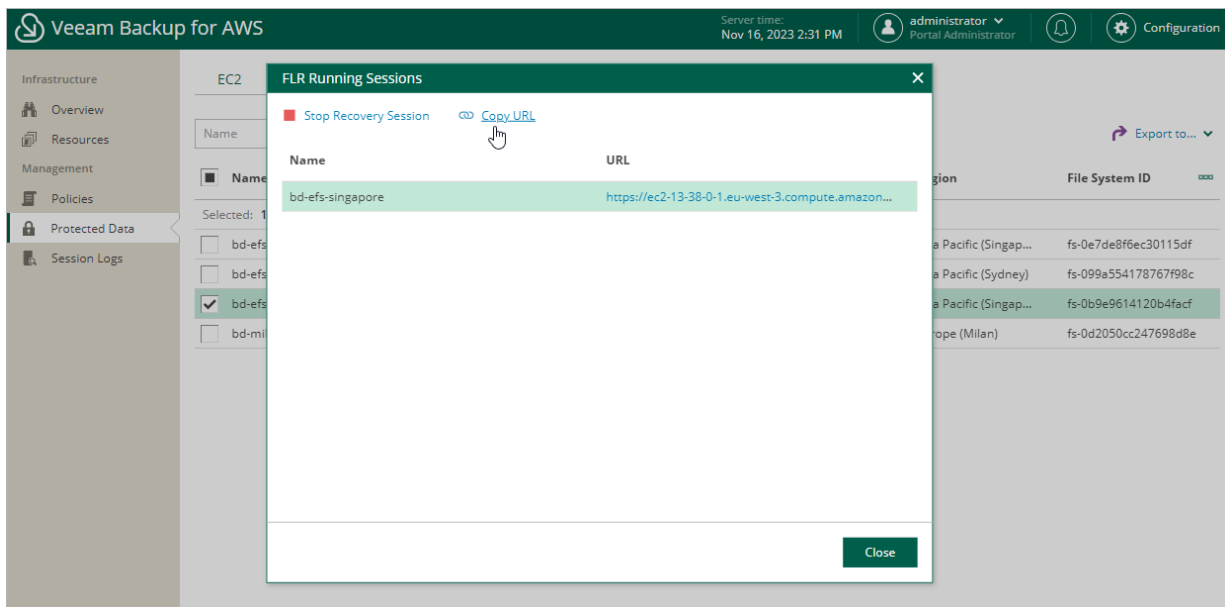
- Click the link to open the file-level recovery browser on your local machine while the recovery session is running.
- Copy the link, close the **FLR Running Sessions** window and open the file-level recovery browser on another machine.

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.

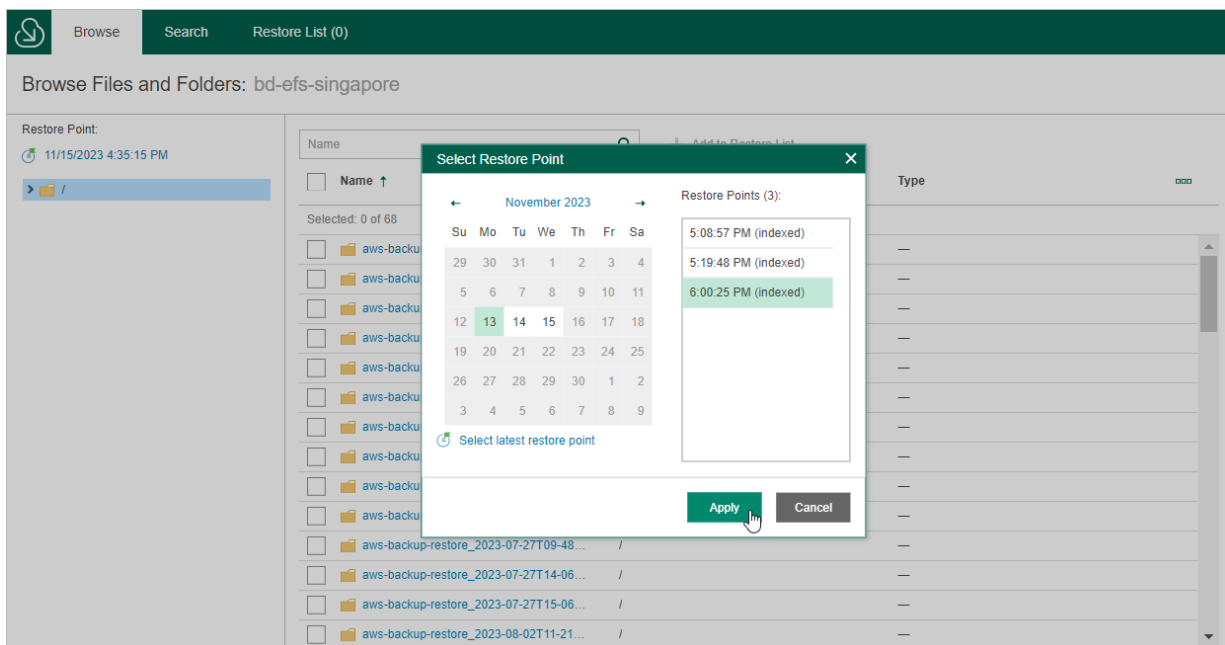
To select a restore point in the file-level recovery browser, do the following:

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 and folders.
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

Restore Remove

<input checked="" type="checkbox"/>	Name ↑	Location	Type	Restore Point	Restore Date	Restore Status	☰
Selected: All 2 items							
<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	—	—	

Session Log

Status: All

Action	Status	Start Time	End Time	Duration	☰
Select a single item to view sessions details					

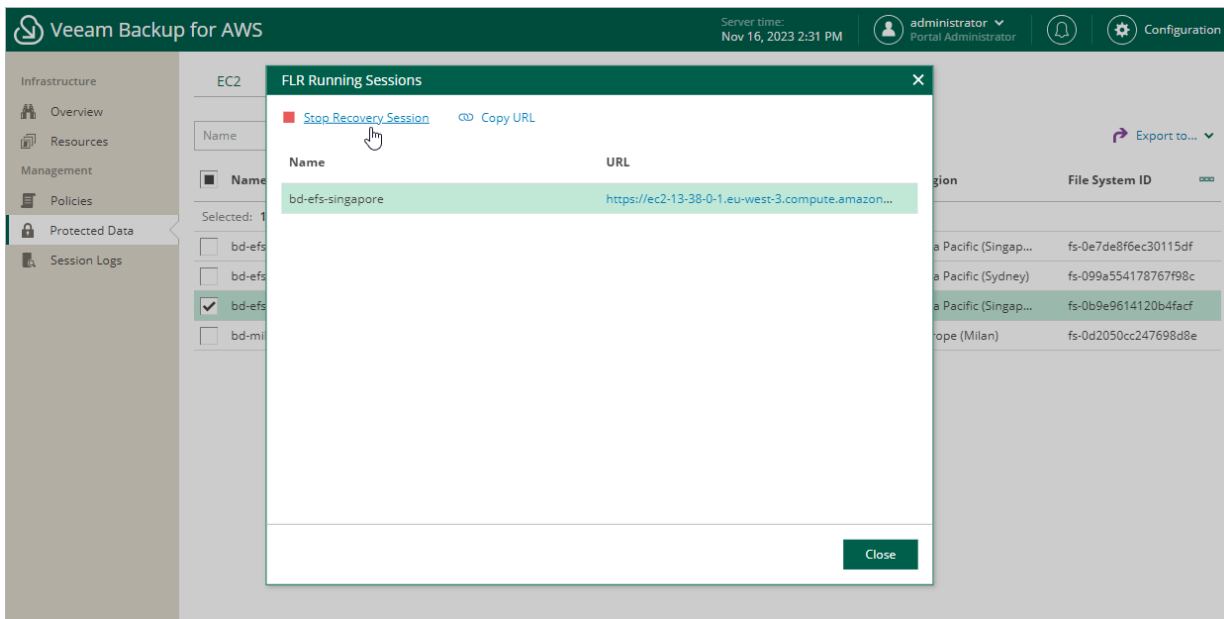
Step 11. Stop Recovery Session

[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 > EFS** and click the link in the **File-Level Recovery URL** column to open the window again.



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

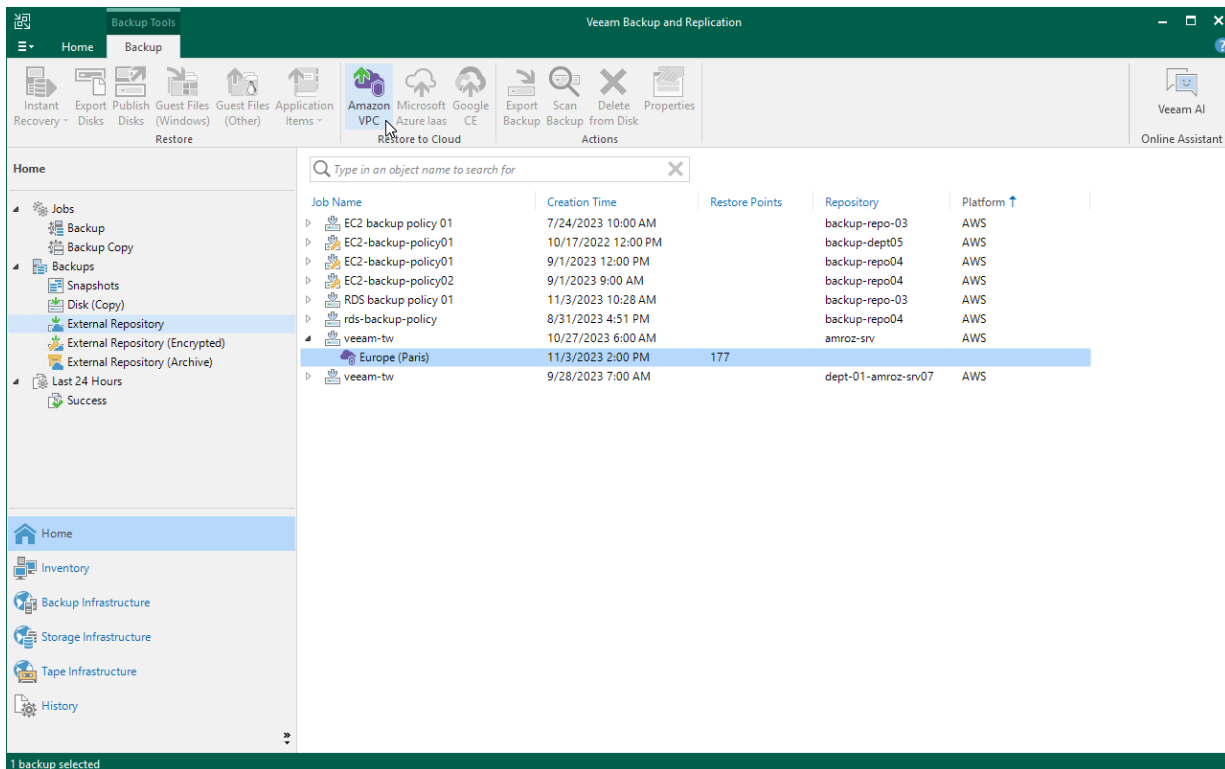
IMPORTANT

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](#).

Veeam Backup & Replication allows you to restore an entire Amazon VPC configuration from a VPC configuration backup to any available restore point. To learn how entire VPC configuration restore works, see [Entire VPC Configuration Restore](#).

To restore a VPC configuration, do the following:

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 to want restore and click **Amazon VPC** on the ribbon.
4. Complete the **VPC Restore** wizard as described in section [VPC Configuration Restore](#).



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

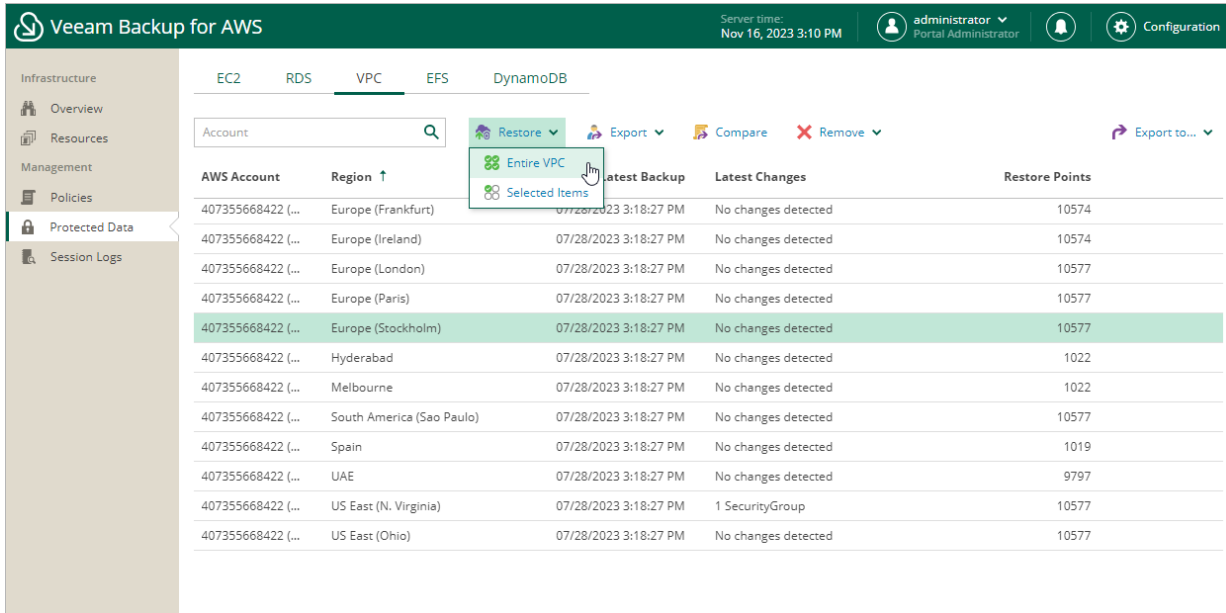
To restore the entire VPC configuration, do the following:

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**.

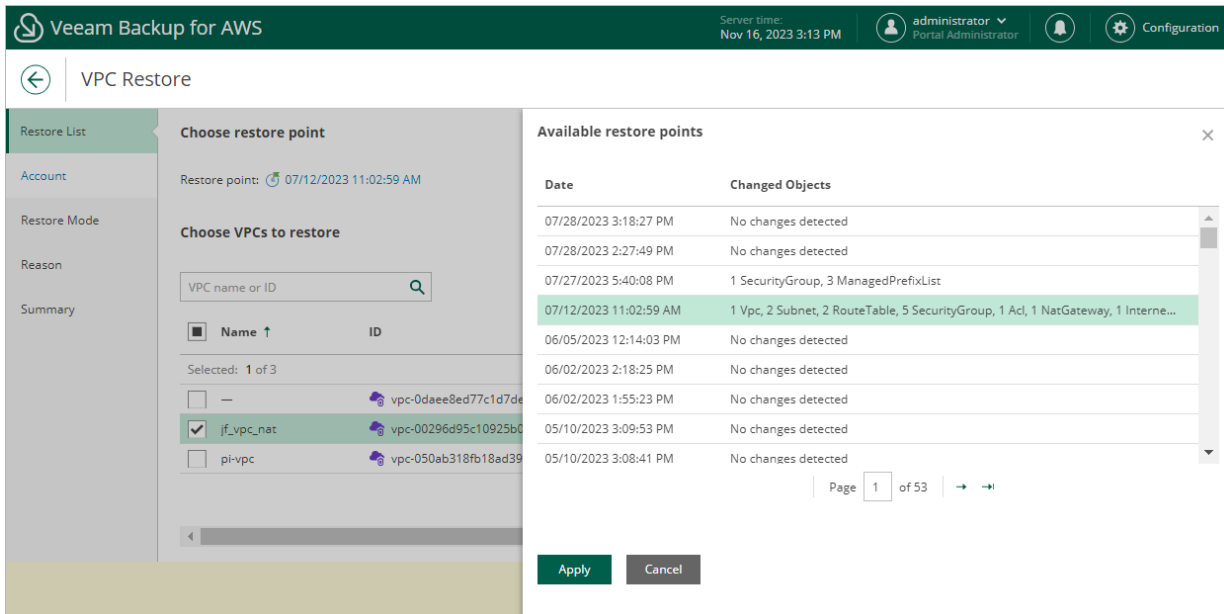


Step 2. Select Restore Point

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, the text 'Veeam Backup for AWS', the server time 'Nov 16, 2023 3:13 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 three VPCs, one of which 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

Page 1 of 53

Apply 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 [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 add an IAM role, click **Add** and complete the [Add IAM Role](#) wizard.

IMPORTANT

It is recommended that you check whether the selected IAM role has all the required permissions 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. The top navigation bar includes the Veeam logo, the text 'Veeam Backup for AWS', the server time 'Nov 16, 2023 3:15 PM', and the user 'administrator Portal Administrator'. The main content area is titled 'VPC Restore' and has a sidebar with options: 'Restore List', 'Account' (selected), 'Restore Mode', 'Reason', and 'Summary'. The 'Account' section is titled 'Select IAM role' and 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', there is a dropdown menu showing 'acc_5393', a '+ Add' button, and a 'Check Permissions' button. Under 'Temporary access keys', there are two text input fields labeled 'Access key:' and 'Secret key:'. An information icon (i) is present with the text: '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 series of dots. 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 for the 'VPC Restore' wizard. The top navigation bar includes the Veeam logo, 'Veeam Backup for AWS', server time 'Nov 16, 2023 3:17 PM', user 'administrator Portal Administrator', and a 'Configuration' icon. The main content area is titled 'VPC Restore' and features a left-hand navigation pane with options: Restore List, Account, Restore Mode (highlighted), Availability Zones, Peering Connection, Reason, and Summary. The 'Restore Mode' section contains the following text and options:

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.

Below the second option, there 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

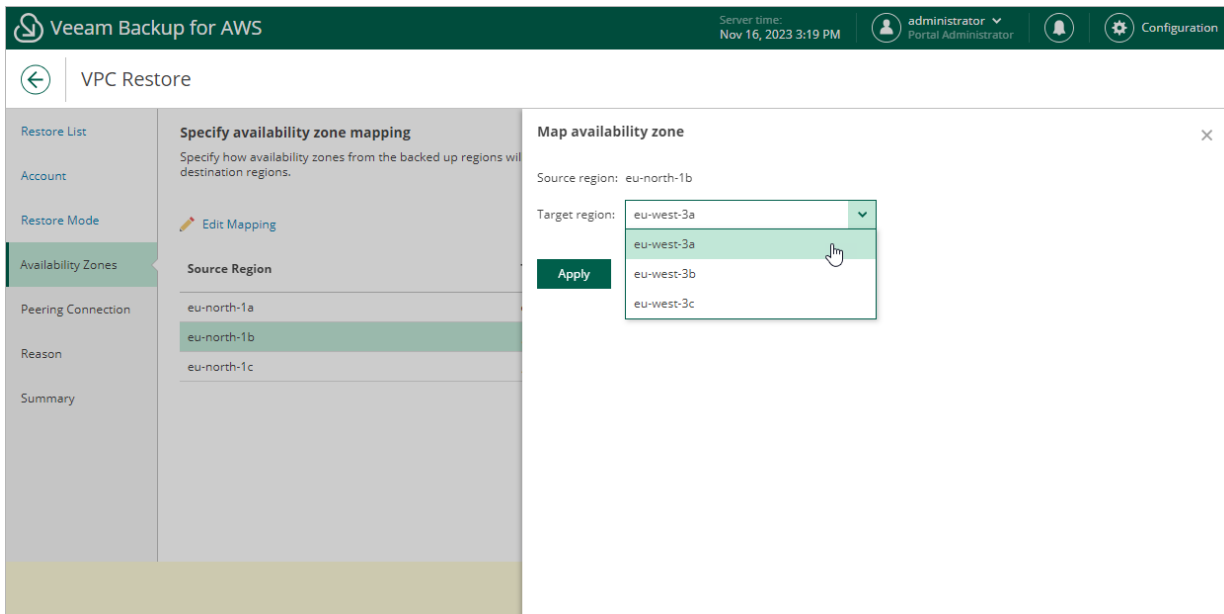
[This step applies only if you have selected the **Restore to new location, or with different settings** option at the **Restore Mode** step of the wizard]

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'.

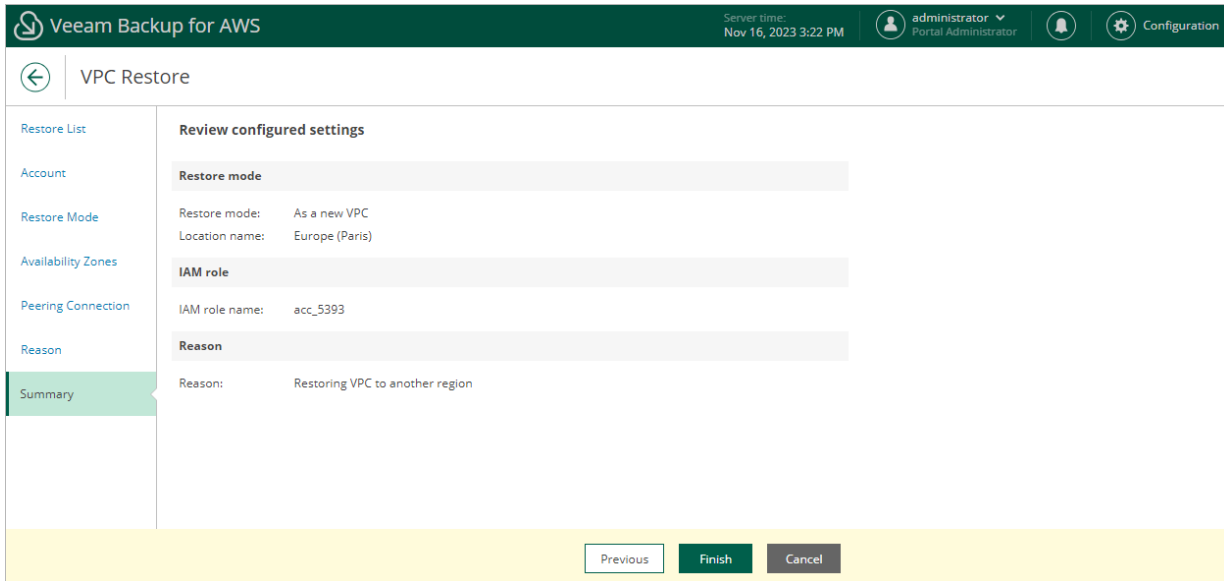
Step 7. Specify Restore Reason

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 "Nov 16, 2023 3:22 PM", the user "administrator Portal Administrator", and a "Configuration" link. 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 (Nov 16, 2023 3:22 PM), and user information (administrator, Portal Administrator). 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

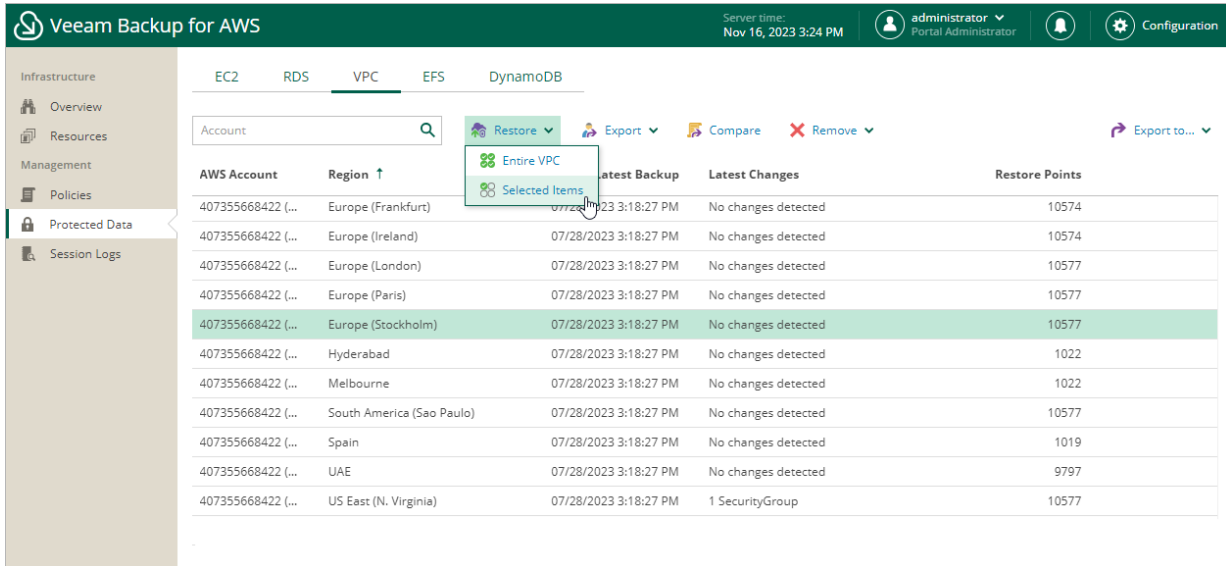
To restore specific items of the VPC configuration, do the following:

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**.



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 shows the Veeam Backup for AWS interface. The main window is titled 'VPC Restore'. On the left, there is a sidebar with 'Restore List' selected. The main content area is divided into two sections: 'Choose restore point' and 'Create restore list'. The 'Create restore list' section is currently active, showing a table of VPC resources. The table has columns for Name, ID, Type, and State. Two VPCs are listed: 'pi-vpc' and 'jf_vpc_nat', both with a state of 'Created'. The 'Apply' button is highlighted.

Name	ID	Type	State
pi-vpc	vpc-050ab318fb18ad392	VPC	Created
jf_vpc_nat	vpc-00296d95c10925b0f	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.

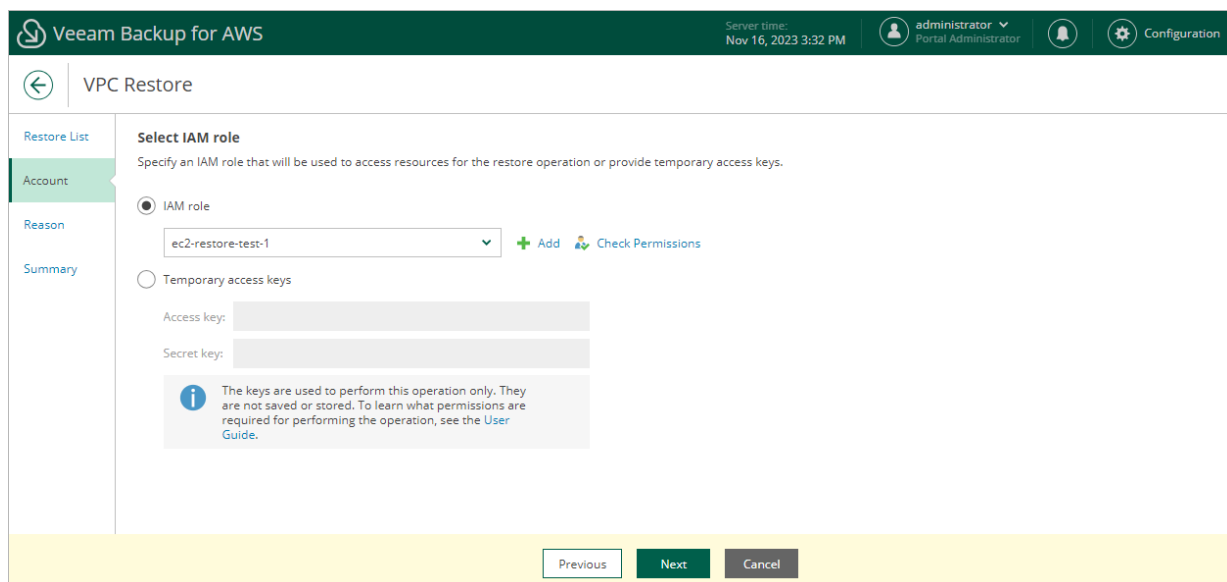
Specifying IAM Role

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 add an IAM role, click **Add** and complete the [Add IAM Role](#) wizard.

IMPORTANT

It is recommended that you check whether the selected IAM role has all the required permissions 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, product name, server time (Nov 16, 2023 3:32 PM), and user information (administrator, Portal Administrator). The main content area is titled "VPC Restore" and has a sidebar with "Restore List", "Account" (selected), "Reason", and "Summary". The "Select 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", there is a dropdown menu showing "ec2-restore-test-1" and buttons for "+ Add" and "Check Permissions". Under "Temporary access keys", there are input fields for "Access key:" and "Secret key:". A blue information icon with a text box states: "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, 'Veeam Backup for AWS', the server time 'Nov 16, 2023 3:33 PM', and the user 'administrator Portal Administrator'. A left sidebar contains navigation options: 'Restore List', 'Account' (highlighted), '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', a dropdown menu shows 'ec2-restore-test-1' with '+ Add' and 'Check Permissions' icons. Under 'Temporary access keys', there are input fields for 'Access key:' (containing 'AKIAY4ZWOU4WMVRAGEVN') and 'Secret key:' (masked with dots). 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 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. At the top, the header includes the Veeam logo, the text "Veeam Backup for AWS", the server time "Nov 16, 2023 3:34 PM", the user "administrator" (Portal Administrator), and a "Configuration" link. Below the header, the main area is titled "VPC Restore". On the left, a navigation pane lists "Restore List", "Account", "Reason" (which is highlighted in green), and "Summary". The main content area is titled "Restore reason" and contains the instruction "Specify a reason for performing the restore operation." Below this, there is a text input field with the text "Restoring VPCs and security groups". At the bottom of the wizard, there are three buttons: "Previous", "Next" (which is 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 Veeam Backup for AWS interface. The top navigation bar includes the Veeam logo, the text "Veeam Backup for AWS", the server time "Nov 16, 2023 3:35 PM", the user "administrator Portal Administrator", and a "Configuration" link. The main content area is titled "VPC Restore" and shows 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" (Original location), "IAM role" (IAM role name: test-rds-restore-role), and "Reason" (Restoring VPCs and security groups). At the bottom, there are three buttons: "Previous", "Finish", 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

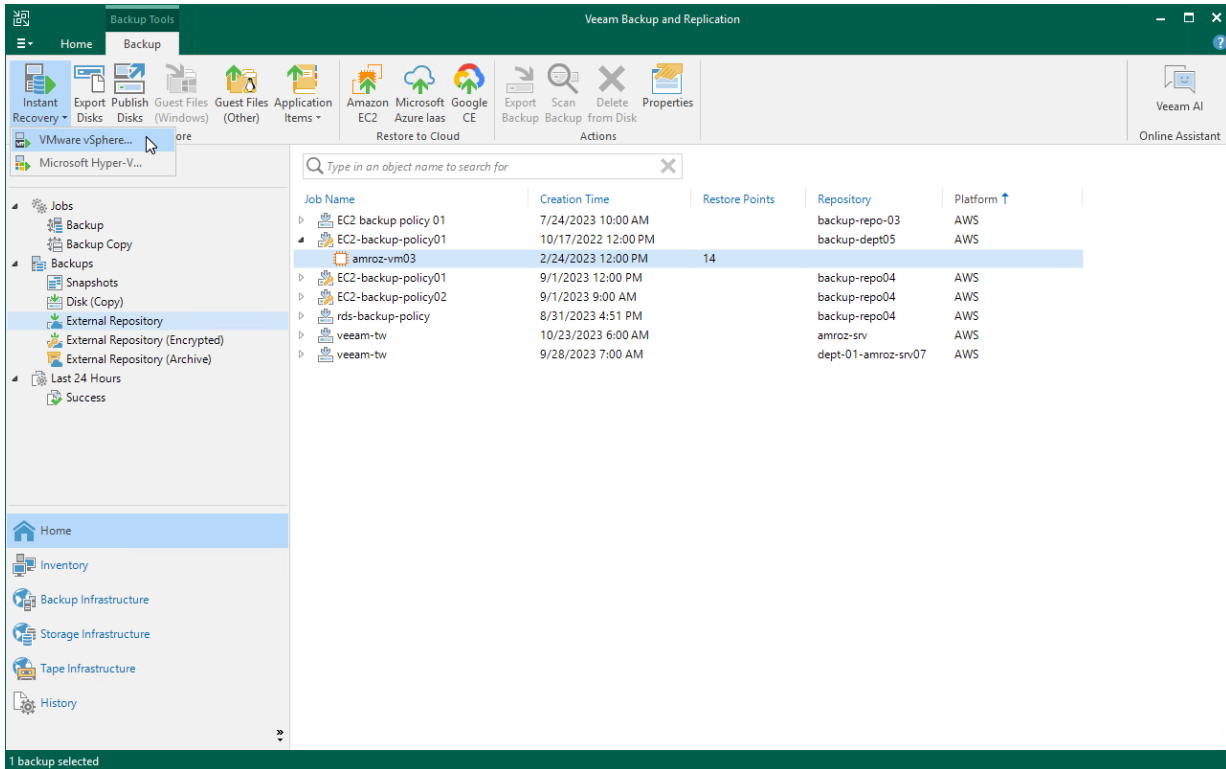
Instant Recovery can be performed only using backup files stored in standard backup repositories for which you have specified access keys of an IAM user whose permissions are used to access the repositories. To learn how to specify credentials for the repositories, see sections [Creating New Repositories](#) 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

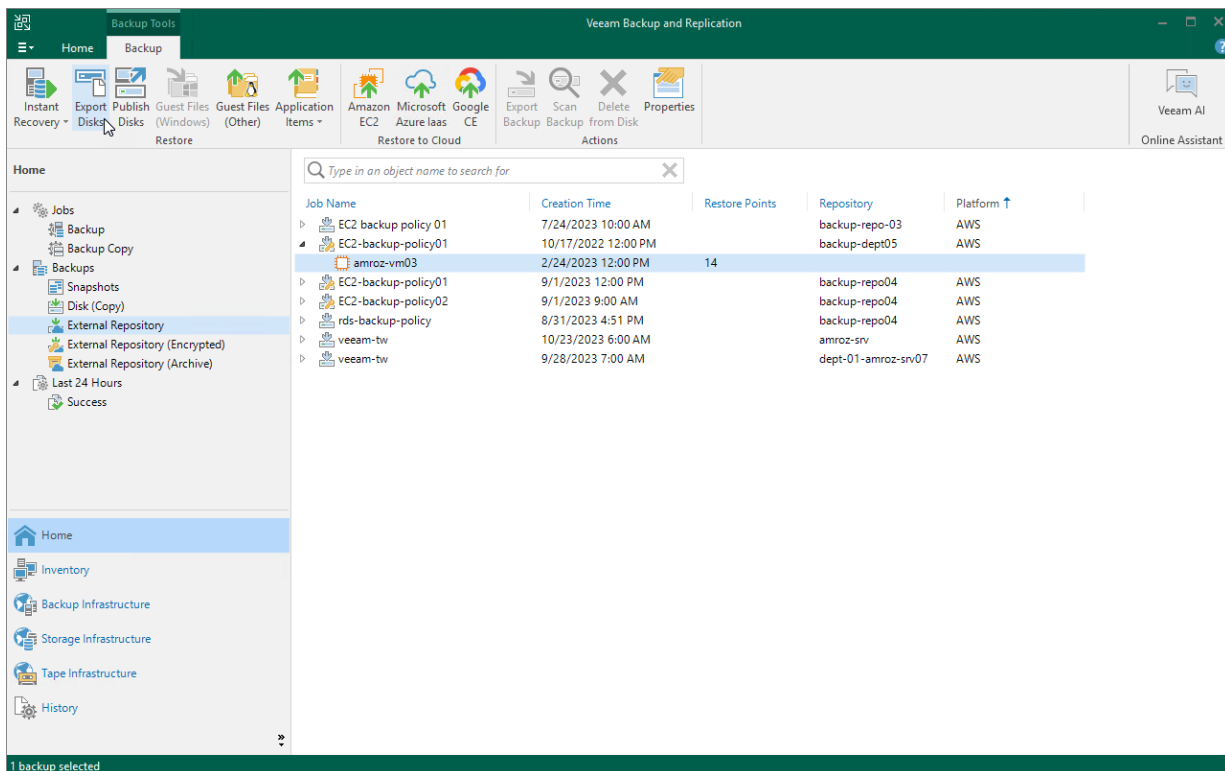
Veeam Backup & Replication allows you to export disks, that is, to restore EBS volumes of EC2 instances from image-level backups created by Veeam Backup for AWS and to convert them to the VMDK, VHD and VHDX formats. You can save the converted disks to any server added to the backup infrastructure or 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

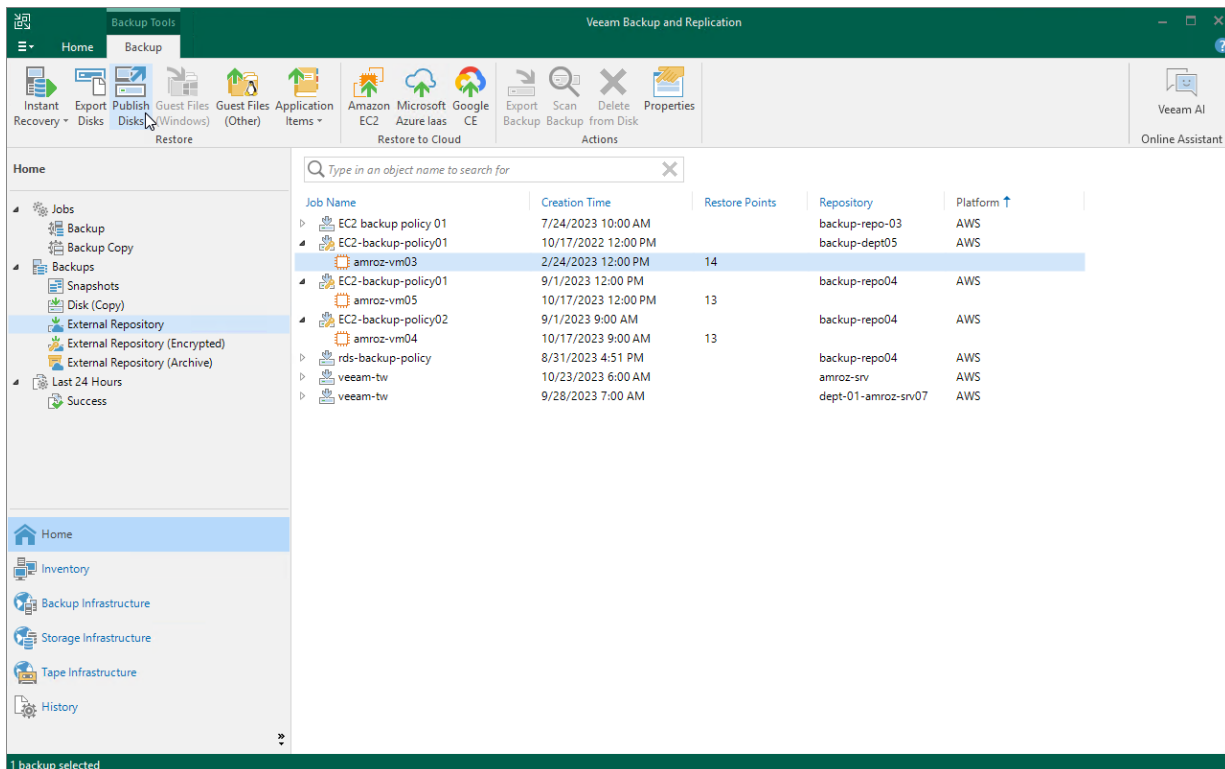
Veeam Backup & Replication allows you to publish point-in-time disks, that is, to mount specific EBS volumes of backed-up EC2 instances to any server to instantly access data in the read-only mode. You can copy the necessary files and folders to the target server, and perform an antivirus scan of the 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

Consider the following:


- 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, check the limitations and prerequisites described in the Veeam Backup & Replication User Guide, section [Before You Begin](#).

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](#).

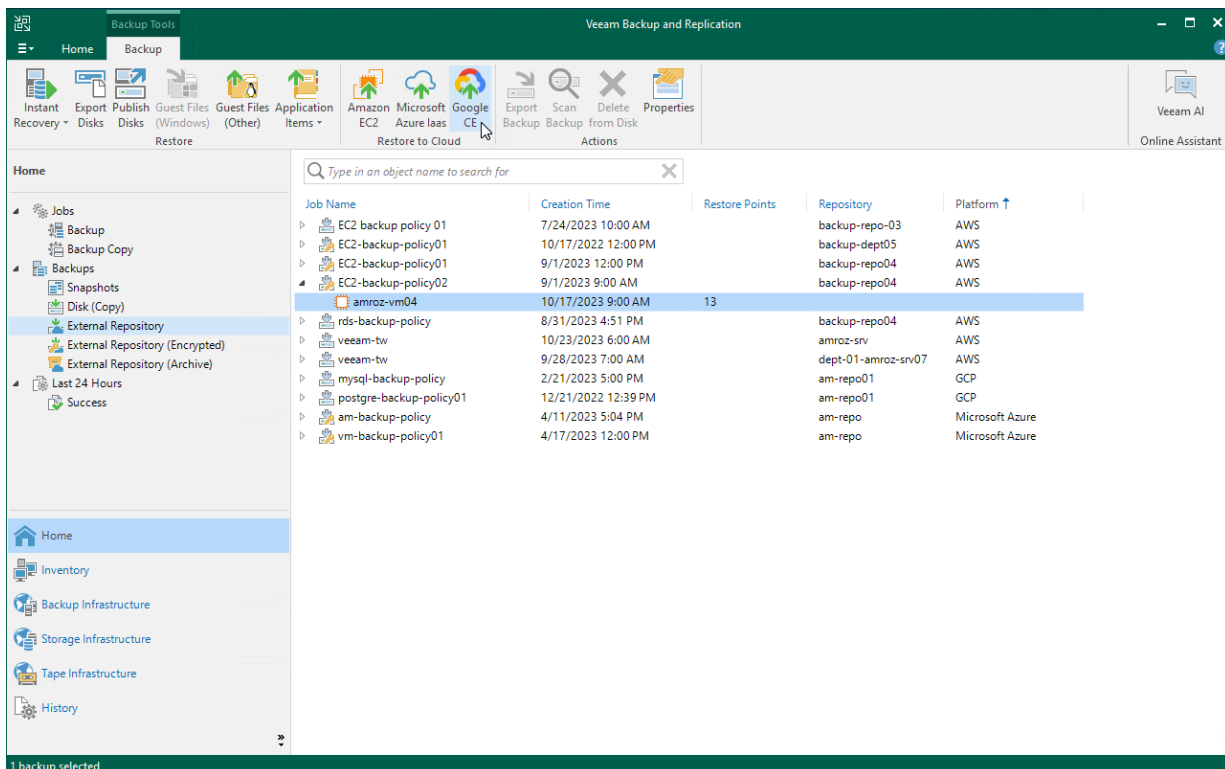
IMPORTANT

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	2 ↑
Warning	0 →
Success	42 ↑
Running now	0

Successful Policy Tasks (Last 24 hours)

- Snapshots: No snapshots created
- Replicas: No replicas created
- Backups: No backups created
- Archives: No archives created

Protected Workloads (Last 24 hours)

EC2 Instances	0 of 5	0%
RDS Instances	2 of 2	100%
EFS file systems	0 of 1	0%
DynamoDB tables	0 of 3	0%

Storage Usage

- Snapshots: 63
- Replicas: 46
- Backups: 10 GB
- Archives: 0 Bytes

Total Storage Usage: 10 GB

- S3 Standard: 10 GB
- S3 Glacier: 0 Bytes
- S3 Glacier Deep Archive: 0 Bytes

Top Policies (Type: Backup)

Policy	Duration	Start Time	Percentage
DynamoDB backup policy	2 min 59 sec	10/20 07:00 AM	-15%
EC2 backup policy 01	5 min 13 sec	10/19 10:00 AM	-19%
EFS backup policy	4 sec	10/19 09:05 AM	-94%
RDS backup policy 02	10 sec	10/19 09:00 AM	-100%

Bottlenecks Overview

- Policy sizing: 10/30/2023 3:06 PM OK
- CPU quota: 10/19/2023 10:05 AM Available
- Appliance disk usage: 10/30/2023 3:00 PM 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](#).

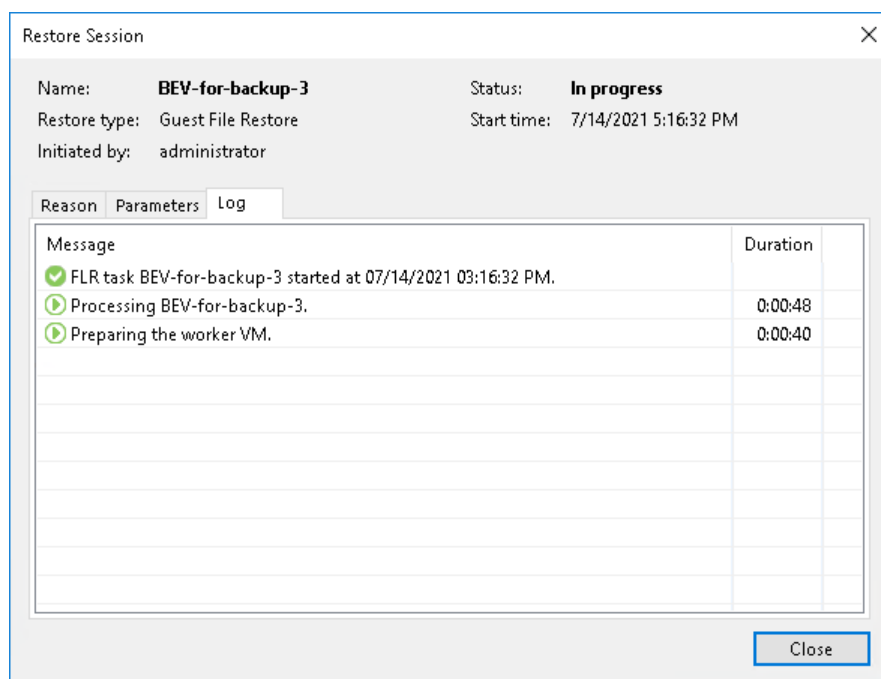
NOTE

Veeam Backup & Replication does not show statistics of EFS indexing sessions. For more information on indexing, see [Enable EFS Indexing](#).

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

- To filter operations by status, session, and workload type, click **Filter** and select the required options.
- 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 web interface. The main window title is "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.	—

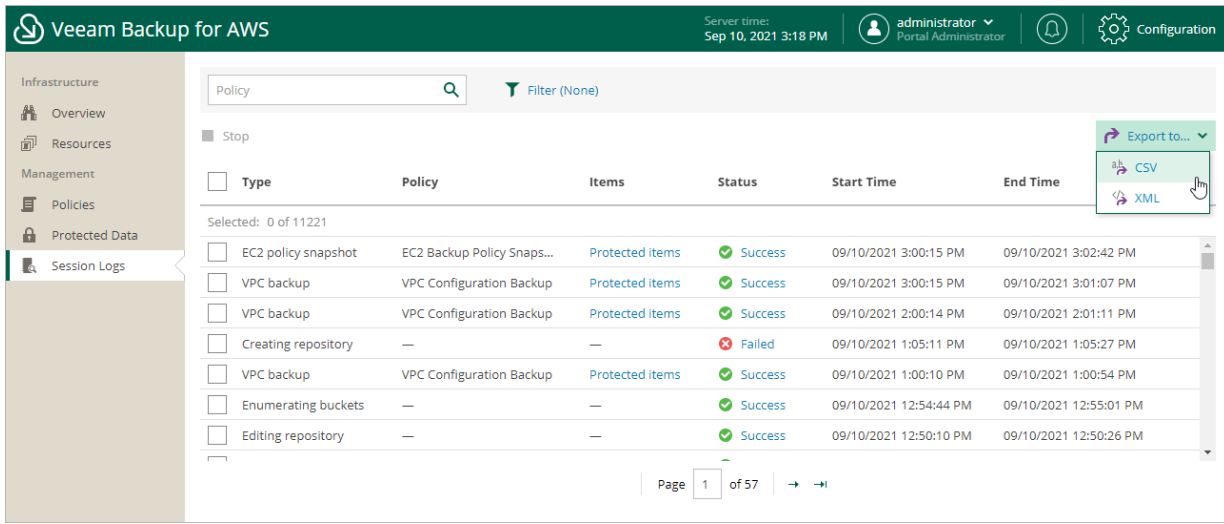
A "Close" button is located at the bottom right of the dialog box.

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 'Sep 10, 2021 3:18 PM', the user 'administrator Portal Administrator', and a 'Configuration' gear icon. A left-hand navigation menu lists categories like Infrastructure, Overview, Resources, Management, Policies, Protected Data, and Session Logs. The main content area features a search bar for 'Policy' and a 'Filter (None)' button. Below this is a table with columns for 'Type', 'Policy', 'Items', 'Status', 'Start Time', and 'End Time'. The table contains several rows of backup logs, including 'EC2 policy snapshot', 'VPC backup', and 'Creating repository'. An 'Export to...' dropdown menu is open in the top right corner of the table, showing options for 'CSV' and 'XML'. The page footer indicates 'Page 1 of 57'.

Type	Policy	Items	Status	Start Time	End Time
<input type="checkbox"/>	EC2 policy snapshot	EC2 Backup Policy Snaps...	Protected items ✓ Success	09/10/2021 3:00:15 PM	09/10/2021 3:02:42 PM
<input type="checkbox"/>	VPC backup	VPC Configuration Backup	Protected items ✓ Success	09/10/2021 3:00:15 PM	09/10/2021 3:01:07 PM
<input type="checkbox"/>	VPC backup	VPC Configuration Backup	Protected items ✓ Success	09/10/2021 2:00:14 PM	09/10/2021 2:01:11 PM
<input type="checkbox"/>	Creating repository	—	✗ Failed	09/10/2021 1:05:11 PM	09/10/2021 1:05:27 PM
<input type="checkbox"/>	VPC backup	VPC Configuration Backup	Protected items ✓ Success	09/10/2021 1:00:10 PM	09/10/2021 1:00:54 PM
<input type="checkbox"/>	Enumerating buckets	—	✓ Success	09/10/2021 12:54:44 PM	09/10/2021 12:55:01 PM
<input type="checkbox"/>	Editing repository	—	✓ Success	09/10/2021 12:50:10 PM	09/10/2021 12:50:26 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.

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 7.0 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. 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 7.0 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 7.0 from Version 6.0 or Earlier](#).



Upgrading to 7.0 from Version 6.0 or Earlier

To upgrade Veeam Backup for AWS to version 7.0, 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 [Upgrading Appliances Using Console](#).

When you perform upgrade to version 7.0 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 7.0 Works

When upgrading backup appliances to version 7.0 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. Launches a new EC2 instance from Veeam Backup for AWS 7.0 AMI that contains Ubuntu 22.04 LTS as an operating system.
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 file <https://www.postgresql.org/media/keys/ACCC4CF8.asc>.
- Outbound internet access must be allowed from the backup appliance to the [Veeam Update Notification Server](#) through port **443** over the HTTPS protocol.

- Outbound internet access must be allowed from the backup appliance to the Ubuntu Security Update 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 7.0 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 permissions 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 7.0 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 becomes invalid. 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 of Managed Backup Appliances](#).

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 [Upgrading Appliances Using Console](#).
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 DynamoDB tables. For more information, see [Integration with Veeam Backup & Replication](#).

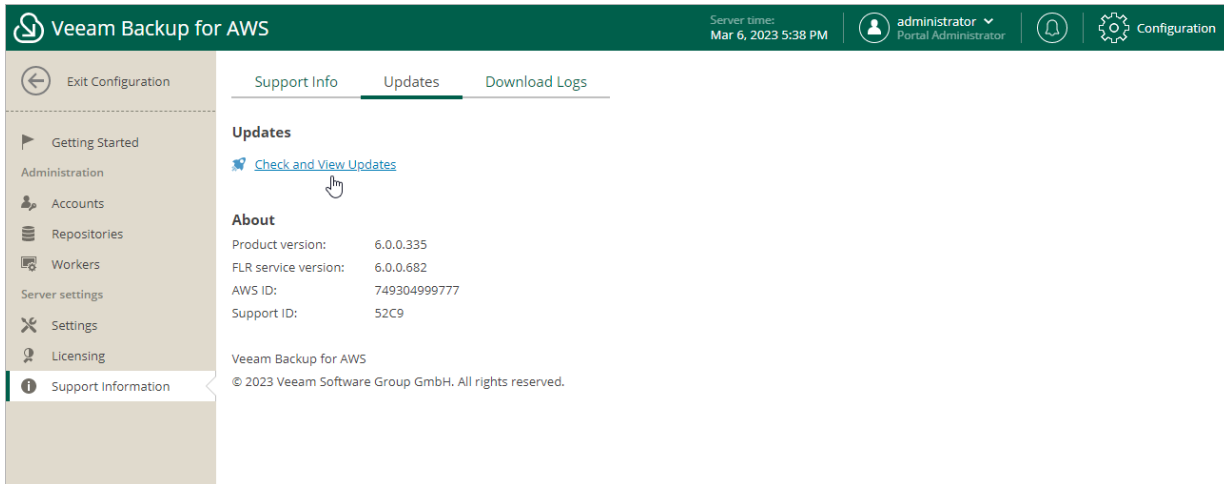
NOTE

When you upgrade to Veeam Backup for AWS version 7.0 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 7.0 from Version 6.0 or Earlier](#).

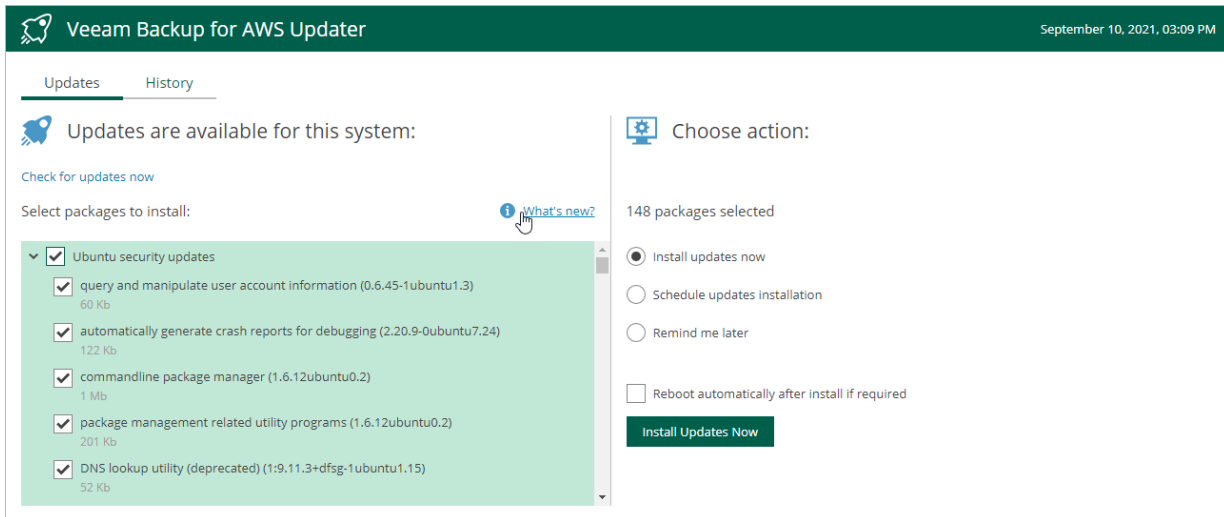
Checking for Updates

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 Backup for AWS 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 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 [Upgrading Appliances Using Console](#). Updating managed backup appliances using the Veeam updater service is not supported.

Installing Updates

IMPORTANT

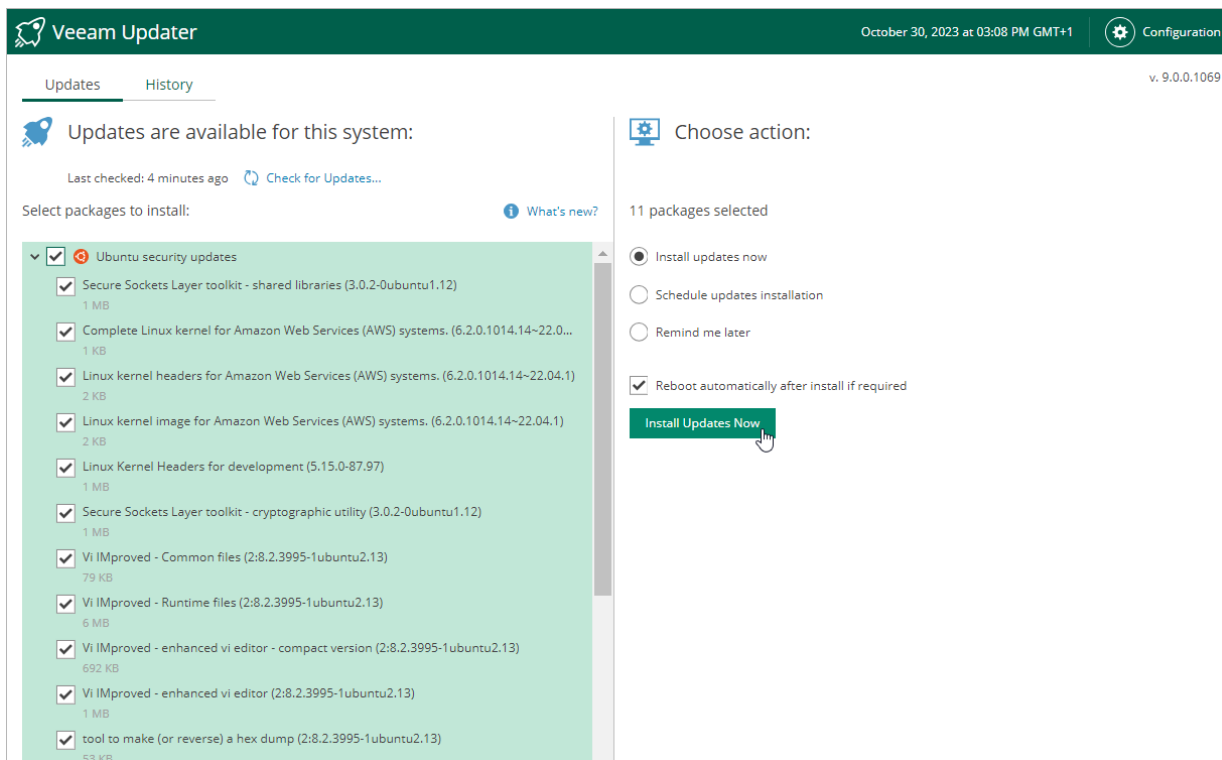
Before you install a product update, make sure that all backup policies are both disabled and stopped, and no restore tasks are currently executing. Otherwise, the update process will interrupt the running activities, which may result in data loss.

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 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

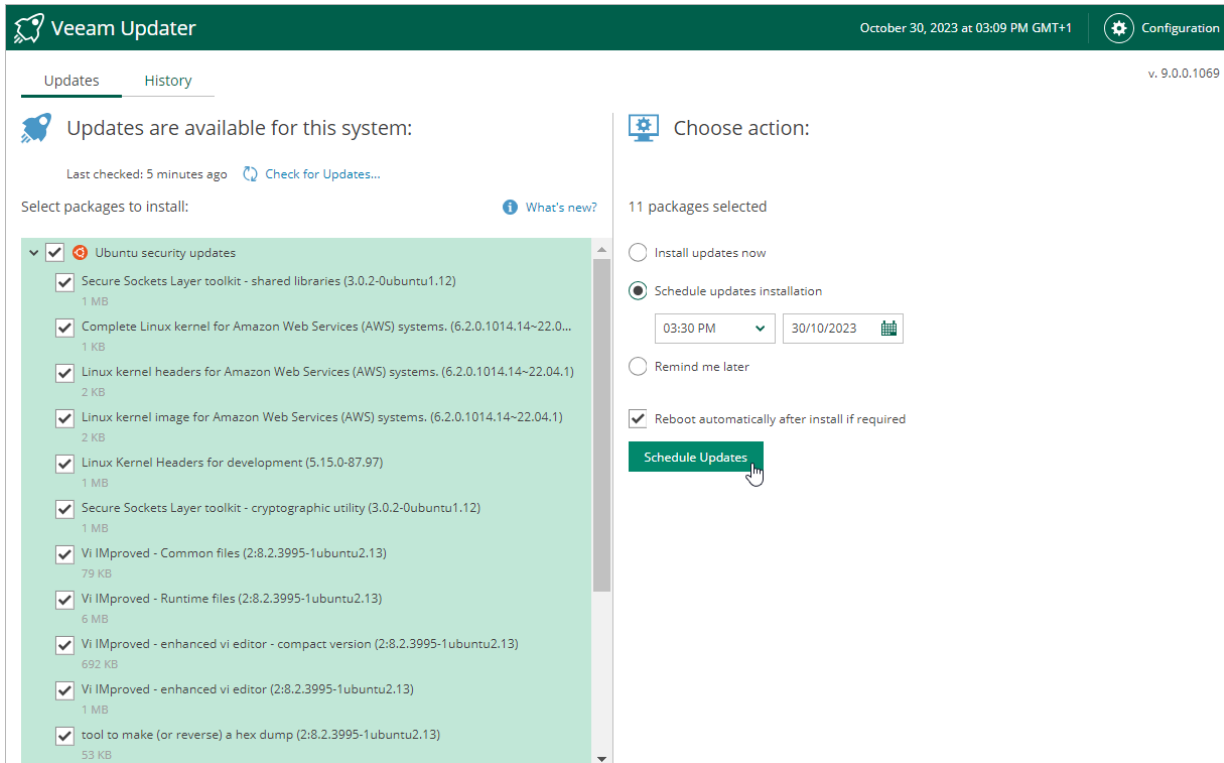
You can instruct Veeam Backup for AWS to automatically download and install available product versions and package updates on a specific date at a specific time:

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 the following Monday.

2. Click Remind me later.

The screenshot displays the Veeam Updater application interface. At the top, the title bar shows 'Veeam Updater' on the left, the date and time 'October 30, 2023 at 03:09 PM GMT+1' in the center, and a 'Configuration' button on the right. Below the title bar, there are two tabs: 'Updates' (active) and 'History'. The version number 'v. 9.0.0.1069' is visible in the top right corner.

The main content area is divided into two sections:

- Updates are available for this system:** This section includes a 'Last checked: 5 minutes ago' status and a 'Check for Updates...' button. Below this, it says 'Select packages to install:' and lists a group of updates under 'Ubuntu security updates'. Each update item has a checkbox and details such as the package name and size. The items listed are:
 - 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
- Choose action:** This section shows '0 packages selected' and three radio button options: 'Install updates now', 'Schedule updates installation', and 'Remind me later'. The 'Remind me later' option is selected. Below these options is a dropdown menu currently showing 'Next Week'. A green button labeled 'Remind me later' is highlighted with a mouse cursor.

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

After every product update, Veeam Backup for AWS checks if the [Default Backup Restore IAM role](#) created while installing the solution has all necessary permissions to perform backup and restore operations. If some of the permissions are missing, you will receive a warning in the notification area. For more information on permissions required for the *Default Backup Restore* IAM role after you update Veeam Backup for AWS to version 7.0, see [Full List of IAM Permissions](#).

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**.

The screenshot shows the Veeam Backup for AWS console interface. The top navigation bar includes the Veeam logo, the product name, the server time (Sep 9, 2021 2:16 PM), the user (administrator), and the Configuration menu. 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 titled 'IAM Roles' and contains a message about IAM roles and a table of roles. An 'IAM Roles Update' dialog box is open in the foreground, prompting the user to provide temporary credentials (Access key and Secret key) and offering 'Apply' and 'Cancel' buttons. The background table shows a list of IAM roles with columns for 'Created', 'Description', and 'Actions'.

Created	Description	Actions
2021 2:34:14 PM	Default Backup Restore	
2021 1:24:35 PM	acc2	

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, completed with warnings 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 and 'Veeam Updater' on the left, and the date 'October 30, 2023 at 03:12 PM GMT+1' and a 'Configuration' button on the right. Below the header, there are two tabs: 'Updates' and 'History', with 'History' being the active tab. The main content area is divided into two sections. On the left, there is a 'Update sessions history' section with a clock icon and a 'Date' column with a downward arrow. The list of sessions includes dates from August 25, 2023, to October 30, 2023. The top session, 'October 30, 2023 at 03:00 PM', is highlighted in green. On the right, there is a 'View Full Log' button and a table with three columns: 'Package', 'Status', and an empty column. The table contains four rows of update details, all with a 'Success' status.

Date ↓	Package	Status
October 30, 2023 at 03:00 PM	Preparing for updates	Success
October 26, 2023 at 01:35 PM	Veeam Backup for AWS (7.0.0.580)	Success
October 13, 2023 at 12:50 PM	File level recovery for Veeam backup (7.0.0.795)	Success
October 3, 2023 at 04:19 PM	Finalizing updates	Success
September 26, 2023 at 06:08 PM		
September 24, 2023 at 12:07 AM		
September 22, 2023 at 04:06 PM		
September 20, 2023 at 12:05 PM		
September 19, 2023 at 02:04 PM		
September 15, 2023 at 05:13 PM		
September 13, 2023 at 12:49 PM		
September 11, 2023 at 11:36 AM		
September 7, 2023 at 02:47 PM		
September 5, 2023 at 12:46 PM		
August 31, 2023 at 09:45 PM		
August 30, 2023 at 09:30 AM		
August 25, 2023 at 03:57 PM		

Configuring Web Proxy

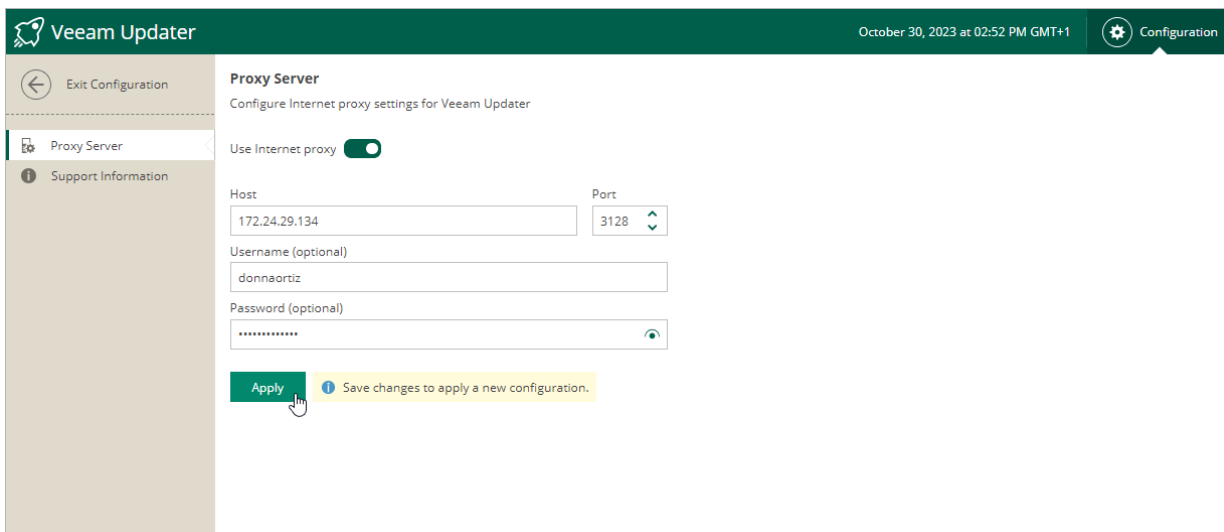
To check for available package updates for Veeam Backup for AWS, the Veeam Updater service running on the backup appliance connects to the Veeam Update repository over the internet. If the backup appliance is not connected to the internet, you can instruct the Veeam Updater service to use a web proxy that will provide access to the required resources.

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 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 **Apply**.

IMPORTANT

You cannot modify the web proxy settings during checking for updates.



The screenshot shows the Veeam Updater Configuration page for the Proxy Server. The page title is "Proxy Server" and the subtitle is "Configure Internet proxy settings for Veeam Updater". The "Use Internet proxy" toggle is turned on. The "Host" field contains "172.24.29.134" and the "Port" field contains "3128". The "Username (optional)" field contains "donnaortiz" and the "Password (optional)" field is masked with "*****". A green "Apply" button is visible at the bottom, with a tooltip that says "Save changes to apply a new configuration." The page header shows "Veeam Updater" and "October 30, 2023 at 02:52 PM GMT+1". The left sidebar shows "Exit Configuration", "Proxy Server", and "Support Information".

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](#).

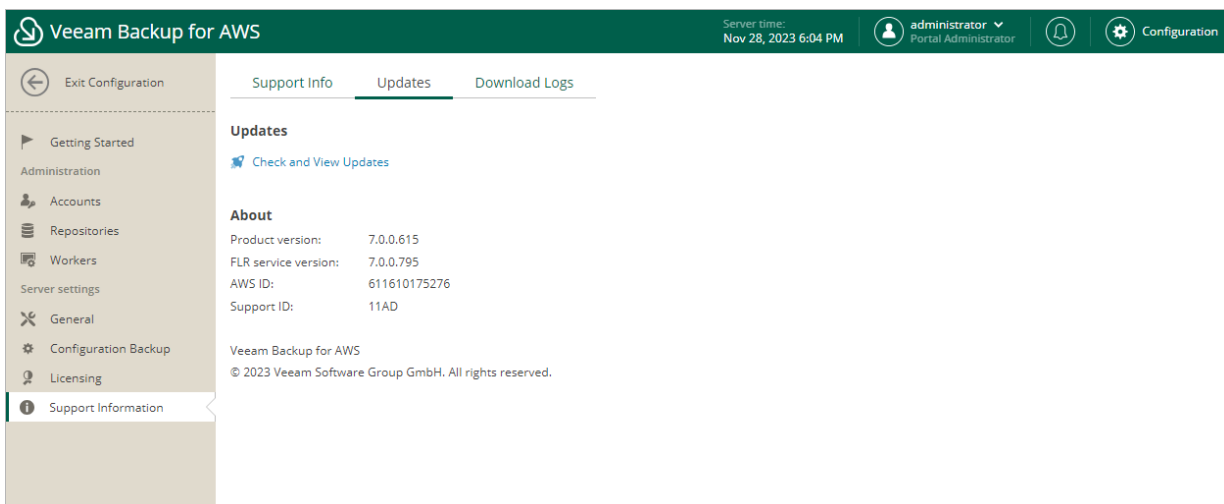
Viewing Product Details Using Web UI

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.



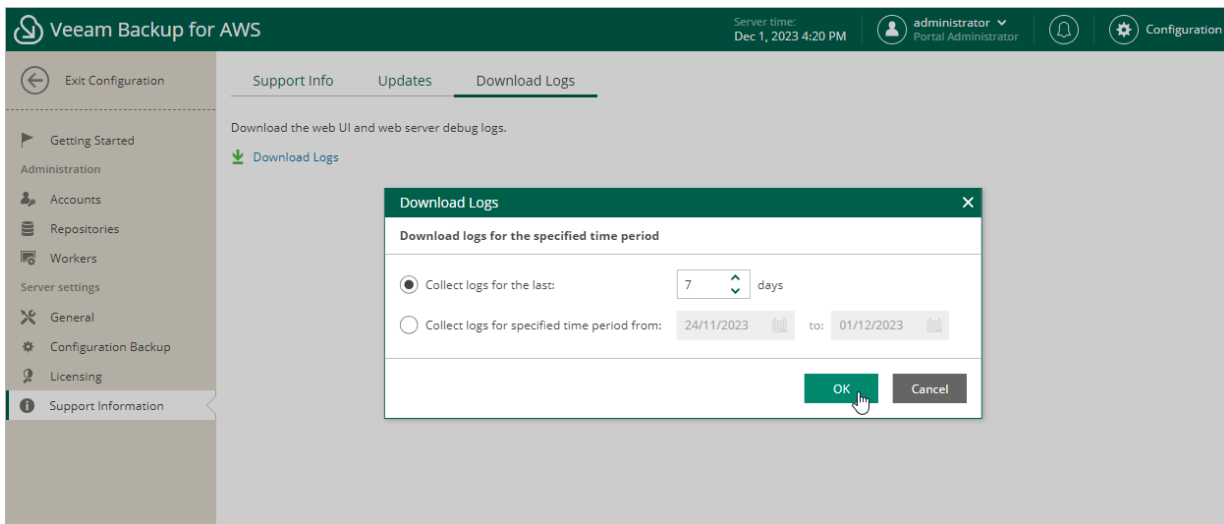
Downloading Product Logs Using Web UI

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 **AWS service** option. Then, in the **Use case** section, select an AWS service for which you plan to use the role and a specific use case for the service.
 - 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](#).

IMPORTANT

After the IAM role is created, you must configure trust relationships for the role to allow the Veeam Backup for AWS to use the IAM to perform operations in your infrastructure, as described in section [Before You Begin](#).

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 launch 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 com.amazonaws.<region>.sts 	<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 launched.

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 launched 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 launch 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 launch 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 launched.
 - 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 launched 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 launch 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 launch to perform specific operations, see [Architecture Overview](#).

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 launched.
 - 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 by 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
```