



Veeam Backup for Google Cloud

Version 5.0

User Guide

June, 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	7
ABOUT THIS GUIDE	8
OVERVIEW	9
Integration with Veeam Backup & Replication	10
Solution Architecture	11
Backup Server	12
Google Cloud Plug-in for Veeam Backup & Replication	12
Backup Appliances	12
Backup Repositories	14
Worker Instances	15
Additional Repositories and Tape Devices	18
Gateway Servers	18
Protecting VM Instances	19
VM Backup	20
VM Restore	26
Protecting Cloud SQL Instances	30
SQL Backup	31
SQL Restore	37
Protecting Cloud Spanner Instances	39
Spanner Backup	40
Spanner Restore	45
Retention Policies	47
Data Encryption	48
Storage Bucket Encryption	48
Cloud KMS Encryption	48
PLANNING AND PREPARATION	52
System Requirements	53
Ports	54
Plug-In Permissions	56
Service Account Permissions	59
Default Permissions	59
Repository Permissions	59
Worker Permissions	60
Snapshot Permissions	69
Backup Permissions	73
Restore Permissions	77
Permissions Changelog	82

Google Cloud APIs	85
Considerations and Limitations	86
Sizing and Scalability Guidelines	89
Backup Appliance	89
Object Storage	91
Backup Policies.....	92
Worker Instances	93
DEPLOYMENT	96
Deploying Plug-In	97
Installing Plug-In in Unattended Mode.....	98
Upgrading Plug-In	99
Uninstalling Plug-In.....	100
Deploying Backup Appliance.....	101
Deploying Backup Appliance from Console.....	101
Deploying Backup Appliance from Google Cloud Marketplace.....	112
Uninstalling Veeam Backup for Google Cloud	117
LICENSING	120
Limitations.....	120
Scenarios	121
Backup Appliance Licensing	122
Installing and Removing Backup Appliance License	123
Viewing License Information.....	125
Revoking License Units.....	129
ACCESSING VEEAM BACKUP FOR GOOGLE CLOUD	131
Accessing Web UI from Console	132
Accessing Web UI from Workstation.....	133
CONFIGURING VEEAM BACKUP FOR GOOGLE CLOUD	135
Managing Backup Appliances.....	136
Adding Appliances	137
Editing Appliance Settings	149
Rescanning Appliances	150
Removing Appliances	151
Managing Backup Repositories	153
Adding Backup Repositories Using Console	154
Adding Backup Repositories Using Web UI	164
Editing Backup Repositories.....	173
Rescanning Backup Repositories	176
Removing Backup Repositories	177
Managing Service Accounts	179
Adding Service Accounts.....	180

Editing Service Accounts.....	187
Removing Service Accounts	188
Managing Projects and Folders	189
Adding Projects and Folders	190
Editing Projects and Folders	199
Removing Projects and Folders	200
Managing User Accounts	201
Adding User Accounts.....	203
Editing User Accounts.....	208
Changing User Passwords	209
Enabling Multi-Factor Authentication	210
Managing Cloud SQL Accounts.....	211
Adding Cloud SQL Accounts.....	212
Editing Cloud SQL Accounts.....	217
Removing Cloud SQL Accounts	218
Managing Worker Instances	219
Managing Worker Configurations	220
Managing Worker Profiles.....	230
Assigning Worker Instance Labels	237
Configuring General Settings.....	238
Configuring Global Retention Settings	239
Configuring Global Notification Settings	241
Replacing Security Certificates.....	249
Changing Time Zone	250
Registering Application.....	251
Performing Configuration Backup and Restore	253
Performing Configuration Backup	254
Performing Configuration Restore.....	258
VIEWING AVAILABLE RESOURCES	276
Adding Resources to Policies	277
PERFORMING BACKUP.....	278
Performing Backup Using Console	280
Creating Backup Policies	281
Editing Backup Policy Settings.....	282
Enabling and Disabling Backup Policies	283
Starting and Stopping Backup Policies.....	284
Deleting Backup Policies.....	285
Creating Backup Copy Jobs	286
Copying Backups to Tapes	287
Performing Backup Using Web UI.....	288

Performing VM Backup	289
Performing SQL Backup	322
Performing Spanner Backup	358
Managing Backup Policies	389
MANAGING BACKED-UP DATA	396
Managing Backed-Up Data Using Console	397
Managing Backed-Up Data Using Web UI	400
PERFORMING RESTORE	403
VM Restore	404
VM Restore Using Console	405
VM Restore Using Web UI	419
SQL Restore	463
SQL Restore Using Console	464
SQL Restore Using Web UI	476
Spanner Restore	504
Spanner Restore Using Console	505
Spanner Restore Using Web UI	506
Instant Recovery	531
Exporting Disks	533
Publishing Disks	534
Restoring to AWS	535
Restoring to Microsoft Azure	536
Restoring to Nutanix AHV	537
REVIEWING DASHBOARD	539
VIEWING SESSION STATISTICS	541
COLLECTING OBJECT PROPERTIES	544
UPDATING VEEAM BACKUP FOR GOOGLE CLOUD	545
Updating Appliances Using Console	546
Updating Appliances Using Web UI	547
Upgrading Appliances	548
Checking for Updates	549
Installing Updates	551
Viewing Updates History	555
Configuring Web Proxy	556
GETTING TECHNICAL SUPPORT	557
APPENDIX. CONFIGURING DEPLOYMENT MODE	561

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](https://www.veeam.com/documentation-guides-datasheets.html)
- Veeam R&D Forums: forums.veeam.com

About This Guide

This guide is designed for IT professionals who plan to use Veeam Backup for Google Cloud. The guide includes system requirements, licensing information and step-by-step deployment instructions. It also provides a comprehensive set of features to ensure easy execution of protection and disaster recovery tasks in Google Cloud environments.

Overview

NOTE

Starting from Veeam Backup for Google Cloud version 5.0, Google Cloud Plug-in for Veeam Backup & Replication is part of the Veeam Backup for Google Cloud architecture. That is why the [Google Cloud Plug-in for Veeam Backup & Replication User Guide](#) has been merged into the main product guide.

Veeam Backup for Google Cloud is a solution developed for protection and disaster recovery tasks for Google Cloud environments. With Veeam Backup for Google Cloud, you can perform the following operations:

- Create image-level backups and cloud-native snapshots of VM instances.
- Create image-level backups and cloud-native snapshots of Cloud SQL instances.
- [Available only for backup appliances managed by Veeam Backup & Replication] Create image-level backups and cloud-native snapshots of Cloud Spanner instances.
- Keep the backed-up data in cost-effective, long-term Google storage buckets.
- Restore entire Cloud SQL instances and specific Cloud SQL databases.
- Restore entire VM instances, individual persistent disks, and guest OS files and folders.
- [Available only for backup appliances managed by Veeam Backup & Replication] Restore entire VM instances to Microsoft Azure, AWS and Nutanix AHV.
- [Available only for backup appliances managed by Veeam Backup & Replication] Perform Instant Recovery of VM instances to VMware vSphere and Hyper-V environments, and to Nutanix AHV clusters.
- Restore entire Cloud Spanner instances and specific Cloud Spanner databases.

IMPORTANT

Starting from version 5.0, Veeam Backup for Google Cloud is part of the Veeam Backup & Replication solution, and some 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 Veeam Backup for Google Cloud 5.0, Google Cloud Plug-in for Veeam Backup & Replication is part of the Veeam Backup for Google Cloud solution. Veeam Backup for Google Cloud extends the Veeam Backup & Replication functionality and allows you to add backup appliances to Veeam Backup & Replication. With Google Cloud 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 5.0 comes with a major feature — the ability to protect Cloud Spanner resources — that is available only for those backup appliances managed by a Veeam Backup & Replication server. To unlock the full functionality, you must [install Google Cloud Plug-in for Veeam Backup & Replication](#) on the server and [add your appliances](#) to the backup infrastructure.

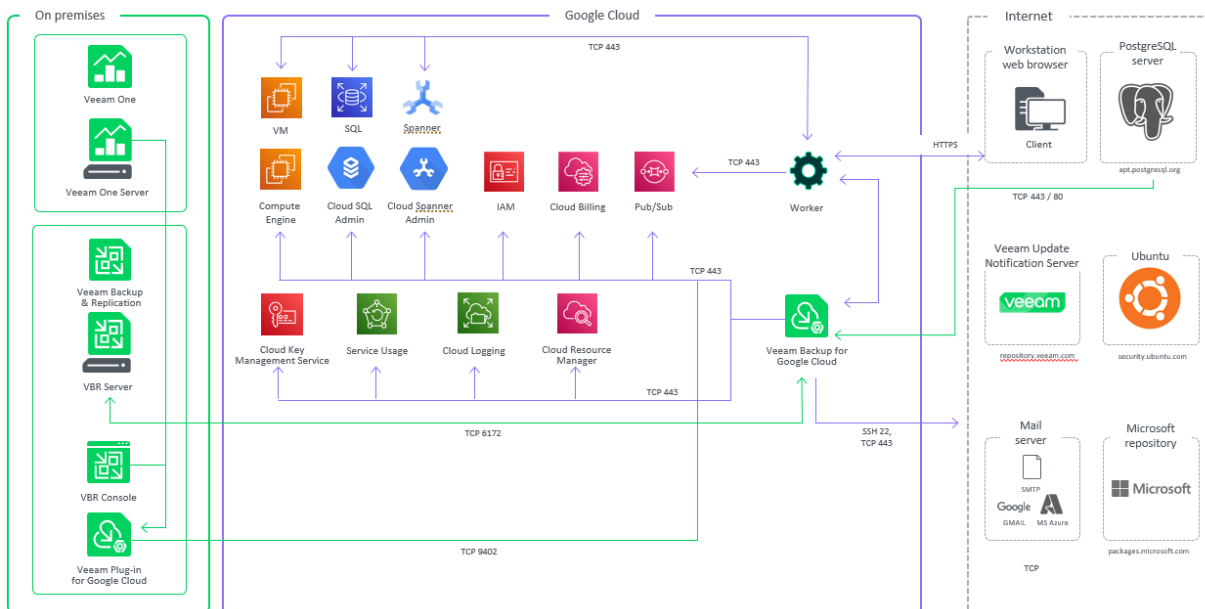
IMPORTANT

- If you remove a backup appliance from the backup infrastructure, you will no longer be able to add, enable and start Spanner backup policies. Creating Cloud Spanner snapshots 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 protect Cloud Spanner instances.

Solution Architecture

The Veeam Backup for Google Cloud architecture includes the following components:

- Backup server
- Google Cloud 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 of the backup infrastructure. For more information, see the Veeam Backup & Replication User Guide, section [Backup Server](#).

Google Cloud Plug-in for Veeam Backup & Replication

Plug-in is an architecture component that extends the Veeam Backup & Replication functionality and allows you to add backup appliances to the backup infrastructure. With Google Cloud Plug-in for Veeam Backup & Replication, you can manage data protection and disaster recovery operations from the Veeam Backup & Replication console.

Backup Appliances

A backup appliance is a Linux-based VM instance where Veeam Backup for Google Cloud is installed. 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.

The backup appliance also maintains the configuration database that stores data collected from Veeam Backup for Google Cloud for the existing backup policies, protected VM, Cloud SQL and Cloud Spanner instances, deployed worker instances, connected Google Cloud projects and so on.

TIP

If you have multiple backup appliances deployed in Google Cloud, 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 Google Cloud operations. For more information on the Veeam Backup & Replication console, see the [Veeam Backup & Replication User Guide](#).

Backup Appliance Software

The VM instance running Veeam Backup for Google Cloud is deployed with the pre-installed set of software components:

- Ubuntu 20.04
- ASP.NET Core Runtime 6.0
- PostgreSQL 12
- nginx 1.24.0
- libpam-google-authenticator 20170702-2

- Veeam Backup for Google Cloud 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 Google Cloud](#).

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 Google Cloud.
- **Web UI** — provides a web interface that allows user to access to the Veeam Backup for Google Cloud functionality.
- **Updater service** — allows Veeam Backup for Google Cloud to check, view and install product and package updates.
- **Self Backup service** — allows Veeam Backup for Google Cloud to back up and restore the configuration database of the backup appliance.
- **REST API service** — allows users to perform operations with Veeam Backup for Google Cloud entities using HTTP requests and standard HTTP methods. For details, see the [Veeam Backup for Google Cloud REST API Reference](#).

Backup Repositories

A backup repository is a subdirectory in a Google Cloud storage bucket where Veeam Backup for Google Cloud stores backups of protected VM instances, Cloud SQL instances and Cloud Spanner instances.

To communicate with a backup repository, Veeam Backup for Google Cloud 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 Google Cloud communicates with backup repositories, see [Managing Backup Repositories](#).

IMPORTANT

Backups are stored in backup repositories in the native Veeam format and must be modified neither manually nor by 3rd party tools, including native Google Cloud capabilities (for example, the [Autoclass feature](#)). Otherwise, Veeam Backup for Google Cloud may fail to restore the backed-up data.

Encryption on Repositories

For enhanced data security, Veeam Backup for Google Cloud allows you to enable encryption at the repository level. Veeam Backup for Google Cloud uses the same encryption standards as Veeam Backup & Replication to encrypt backups stored in backup repositories. To learn what encryption standards Veeam Backup & Replication uses to encrypt its data, see the Veeam Backup & Replication User Guide, section [Encryption Standards](#).

To learn how to enable encryption at the repository level, see [Data Encryption](#).

Limitations for Repositories

To use a storage bucket as a target location for backups, you must connect to a project in which this bucket resides as described in section [Adding Backup Repositories](#).

Veeam Backup for Google Cloud allows you to store backups in the *Standard*, *Nearline* and *Archive* storage classes. The *Coldline* storage class is not supported. For more information on storage classes offered by Cloud Storage, see [Google Cloud documentation](#).

Worker Instances

A worker instance is an auxiliary Linux-based VM instance that is responsible for the interaction between the backup appliance and other components of the Veeam Backup for Google Cloud architecture. Worker instances process backup workload and distribute backup traffic when transferring data to and from backup repositories.

Worker Instance Components

A worker instance uses the following components:

- **Veeam Data Mover** – the service that performs data processing tasks. During backup, Veeam Data Mover retrieves data from snapshots and stores the retrieved data to backup repositories. During restore, Veeam Data Mover transfers backed-up data from backup repositories to the target location.
- **File-Level Restore Browser** – the web service that allows you to find and save files and folders of a backed-up VM instance to a local machine or to the original location. The File-Level Restore browser is installed automatically on every worker instance that is launched for file-level recovery.

For more information on recovering files of VM instances with the File-Level Restore browser, see [Performing File-Level Recovery](#).

Security Certificates for Worker Instances

Veeam Backup for Google Cloud uses self-signed TLS certificates to establish secure communication between the web browser on a user workstation and the File-Level Restore browser running on a worker instance during the file-level recovery process. A self-signed certificate is generated automatically on the worker instance when the recovery session starts.

How Worker Instances Work

Veeam Backup for Google Cloud automatically deploys a worker instance in Google Cloud for the duration of a backup or restore process, and removes it immediately as soon as the process is over. To minimize cross-region traffic charges and to speed up the data transfer, depending on the performed operation, Veeam Backup for Google Cloud deploys the worker instance in the following location:

Operation	Worker Instance Location	Default Worker Machine Types
Creating image-level backups of VM instances	Google Cloud region in which a processed VM instance resides	e2-highcpu-8, with an additional empty standard persistent (pd-standard) disk up to 4000 GB in size
Creating image-level backups of Cloud SQL instances	Google Cloud region in which a processed Cloud SQL instance resides	e2-highcpu-8
Creating image-level backups of Cloud SQL instances using a staging server	Google Cloud region in which a source Cloud SQL instance resides	e2-highcpu-8

Operation	Worker Instance Location	Default Worker Machine Types
Creating image-level backups of Cloud Spanner instances	Either the Google Cloud region in which a target backup repository resides, or the region in which read-write and read-only replicas are located, or any other region defined by the Google Cloud logic	e2-highcpu-8
Creating archived image-level backups of VM instances	Google Cloud region in which a target backup repository of the <i>Standard</i> or <i>Nearline</i> storage class resides.	e2-standard-4
Creating archived image-level backups of Cloud SQL instances	Google Cloud region in which a target backup repository of the <i>Standard</i> or <i>Nearline</i> storage class resides	e2-standard-4
Creating archived image-level backups of Cloud Spanner instances	Google Cloud region in which a target backup repository of the <i>Standard</i> or <i>Nearline</i> storage class resides	e2-standard-4
Performing health check for created restore points	Google Cloud region in which a target backup repository of the <i>Standard</i> or <i>Nearline</i> storage class resides	e2-standard-4
Applying retention policy settings to created restore points	Google Cloud region in which a backup repository with backed-up data resides	e2-highcpu-8
Restoring VM instances	Google Cloud region to which a VM instance is restored	e2-highcpu-4, with an additional empty standard persistent (pd-standard) disk up to 1500 GB in size
Restoring Cloud SQL instances	Google Cloud region in which a backup repository with backed-up data resides (for MySQL instances); Google Cloud region in which the restored Cloud SQL instance will reside (for PostgreSQL instances)	e2-highcpu-8
Restoring Cloud Spanner instances	Either the Google Cloud region to which a Cloud Spanner instance is restored, or the region in which read-write replicas are located	e2-highcpu-8

Operation	Worker Instance Location	Default Worker Machine Types
Restoring individual persistent disks of VM instances	Google Cloud region to which the persistent disks of a processed VM instance are restored	e2-highcpu-4, with an additional empty standard persistent (pd-standard) disk up to 1500 GB in size
Restoring specific Cloud SQL databases	Google Cloud region in which a backup repository with backed-up data resides (for MySQL databases); Google Cloud region in which the target Cloud SQL instance resides (for PostgreSQL databases)	e2-highcpu-8
Restoring specific Cloud Spanner databases	Either the Google Cloud region to which the databases of a processed Cloud Spanner instance are restored, or the region in which read-write replicas are located	e2-highcpu-8
File-level recovery from cloud-native snapshots	Google Cloud region in which a source VM instance resides	e2-highcpu-4
File-level recovery from image-level backups	Google Cloud region in which a backup repository with backed-up data resides	e2-highcpu-4

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

IMPORTANT

For Veeam Backup for Google Cloud to deploy the number of worker instances required for a backup or restore process, you must have enough resource quotas allocated between your projects. To learn how to check your quotas, see [Google Cloud documentation](#).

For the list of network ports that must be open to ensure proper communication of worker instances with other components of the Veeam Backup for Google Cloud architecture, see [Ports](#).

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

A gateway server is an auxiliary backup infrastructure components that provide access from the backup server to repositories. By default, the role of a gateway server is assigned to the backup server.

Gateway servers cache 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 VM Instances

To produce cloud-native snapshots and image-level backups of VM instances, Veeam Backup for Google Cloud 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, and so on.

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

How to Protect VM Instances

To create a VM backup policy, complete the following steps:

1. [Check limitations and prerequisites.](#)
2. [Add service accounts.](#)
3. [Connect projects and folders.](#)
4. [Add backup repositories.](#)
5. [Configure worker instance settings.](#)
6. [Configure global retention and email notification settings.](#)
7. [Complete the Add VM Policy wizard.](#)

VM Backup

Veeam Backup for Google Cloud performs VM instance backup in the following way:

1. Creates snapshots of persistent disks that are attached to the processed VM instance.

PD snapshots are assigned resource labels upon creation. Keys and values of resource labels contain encrypted metadata that helps Veeam Backup for Google Cloud identify the related PD snapshots and treat them as a single unit — a cloud-native snapshot.
2. If you enable image-level backup for the backup policy, Veeam Backup for Google Cloud performs the following operations:
 - a. Deploys a worker instance within the worker project in the Google Cloud region in which the processed VM instance resides. For more information, see [Managing Worker Instances](#).
 - b. Re-creates the persistent disks from the cloud-native snapshot created at step 1 and attaches them to the worker instance.

Note that the cloud-native snapshot used as a source for image-level backup is not a temporary snapshot — when the backup session completes, this snapshot remains in the snapshot chain and is deleted later according to the specified [policy scheduling settings](#).
 - c. Reads data from the persistent disks on the worker instance, transfers the data to the target standard or nearline repository, and stores it in the native Veeam format.

Veeam Backup for Google Cloud encrypts and compresses data saved to storage buckets. For more information, see [Enabling Data Encryption](#).
 - d. Removes the worker instance when the backup session completes.
3. If you enable the [backup archiving mechanism](#), Veeam Backup for Google Cloud performs the following operations:
 - a. Deploys a worker instance within the worker project in the Google Cloud region in which the processed VM instance resides.

For more information on how to specify a project for worker instances, see [Managing Worker Configurations](#).
 - b. Retrieves data from the target standard or nearline repository, and transfers it to the target archive repository.
 - c. Removes the worker instance when the archive session completes.

Snapshot Chain

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

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

1. During the first backup session, Veeam Backup for Google Cloud creates a snapshot of all instance data and, by default, saves it in the multi-regional location closest to the region in which the original 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 Google Cloud copies the whole image of the instance.

TIP

You can change the default location of cloud-native snapshots created for VM instances in the [backup policy settings](#).


2. During subsequent backup sessions, Veeam Backup for Google Cloud 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 processed data.

For more information on how incremental VM snapshots work, see [Google Cloud documentation](#).

Cloud-native snapshots in the snapshot chain are assigned encrypted labels. These labels store information about the protected instances and the backup policies that created the snapshots. Veeam Backup for Google Cloud uses the encrypted labels to identify outdated snapshots, 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.

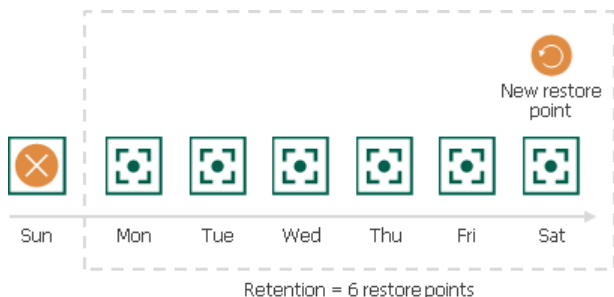


The number of cloud-native snapshots kept in the snapshot chain is defined by retention policy settings. For more information, see [VM Snapshot Retention](#).

VM Snapshot Retention

For cloud-native snapshots, Veeam Backup for Google Cloud retains the number of latest restore points defined in backup scheduling settings as described in section [Creating VM Policies](#).

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



NOTE

Retention policy settings configured when creating backup policies do not apply to cloud-native snapshots created manually. To learn how to remove these snapshots, see [Managing Backed-Up Data](#).

Backup Chain

If you enable image-level backups for a backup policy, Veeam Backup for Google Cloud creates a new backup in a standard or nearline repository during every backup session. A sequence of backups created during a set of backup sessions makes up a regular backup chain.


The regular backup chain includes backups of the following types:

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

To create a regular backup chain for a VM instance protected by a backup policy, Veeam Backup for Google Cloud implements the forever forward incremental backup method:

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

Veeam Backup for Google Cloud creates incremental backups based on the Veeam proprietary filtering mechanism that filters out unchanged data blocks by calculating a checksum for every block. The Google Cloud changed block tracking (CBT) mechanism that would allow tracking changed blocks of data and would increase the efficiency of incremental backups is not implemented at the moment.



Full and incremental backups act as restore points for backed-up instances that let you roll back instance data to the necessary state. To recover an 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 regular backup chain is missing, you will not be able to roll back to the necessary state. For this reason, you must not delete individual backups from the backup repository manually. Instead, you must specify retention policy settings that will let you maintain the necessary number of backups in the repository. For more information, see [VM Backup Retention](#).

Archive Backup Chain

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


The archive backup chain includes backups of the following types:

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

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

1. During the first archive session, Veeam Backup for Google Cloud detects backed-up data that is stored in the full backup and all incremental 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 Google Cloud checks the regular 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 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.

Veeam Backup for Google Cloud creates incremental backups based on the Veeam proprietary filtering mechanism that filters out unchanged data blocks by calculating a checksum for every block. The Google Cloud changed block tracking (CBT) mechanism that would allow tracking changed blocks of data and would increase the efficiency of incremental backups is not implemented at the moment.



Full and incremental archive backups act as restore points for backed-up instances that let you roll back instance data to the necessary state. To recover an 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 backups from the archive repository manually. Instead, you must specify retention policy settings that will let you maintain the necessary number of backups in the archive repository. For more information, see [Retention Policy for Archived Backups](#).


VM Backup Retention

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


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

1. Veeam Backup for Google Cloud checks the configuration database to detect standard and nearline repositories that contain outdated restore points.
2. If an outdated restore point exists in a backup repository, Veeam Backup for Google Cloud deploys a worker instance in a Google Cloud region in which the repository with backed-up data resides.


3. Veeam Backup for Google Cloud transforms the regular backup chain in the following way:
 - a. Rebuilds the full backup to include the data of the incremental backup that follows the full backup. To do that, Veeam Backup for Google Cloud injects into the full backup data blocks from the earliest incremental backup in the chain. This way, the full backup 'moves' forward in the regular backup chain.



- b. Removes the earliest incremental backup from the chain as redundant — this data has already been injected into the full backup.



4. Veeam Backup for Google Cloud repeats step 2 for all other outdated restore points found in the regular 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 Google Cloud ensures that the regular backup chain is not broken and that you will be able to recover your data when needed.



5. Veeam Backup for Google Cloud removes the worker instance when the retention session completes.

NOTE

Each worker instance can process only one retention task at a time, and Veeam Backup for Google Cloud can simultaneously deploy maximum 10 worker instances to process retention tasks. If the number of retention tasks that must be processed by worker instances is more than the specified limit, the tasks exceeding this limit are queued.

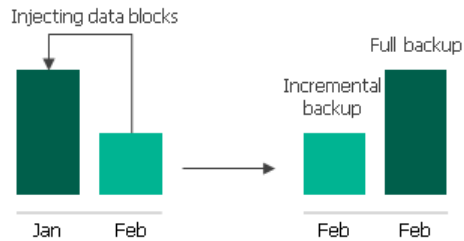
Retention Policy for Archived Backups

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


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

1. Veeam Backup for Google Cloud checks the configuration database to detect archive backup repositories that contain outdated restore points.


2. If an outdated restore point exists in a backup repository, Veeam Backup for Google Cloud deploys a worker instance in a Google Cloud region in which the repository with backed-up data resides.
3. Veeam Backup for Google Cloud transforms the archive backup chain in the following way:
 - a. Rebuilds the full archive backup to include the data of the incremental archive backup that follows the full archive backup. To do that, Veeam Backup for Google Cloud 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.



- b. Removes the earliest incremental archive backup from the chain as redundant — this data has already been injected into the full archive backup.



4. Veeam Backup for Google Cloud 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 Google Cloud ensures that the archive backup chain is not broken and that you will be able to recover your data when needed.



5. Removes the worker instance when the retention session completes.

NOTES

Each worker instance can process only one retention task at a time, and Veeam Backup for Google Cloud can simultaneously deploy maximum 10 worker instances to process retention tasks. If the number of retention tasks that must be processed by worker instances is more than the specified limit, the tasks exceeding this limit are queued.

VM Restore

Veeam Backup for Google Cloud offers the following restore operations:

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

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

Instance Restore

To restore a VM instance from a cloud-native snapshot, Veeam Backup for Google Cloud uses [native Google Cloud capabilities](#). To restore a VM instance from an image-level backup, Veeam Backup for Google Cloud performs the following steps:

1. Deploys a worker instance within the worker project in the Google Cloud region in which the restored VM instance will reside.

For more information on how to specify a project for worker instances, see [Managing Worker Configurations](#).
2. Creates empty persistent disks and attaches them to the worker instance.

The number of empty persistent disks equals the number of persistent disks attached to the backed-up VM instance.
3. Restores backed-up data to the empty persistent disks on the worker instance.
4. Takes cloud-native snapshots of the persistent disks with the restored data.
5. Creates disks from the snapshots in the target location (that is, the project and region specified for the restore operation).
6. Removes the worker instance from Google Cloud.
7. Removes all the created snapshots from Google Cloud Storage.
8. Creates a VM instance in the target location and attaches the created persistent disks with the restored data to the VM instance.
9. [This step applies only if you perform restore to the original location and if the source VM instance is still present in the location] Powers off the source VM instance, removes the source VM instance from Google Cloud and then renames the restored VM instance.

IMPORTANT

To allow Veeam Backup for Google Cloud to perform restore to the original location while source VM instances still exist there, the [deletion protection](#) setting must be disabled for the source instance, and the service account whose permissions are used for the operation must be assigned the `compute.instances.setName` permission.

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

Disk Restore

To restore persistent disks from a cloud-native snapshot, Veeam Backup for Google Cloud uses [native Google Cloud capabilities](#). To restore persistent disks from an image-level backup, Veeam Backup for Google Cloud performs the following steps:

1. Deploys a worker instance within the worker project in the Google Cloud region in which the restored persistent disks will reside.

For more information on how to specify a project for worker instances, see [Managing Worker Configurations](#).

2. Creates empty persistent disks and attaches them to the worker instance.

The number of empty persistent disks equals the number of disks you want to restore.

NOTE

Every time before creating persistent disks, Veeam Backup for Google Cloud checks whether the total size of pd-standard disks breaches the zone quota for the project in which the worker instance is deployed. If the total disk size is less than 1500 GB, Veeam Backup for Google Cloud temporarily attaches an additional empty disk to the worker instance — but only for the duration of the restore process and if the quota allows attaching the disk. This allows Veeam Backup for Google Cloud to speed up the data transfer to reduce your restore costs.

3. Restores backed-up data to the empty persistent disks on the worker instance.
4. Takes cloud-native snapshots of the persistent disks with the restored data.
5. Creates disks from snapshots in the target location (that is, the project and region specified for the restore operation).
6. Removes the worker instance from Google Cloud.
7. Removes all the created snapshots from Google Cloud Storage.

NOTE

Veeam Backup for Google Cloud does not attach the restored persistent disks to any VM instances — the disks are placed to the specified location as standalone persistent disks.

To learn how to restore persistent disks attached to a VM instance from a cloud-native snapshot or an image-level backup, see [Performing Disk Restore](#).

File-Level Recovery

Veeam Backup for Google Cloud allows you to recover the files and folders of a backed-up VM instance to a local machine or to the original location.

File-Level Recovery to Local Machine

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

1. Deploys a worker instance within the worker project in either of the following Google Cloud regions.
 - o To recover files and folders from a cloud-native snapshot, the worker instance is deployed in the region in which the VM instance resides.
 - o To recover files and folders from an image-level backup, the worker instance is deployed in the region in which the storage bucket with backed-up data resides.

2. When recovering files and folders from a cloud-native snapshot, Veeam Backup for Google Cloud copies the persistent disks of the VM instance from the snapshot and attaches them to the worker instance.

When recovering files and folders from an image-level backup, the disks are not physically extracted from the backup — Veeam Backup for Google Cloud emulates their presence on the worker instance. The source backup itself remains in the read-only state.

3. Launches the File-Level Restore browser.

The File-Level Restore browser displays the directory structure of the backed-up VM instance. In the browser, you select the necessary files and folders to recover.

4. Saves the selected files and folders to the local machine.
5. Removes the worker instance from Google Cloud.

File-Level Recovery to Original Location

To recover files and folders of a backed-up VM instance to the original location, Veeam Backup for Google Cloud performs the following steps:

1. When recovering files and folders from a cloud-native snapshot, Veeam Backup for Google Cloud copies the persistent disks of the VM instance from the snapshot and attaches them to the worker instance.

When recovering files and folders from an image-level backup, the disks are not physically extracted from the backup — the source backup itself remains in the read-only state.

2. Deploys a worker instance within the worker project in either of the following Google Cloud regions.
 - o To recover files and folders from a cloud-native snapshot, the worker instance is deployed in the region in which the target VM instance resides.
 - o To recover files and folders from an image-level backup, the worker instance is deployed in the region in which the storage bucket with backed-up data resides.

3. When recovering files and folders from a cloud-native snapshot, Veeam Backup for Google Cloud attaches the copied persistent disks to the worker instance.

When recovering files and folders from an image-level backup, Veeam Backup for Google Cloud emulates disk presence on the worker instance.

4. [For Linux-operated instances] Generates an SSH key for *veeam_restore_user* and uploads the key to the target VM instance using Compute Engine API.

[For Windows-operated instances] Creates credentials for *veeam_restore_user* on the target VM instance using Compute Engine API.

5. Establishes an [encrypted IAP tunnel](#) between the backup appliance and the target VM instance to enable administrative access to the instance.

6. Creates a storage bucket with the *veeam-transfer-files-{GUID}* name in the region where the target VM instance resides, which is required to copy and launch the restore utilities.

Veeam Backup for Google Cloud will use the storage bucket created during the first file-level recovery session for all the subsequent recovery sessions — unless you delete the bucket from Google Cloud manually.

7. Launches the File-Level Restore browser.

The File-Level Restore browser displays the directory structure of the backed-up VM instance. In the browser, you select the necessary files and folders to recover.

8. Recovers the selected items to the target VM instance using the [Pub/Sub service](#).

9. Removes the restore utilities from the storage bucket.

10. Removes the worker instance from Google Cloud.

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

Protecting Cloud SQL Instances

To produce cloud-native snapshots and image-level backups of Cloud SQL instances, Veeam Backup for Google Cloud 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, and so on.

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

How to Protect Cloud SQL Instances

To create an SQL backup policy, complete the following steps:

1. [Check limitations and prerequisites.](#)
2. [Add service accounts.](#)
3. [Connect projects and folders.](#)
4. [Add backup repositories.](#)
5. [Configure worker instance settings.](#)
6. [Configure global retention and email notification settings.](#)
7. [Complete the Add Cloud SQL Policy wizard.](#)

SQL Backup

When processing a Cloud SQL instance added to a backup policy, Veeam Backup for Google Cloud can create a restore point for the instance and transfer the point directly to a backup repository, or can copy the instance to a staging server first, create a restore point and then transfer it to a repository.

Veeam Backup for Google Cloud performs Cloud SQL instance backup in the following way:

1. Creates a cloud-native snapshot of the processed Cloud SQL instance.
2. If you enable image-level backup for the backup policy, Veeam Backup for Google Cloud performs the following operations:
 - a. If you choose to perform backup using a staging server:
 - i. Launches a staging server instance in a Google Cloud region in which the source Cloud SQL instance resides.
 - ii. Reverts the staging server instance to the cloud-native snapshot created at step 1.

Note that the cloud-native snapshot used as a source for image-level backup is not a temporary snapshot – when the backup session completes, this snapshot remains in the snapshot chain and is deleted later according to the specified [policy scheduling settings](#).

IMPORTANT

Veeam Backup for Google Cloud launches staging server instances without public IP addresses. To allow backup operations to complete successfully, you must configure private services access for these instances manually, as described in [Google Cloud documentation](#).

- b. Deploys a worker instance within the worker project in the Google Cloud region in which the processed Cloud SQL instance resides.

By default, Veeam Backup for Google Cloud deploys worker instances with private IP addresses regardless of the network configurations specified for the processed Cloud SQL instances (that is, if a Cloud SQL instance has either a public IP address or both a public and a private IP address, the worker instance will still have a private IP address). However, you can add specific worker configurations. For more information, see [Managing Worker Instances](#).

IMPORTANT

If you plan to back up Cloud SQL instances, you must configure network access between the subnets of the worker instances and the subnets of the processed Cloud SQL instances. Alternatively, you can configure the worker instances to allow public IP access as described in section [Configuring Deployment Mode](#).

- c. Uses the worker instance to retrieve databases, views, triggers, stored procedures and users of the processed Cloud SQL instance, transfers the retrieved data to the target backup repository and stores the data in the native Veeam format.
 - d. Removes the staging server instance (if launched at step 2a).
 - e. When the backup session completes, removes the worker instance from Google Cloud.

3. If you enable the [backup archiving mechanism](#), Veeam Backup for Google Cloud performs the following operations:
 - a. Deploys a worker instance within the worker project in the Google Cloud region in which the target standard or nearline repository is located.

For more information on how to specify a project for worker instances, see [Managing Worker Configurations](#).
 - b. Retrieves data from the target standard or nearline repository, and transfers it to the target archive repository.
 - c. Removes the worker instance when the archive session completes.

Snapshot Chain

During every backup session, Veeam Backup for Google Cloud creates a cloud-native snapshot of each Cloud SQL instance added to a backup policy. The cloud-native snapshot itself is a single 'backup' that Veeam Backup for Google Cloud creates using native Google Cloud capabilities.

NOTE

Cloud-native snapshots of Cloud SQL instances are referred to as backups in Google Cloud documentation. However, since all 'backups' of a Cloud SQL instance are automatically deleted after you remove the instance itself, 'backups' of Cloud SQL instances are referred to as snapshots in this guide. In terms of Veeam logic, backups are independent files that are stored in backup repositories and that are not affected by any actions performed with the original instances whatsoever.

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

1. During the first backup session, Veeam Backup for Google Cloud creates a snapshot of all instance data, and saves it in the multi-regional location closest to the region in which the original 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 Google Cloud processes all the instance databases.

2. During subsequent backup sessions, Veeam Backup for Google Cloud 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 processed data.

For more information on how incremental Cloud SQL snapshots work, see [Google Cloud SQL documentation](#).

Each cloud-native snapshot in the snapshot chain contains metadata. Metadata stores information about the backup policy that created the snapshot, but does not contain any information about the protected instance.

Veeam Backup for Google Cloud uses metadata to identify outdated snapshots only; information about the protected instance is stored separately, in the internal Veeam Backup for Google Cloud database.

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.

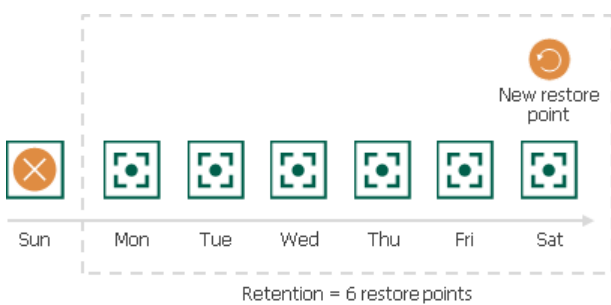


The number of cloud-native snapshots kept in the snapshot chain is defined by retention policy settings. For more information, see [SQL Snapshot Retention](#).

SQL Snapshot Retention

For cloud-native snapshots, Veeam Backup for Google Cloud retains the number of latest restore points defined in backup scheduling settings as described in section [Creating SQL Policies](#).

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



NOTE

Retention policy settings configured when creating backup policies do not apply to cloud-native snapshots created manually. To learn how to remove these snapshots, see [Removing Backups and Snapshots](#).

Backup Chain

If you enable image-level backups for a backup policy, Veeam Backup for Google Cloud creates a new backup in a standard or nearline repository during every backup session. A sequence of backups created during a set of backup sessions makes up a regular backup chain.

Each Cloud SQL backup in the backup chain contains metadata that stores information about the protected instance, the backup policy that created the backup, as well as the date, time and configured retention settings.

Veeam Backup for Google Cloud uses metadata to identify outdated backups, to retrieve information on the source instance configuration during recovery operations, and so on.

NOTE

The [forever forward incremental backup](#) method is not fully implemented for Cloud SQL instances — during every backup session, Veeam Backup for Google Cloud creates a full backup in the regular backup chain (that is, every incremental backup contains the full instance data set).

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

Archive Backup Chain

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


The archive backup chain includes backups of the following types:

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

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

1. During the first archive session, Veeam Backup for Google Cloud detects backed-up data that is stored in the full backup and all incremental 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 Google Cloud checks the regular 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 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.

Veeam Backup for Google Cloud creates incremental backups based on the Veeam proprietary filtering mechanism that filters out unchanged data blocks by calculating a checksum for every block. The Google Cloud changed block tracking (CBT) mechanism that would allow tracking changed blocks of data and would increase the efficiency of incremental backups is not implemented at the moment.



Full and incremental archive backups act as restore points for backed-up instances that let you roll back instance data to the necessary state. To recover an 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 backups from the archive repository manually. Instead, you must

specify retention policy settings that will let you maintain the necessary number of backups in the archive repository. For more information, see [Retention Policy for Archived Backups](#).

SQL Backup Retention

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


The forever forward incremental backup method is not fully implemented for Cloud SQL instances – during every backup session Veeam Backup for Google Cloud creates a full backup in the regular backup chain (that is, every incremental backup contains the full instance data set). If Veeam Backup for Google Cloud detects an outdated restore point in a backup repository, it removes this restore point from the backup chain.

Retention Policy for Archived Backups


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

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


1. Veeam Backup for Google Cloud checks the configuration database to detect archive backup repositories that contain outdated restore points.
2. If an outdated restore point exists in a backup repository, Veeam Backup for Google Cloud transforms the archive backup chain in the following way:
 - a. Rebuilds the full archive backup to include the data of the incremental archive backup that follows the full archive backup. To do that, Veeam Backup for Google Cloud 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.



- b. 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 Google Cloud 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 Google Cloud ensures that the archive backup chain is not broken and that you will be able to recover your data when needed.



4. Removes the worker instance when the retention session completes.

SQL Restore

Veeam Backup for Google Cloud offers the following restore operations:

- [Instance restore](#) — start an entire Cloud SQL instance from a restore point.
- [Database restore](#) — restore specific databases of a Cloud SQL instance.

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

Instance Restore

To restore a Cloud SQL instance from a cloud-native snapshot, Veeam Backup for Google Cloud first creates a Cloud SQL instance in the target location and then uses [native Google Cloud capabilities](#) to revert the instance to the snapshot. Restore of Cloud SQL instances from cloud-native snapshots is supported only to a new location or with different settings.

To restore a Cloud SQL instance from an image-level backup, Veeam Backup for Google Cloud performs the following steps:

1. Creates a Cloud SQL instance in the target location (that is, the project and region specified for the restore operation).
2. Deploys a worker instance within the worker project in the Google Cloud region in which the restored Cloud SQL instance will reside.

NOTE

Every time before creating a Cloud SQL instance, Veeam Backup for Google Cloud checks whether the total size of the instance breaches the zone quota for the project in which the worker instance is deployed. If the total instance size is less than 1000 GB, Veeam Backup for Google Cloud temporarily attaches an additional empty disk to the worker instance — but only for the duration of the restore process and if the quota allows attaching the disk. This allows Veeam Backup for Google Cloud to speed up the data transfer to reduce your restore costs.

For more information on how to specify a project for worker instances, see [Managing Worker Configurations](#).

3. Uses the worker instance to retrieve database files, triggers, views, stored procedures and users of the processed Cloud SQL instance from the backup file, and then imports this data to the created Cloud SQL instance.
4. Removes the worker instance from Google Cloud.

To learn how to restore a Cloud SQL instance from a cloud-native snapshot or an image-level backup, see [Performing SQL Instance Restore](#).

Database Restore

To restore a Cloud SQL database from an image-level backup, Veeam Backup for Google Cloud performs the following steps:

1. Creates a temporary snapshot of the Cloud SQL instance that will host the restored database.

2. Deploys a worker instance within the worker project in the Google Cloud region in which the target Cloud SQL instance resides.

For more information on how to specify a project for worker instances, see [Managing Worker Configurations](#).

3. [Applies only to restore points of MySQL instances protected by Veeam Backup for Google Cloud version 4.0 and earlier] Uses native Google Cloud capabilities to import the exported data to the target Cloud SQL instance.
4. Uses the worker instance to retrieve database files, triggers, stored procedures and users of the processed Cloud SQL instance from the backup file, and then imports this data to the target Cloud SQL instance.
5. Removes the worker instance from Google Cloud.
6. Deletes the temporary snapshot.

NOTE

If Veeam Backup for Google Cloud fails to restore the database, the temporary snapshot will not be deleted automatically. You can either delete the snapshot manually or use it to revert the Cloud SQL instance to its initial state. Consider that if the target instance hosts other databases and any write operations to these databases occur during the restore process, the revert operation will result in data loss.

To learn how to restore a Cloud SQL database from an image-level backup, see [Performing Database Restore](#).

Protecting Cloud Spanner Instances

To produce cloud-native snapshots and image-level backups of Cloud Spanner instances, Veeam Backup for Google Cloud 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, and so on.

Veeam Backup for Google Cloud does not install agent software inside instances to back up Cloud Spanner data – it uses native Google Cloud capabilities instead. During every backup session, Veeam Backup for Google Cloud creates a cloud-native snapshot of each database of a specific Cloud Spanner instance added to a backup policy. For more information on how Cloud Spanner instance backup works, see [Spanner Backup](#).

How to Protect Cloud Spanner Instances

To create a Spanner backup policy, complete the following steps:

1. [Check limitations and prerequisites.](#)
2. [Add service accounts.](#)
3. [Connect projects and folders.](#)
4. [Add backup repositories.](#)
5. [Configure worker instance settings.](#)
6. [Configure global retention and email notification settings.](#)
7. [Complete the Add Cloud Spanner Policy wizard.](#)

Spanner Backup

When processing a Cloud Spanner instance added to a backup policy, Veeam Backup for Google Cloud creates a restore point for the instance and transfers the point directly to a backup repository.

Veeam Backup for Google Cloud performs Cloud Spanner instance backup in the following way:

1. Creates a cloud-native snapshot of each database of the processed Cloud Spanner instance.
2. If you enable image-level backup for the backup policy, Veeam Backup for Google Cloud performs the following operations:
 - a. Deploys a worker instance within the worker project in the Google Cloud region depending both on the target backup repository location and the region where read-write and read-only replicas reside.
 - b. Uses the worker instance to retrieve database schema, views, keys and data of the processed Cloud Spanner instance, transfers the retrieved data to the target backup repository and stores the data in the native Veeam format.
 - c. When the backup session completes, Veeam Backup for Google Cloud removes the worker instance from Google Cloud.
3. If you enable the [backup archiving mechanism](#), Veeam Backup for Google Cloud performs the following operations:
 - a. Deploys a worker instance within the worker project in the Google Cloud region in which the target standard or nearline repository is located.

For more information on how to specify a project for worker instances, see [Managing Worker Configurations](#).
 - b. Retrieves data from the target standard or nearline repository, and transfers it to the target archive repository.
 - c. Removes the worker instance when the archive session completes.

Snapshot Chain

During every backup session, Veeam Backup for Google Cloud creates a 'backup' of each database of a Cloud Spanner instance added to a backup policy. The set of 'backups' is a single cloud-native snapshot that Veeam Backup for Google Cloud creates using native Google Cloud capabilities.

NOTE

Cloud-native snapshots of Cloud Spanner instances are referred to as backups in Google Cloud documentation. However, since all 'backups' of a Cloud Spanner instance are stored in the instance itself, and you cannot delete an instance without deleting its snapshots first, 'backups' of Cloud Spanner instances are referred to as snapshots in this guide. In terms of Veeam logic, backups are independent files that are stored in backup repositories and that are not affected by any actions performed with the original instances whatsoever.

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

1. During the first backup session, Veeam Backup for Google Cloud creates snapshots of all the instance databases and saves them in the target location. These snapshots become a starting point in the snapshot chain.

The creation of the first snapshot may take significant time to complete since Veeam Backup for Google Cloud processes all the instance databases.

2. During subsequent backup sessions, Veeam Backup for Google Cloud creates snapshots as described in the [Google Cloud Spanner documentation](#).


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 processed data.

NOTE

The target location of cloud-native snapshots depends on the regional configuration of the processed instance. For more information, see [Google Cloud documentation](#).

Each cloud-native snapshot in the snapshot chain contains metadata. Metadata stores information about the protected instance and the backup policy that created the snapshot. Veeam Backup for Google Cloud uses metadata to identify outdated snapshots, 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 the snapshot chain is defined by retention policy settings. For more information, see [Spanner Snapshot Retention](#).

Spanner Snapshot Retention

For cloud-native snapshots, Veeam Backup for Google Cloud retains the number of latest restore points defined in backup scheduling settings as described in section [Creating Spanner Policies](#).

During every successful backup session, Veeam Backup for Google Cloud creates a new restore point. If Veeam Backup for Google Cloud detects that the number of restore points in the snapshot chain exceeds the retention

limit, it removes the earliest restore point from the chain. For more information on the snapshot deletion process, see [Google Cloud documentation](#).



NOTE

Retention policy settings configured when creating backup policies do not apply to cloud-native snapshots created manually. To learn how to remove these snapshots, see [Removing Backups and Snapshots](#).

Backup Chain

If you enable image-level backups for a backup policy, Veeam Backup for Google Cloud creates a new backup in a standard or nearline repository during every backup session. A sequence of backups created during a set of backup sessions makes up a regular backup chain.

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

NOTE

The [forever forward incremental backup](#) method is not implemented for Cloud Spanner instances – during every backup session Veeam Backup for Google Cloud creates a full backup in the regular backup chain.

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

Archive Backup Chain


The archive backup chain includes backups of the following types:

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

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

1. During the first archive session, Veeam Backup for Google Cloud detects backed-up data that is stored in the full backup 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 Google Cloud checks the regular backup chain to detect data blocks that have changed since the previous archive session, creates incremental archive backups with only those changed data blocks, and copies these backups to the archive 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 instances that let you roll back instance data to the necessary state. To recover an 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 backups from the archive repository manually. Instead, you must specify retention policy settings that will let you maintain the necessary number of backups in the archive repository. For more information, see [Retention Policy for Archived Backups](#).

Spanner Backup Retention

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

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

Retention Policy for Archived Backups


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

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


1. Veeam Backup for Google Cloud checks the configuration database to detect archive backup repositories that contain outdated restore points.

2. If an outdated restore point exists in a backup repository, Veeam Backup for Google Cloud transforms the archive backup chain in the following way:

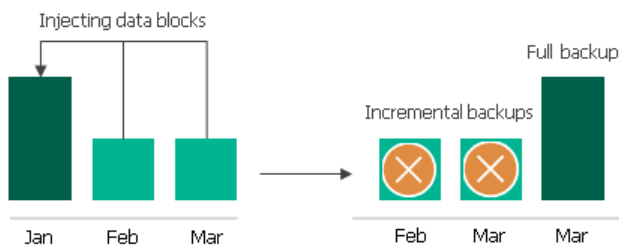
a. Rebuilds the full archive backup to include the data of the incremental archive backup that follows the full archive backup. To do that, Veeam Backup for Google Cloud injects into the full archive backup data blocks from the earliest incremental archive backup in the chain. This way, the full archive backup 'moves' forward in the archive backup chain.



b. Removes the earliest incremental archive backup from the chain as redundant – this data has already been injected into the full archive backup.



4. Veeam Backup for Google Cloud 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 Google Cloud ensures that the archive backup chain is not broken and that you will be able to recover your data when needed.



5. Removes the worker instance when the retention session completes.

Spanner Restore

Veeam Backup for Google Cloud offers the following restore operations:

- [Instance restore](#) — start an entire Cloud Spanner instance from a restore point.
- [Database restore](#) — restore specific databases of a Cloud Spanner instance.

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

Instance Restore

To restore a Cloud Spanner instance from a cloud-native snapshot, Veeam Backup for Google Cloud first creates a Cloud Spanner instance in the target location, copies database snapshots from the source instance, and then uses [native Google Cloud capabilities](#) to revert the databases to their snapshots. Restore of Cloud Spanner instances from cloud-native snapshots is supported only to a new location or with different settings.

To restore a Cloud Spanner instance from an image-level backup, Veeam Backup for Google Cloud performs the following steps:

1. Creates a Cloud Spanner instance with default databases in the target location (that is, the project and region specified for the restore operation).
2. Deploys a worker instance within the worker project in the Google Cloud region in which the target instance resides or the region in which the read-write replicas are located.

For more information on how to specify a project for worker instances, see [Managing Worker Configurations](#).

3. Uses the worker instance to retrieve database schema and data of the processed Cloud Spanner instance from the backup file, and then imports this data to the created Cloud Spanner instance.

If the processed instance contains databases with foreign keys, Veeam Backup for Google Cloud will first restore the database schema without foreign keys, then import the database data, and then restore the foreign keys.

4. Removes the worker instance from Google Cloud.

To learn how to restore a Cloud Spanner instance from a cloud-native snapshot or an image-level backup, see [Performing Spanner Instance Restore](#).

Database Restore

To restore a Cloud Spanner database from a cloud-native snapshot, Veeam Backup for Google Cloud copies database snapshots from the source instance, and then uses [native Google Cloud capabilities](#) to revert the databases to their snapshots.

To restore a Cloud Spanner database from an image-level backup, Veeam Backup for Google Cloud performs the following steps:

1. Creates default databases on the target Cloud Spanner instance.
2. Deploys a worker instance within the worker project in the Google Cloud region closest to the region where the target Cloud Spanner instance resides.

For more information on how to specify a project for worker instances, see [Managing Worker Configurations](#).

3. Uses the worker instance to retrieve database schema and data of the processed Cloud Spanner instance from the backup file, and then transfers this data to the target Cloud Spanner instance.

If the processed instance contains databases with foreign keys, Veeam Backup for Google Cloud will first restore the database schema without foreign keys, then transfer the database data, and then restore the foreign keys.

4. Removes the worker instance from Google Cloud.

To learn how to restore a Cloud Spanner database from an image-level backup or a cloud-native snapshot, see [Performing Database Restore](#).

Retention Policies

Cloud-native snapshots and image-level backups created by backup policies are not kept forever — they are removed according to retention policy settings specified while creating the policies as described in sections [Creating VM Policies](#), [Creating SQL Policies](#) and [Creating Spanner Policies](#).

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

- **In restore points** — for cloud-native snapshots.

The snapshot chain can contain only the allowed number of restore points. If the number of allowed restore points is exceeded, Veeam Backup for Google Cloud removes the earliest restore point from the snapshot chain. For more information, see [VM Snapshot Retention](#), [SQL Snapshot Retention](#) and [Spanner Snapshot Retention](#).

- **In days/months/years** — for image-level backups.

Restore points in the backup chain (either regular or archive) can be stored in the configuration database only for the allowed period of time. If a restore point is older than the specified time limit, Veeam Backup for Google Cloud removes it from the backup chain. For more information, see sections [VM Backup Retention](#), [SQL Backup Retention](#) and [Spanner Backup Retention](#).

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

Data Encryption

For enhanced data security, Veeam Backup for Google Cloud allows you to encrypt backed-up data stored in Google Cloud storage buckets using Veeam encryption mechanisms. Additionally, Veeam Backup for Google Cloud supports native Google Cloud encryption – Google Cloud Key Management Service (Cloud KMS) [customer-managed encryption keys \(CMEKs\)](#).

IMPORTANT

[Customer-supplied encryption keys \(CSEKs\)](#) are not supported.

Storage Bucket Encryption

Veeam Backup for Google Cloud encrypts backups stored in storage buckets the same way Veeam Backup & Replication encrypts backups stored in backup repositories. To learn what algorithms Veeam Backup & Replication uses to encrypt backups, see the Veeam Backup & Replication User Guide, section [Encryption Standards](#).

To enable encryption for a backup repository added to Veeam Backup for Google Cloud, configure the repository settings as described in section [Adding Backup Repositories](#). After you create a backup policy and specify the backup repository as a target location for image-level backups, as described in sections [Performing VM Backup](#), [Performing SQL Backup](#), and [Performing Spanner Backup](#), Veeam Backup for Google Cloud performs the following steps:

1. Generates an encryption key to protect backups 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.

Cloud KMS Encryption

Veeam Backup for Google Cloud allows you to back up and restore data of encrypted Cloud Spanner databases, Cloud SQL instances and VM instances whose persistent disks are encrypted with Google Cloud KMS. Additionally, you can choose to encrypt data with original CMEKs or change CMEKs used to encrypt data when performing the following operations:

- [Restoring entire VM instances to a new location](#)
- [Restoring persistent disks of VM instances to a new location](#)
- [Restoring entire Cloud SQL instances to a new location](#)
- [Restoring entire Cloud Spanner instances to a new location](#)
- [Restoring specific Cloud Spanner databases to a new location](#)

Depending on the operation performed for an encrypted Cloud Spanner databases, Cloud SQL instance or a VM instance that has encrypted persistent disks, the service account that Veeam Backup for Google Cloud uses for the operation may require specific permissions to access Google Cloud KMS resources:

- [Creating cloud-native snapshots](#)
- [Creating image-level backups](#)

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

Creating Cloud-Native Snapshots

The process of creating cloud-native snapshots of an encrypted Cloud Spanner databases, Cloud SQL instance or a VM instance with encrypted persistent disks does not differ from the same process for an unencrypted Cloud Spanner instance, Cloud SQL instance or a VM instance with unencrypted persistent disks. The service account used to encrypt the created snapshots does not require any additional permissions — Veeam Backup for Google Cloud encrypts these snapshots with the same CMEKs with which the source Cloud SQL instance, databases of the source Cloud Spanner instance or persistent disks of the source VM instance are encrypted.

Creating Image-Level Backups

The process of creating image-level backups of a Cloud Spanner instance with encrypted databases, an encrypted Cloud SQL instance or a VM instance with encrypted persistent disks does not depend on the location where the worker instance processing the data is deployed. Regardless of whether the worker instance is deployed in the same Google Cloud project to which the source Cloud Spanner, Cloud SQL or VM instance belongs, Veeam Backup for Google Cloud performs the following steps:

- To back up a Cloud Spanner instance:
 - a. Takes a cloud-native snapshot of the Cloud Spanner instance.
 - b. Uses the worker instance to retrieve databases, views, tables and foreign keys of the processed Cloud Spanner instance, transfers the retrieved data to the target backup repository and stores the data in the native Veeam format.

The service account that is used to retrieve data from the Cloud Spanner instance requires permissions to access CMEKs with which the source Cloud Spanner database is encrypted.
 - c. Removes the worker instance from Google Cloud.
- To back up a Cloud SQL instance:
 - a. Takes a cloud-native snapshot of the Cloud SQL instance.
 - b. Uses the worker instance to export databases, triggers, stored procedures and users of the Cloud SQL instance to the target backup repository.

The service account that is used to retrieve the data requires permissions to access CMEKs with which the source Cloud SQL instance is encrypted.
 - c. Removes the worker instance from Google Cloud.
- To back up a VM instance:
 - a. Takes a cloud-native snapshot of the VM instance.
 - b. Creates persistent disks from the snapshot.

To encrypt the created disks, Veeam Backup for Google Cloud requires permissions of a service account that can access CMEKs with which you want to encrypt these disks.
 - c. Attaches the created persistent disks to the worker instance to read and further transfer the backed-up data to a backup repository.
 - d. Removes the worker instance from Google Cloud.

NOTE

Every time before creating persistent disks from a cloud-native snapshot, Veeam Backup for Google Cloud checks whether the total size of pd-standard disks breaches the zone quota for the project in which the worker instance is deployed. If the total disk size is less than 4000 GB, Veeam Backup for Google Cloud temporarily attaches an additional empty disk to the worker instance — but only for the duration of the backup process and if the quota allows attaching the disk. This allows Veeam Backup for Google Cloud to speed up the data transfer to reduce your backup costs.

Restoring from Cloud-Native Snapshots

The process of restoring a Cloud Spanner, Cloud SQL or VM instance from an encrypted cloud-native snapshot does not differ depending on the location where the restored instance will reside. Regardless of whether the Cloud Spanner, Cloud SQL or VM instance will be restored to the same Google Cloud project to which the cloud-native snapshot belongs, Veeam Backup for Google Cloud performs the following steps:

- To restore a Cloud Spanner instance:
 - a. Creates a Cloud Spanner instance in the target location.

To encrypt the databases of the created instance, Veeam Backup for Google Cloud requires permissions of a service account that can access the CMEK with which you want to encrypt these databases.
 - b. Copies the snapshot of the source Cloud Spanner instance to the target Cloud Spanner instance, and then restores databases from the snapshot to the target instance.
- To restore a Cloud SQL instance:
 - a. Creates a Cloud SQL instance in the target location.

The service account that is used to create the instance requires permissions to access the CMEK with which you want to encrypt this instance.
 - b. Uses native Google Cloud capabilities to revert the created Cloud SQL instance to the snapshot.
- To restore a VM instance:
 - a. Creates persistent disks from the cloud-native snapshot.

To encrypt the created disks, Veeam Backup for Google Cloud requires permissions of a service account that can access the CMEK with which you want to encrypt these disks.
 - b. Creates a VM instance in the target location.
 - c. Attaches the created persistent disks with the restored data to the VM instance.

Restoring from Image-Level Backups

The process of restoring a Cloud Spanner instance with encrypted databases, an encrypted Cloud SQL instance or a VM instance with encrypted persistent disks from an image-level backup does not differ depending on the location where the worker instance processing the data is deployed. Regardless of whether the worker instance

is deployed in the same Google Cloud project to which the restored Cloud Spanner, Cloud SQL or VM instance will belong, Veeam Backup for Google Cloud performs the following steps:

- To restore a Cloud Spanner instance:
 - a. Creates a Cloud Spanner instance in the target location.

To encrypt the databases of the created instance, Veeam Backup for Google Cloud requires permissions of a service account that can access the CMEK with which you want to encrypt these databases.
 - b. Uses the worker instance to transfer database schema, data and foreign keys of the backed-up Cloud Spanner instance to the target instance.
 - c. Removes the worker instance from Google Cloud.
- To restore a Cloud SQL instance:
 - a. Creates a Cloud SQL instance in the target location.

The service account that is used to create the instance requires permissions to access the CMEK with which you want to encrypt this instance.
 - b. Uses the worker instance to transfer databases, triggers, stored procedures and users of the backed-up Cloud SQL instance to the target instance.
 - c. Removes the worker instance from Google Cloud.
- To restore a VM instance:
 - a. Creates empty persistent disks and attaches the disks to the worker instance to restore the backed-up data to the target location.

To encrypt the created disks, Veeam Backup for Google Cloud requires permissions of a service account that can access the CMEK with which you want to encrypt these disks.
 - b. Takes cloud-native snapshots of the persistent disks with the restored data.
 - c. Creates a VM instance in the target location.
 - d. Creates persistent disks from the snapshots, and attaches the disks to the VM instance.

To encrypt the created disks, Veeam Backup for Google Cloud requires permissions of a service account that can access the CMEK with which you want to encrypt these disks.
 - e. Removes the worker instance from Google Cloud.

Planning and Preparation

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

- [Hardware and software requirements](#)
- [Network ports that must be open for data transmission](#)
- [Permissions that must be granted to accounts used for operations started from the Veeam Backup & Replication console](#)
- [Permissions that must be granted to accounts used for operations started from the Veeam Backup for Google Cloud Web UI](#)
- [Google Cloud APIs to which Veeam Backup for Google Cloud must have outbound internet access](#)
- [Considerations and limitations that should be kept in mind before you deploy Veeam Backup for Google Cloud](#)

System Requirements

When you plan to install Veeam Backup for Google Cloud, consider the following hardware and software requirements.

Google Cloud Plug-in for Veeam Backup & Replication

The machine where Google Cloud 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 or later
- Microsoft ASP.NET Core Shared Framework 6.0.24 or later

Backup Server

Google Cloud Plug-in for Veeam Backup & Replication version 5.0 supports integration with Veeam Backup & Replication version 12.1.

Backup Appliance

Google Cloud Plug-in for Veeam Backup & Replication version 5.0 supports integration with Veeam Backup for Google Cloud version 5.x.

Google Cloud APIs

The backup appliance and worker instances must have outbound internet access to a number of Google Cloud APIs. For more information, see [Google Cloud APIs](#).

Web Browsers

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

Ports

As Google Cloud 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, Google Cloud Plug-in for Veeam Backup & Replication also uses ports listed in the following table.

From	To	Protocol	Port	Description
Workstation web browser	Backup appliance	TCP/HTTPS	443	Required to access the Web UI component from a user workstation.
		TCP/HTTPS	13140	Required to communicate with the REST API service running on the backup appliance.
	Worker instance	TCP/HTTPS	443	Required to access the File-Level Restore browser running on a worker instance during the file-level recovery process.
Backup appliance	Ubuntu Security Update Repository (security.ubuntu.com)	TCP/HTTP	80	Required to get OS security updates.
	Veeam Update Notification Server (repository.veeam.com)	TCP/HTTPS	443	Required to download the necessary 3rd party software upon the installation, and further to download information on available product updates.
	SMTP server	TCP	587	Required to send email notifications. Note: You cannot use the TCP port 25 that is most commonly used by SMTP servers – the port is always blocked by Google Compute Engine. For more information, see Google Cloud documentation .
	nginx web server (nginx.org)	HTTPS	80/443	Required to upgrade the backup appliance.
	PostgreSQL repository (apt.postgresql.org)	HTTPS	80/443	Required to get PostgreSQL updates.
	Microsoft Repository (packages.microsoft.com)	HTTPS	80/443	Required to get .NET and ASP.NET updates.

From	To	Protocol	Port	Description
Google Cloud Plug-in for Veeam Backup & Replication	Backup appliance, Google Cloud services	TCP/HTTPS	443	Required to communicate with Google Cloud.
	Backup server	TCP	6172	Required to connect to a component that enables communication with the Veeam Backup & Replication database.
Veeam Backup & Replication console and Veeam ONE server	Backup server	TCP	9403	Required to connect to Google Cloud Plug-in for Veeam Backup & Replication.
Worker instance	Google Cloud services	TCP/HTTPS	443	Required to perform data protection and disaster recovery operations.
	Cloud SQL instances	TCP	3306	
		TCP	5432	
	Cloud Spanner instances	TCP	443	

NOTE

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

Plug-In Permissions

To perform backup and restore operations, accounts that Google Cloud 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 use when installing and working with Veeam Backup & Replication must have the permissions listed in the Veeam Backup & Replication User Guide, section [Installing and Using Veeam Backup & Replication](#).

Veeam Backup for Google Cloud User Account Permissions

A user account that Veeam Backup & Replication uses to authenticate against a 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 automatically creates the necessary user account that is assigned all the required permissions.

Google Cloud Service Account Permissions

Google Cloud Plug-in for Veeam Backup & Replication requires the following service accounts:

- A service account whose permissions are used to create, connect and manage backup appliances. You can create this account manually in Google Cloud or instruct Veeam Backup & Replication to create the account automatically.

If you instruct Veeam Backup & Replication to create the service account automatically, the account is assigned the Owner role with a wide scope of permissions and capabilities. If you create a new service account in Google Cloud manually, consider that the service account must have the following minimal set of permissions:

> List of permissions

```
{
  compute.addresses.list
  compute.disks.create
  compute.disks.createSnapshot
  compute.disks.delete
  compute.disks.get
  compute.disks.setLabels
  compute.disks.use
  compute.firewalls.list
  compute.globalOperations.get
  compute.instances.attachDisk
  compute.instances.detachDisk
  compute.instances.get
  compute.instances.getGuestAttributes
  compute.instances.list
  compute.instances.setMetadata
  compute.instances.start
  compute.instances.stop
  compute.networks.get
  compute.networks.list
  compute.projects.get
  compute.regions.get
  compute.regions.list
  compute.snapshots.create
  compute.snapshots.delete
  compute.snapshots.get
  compute.snapshots.useReadOnly
  compute.subnetworks.get
  compute.subnetworks.list
  compute.zoneOperations.get
  compute.zones.get
  compute.zones.list
  compute.machineTypes.list
  deploymentmanager.deployments.create
  deploymentmanager.deployments.delete
  deploymentmanager.deployments.get
  deploymentmanager.operations.get
  deploymentmanager.resources.list
  iam.roles.create
  iam.serviceAccounts.actAs
  iap.tunnelInstances.accessViaIAP
  iam.serviceAccounts.actAs
  resourcemanager.projects.getIamPolicy
  resourcemanager.projects.setIamPolicy
  storage.buckets.create
}
```

After you create a service account in Google Cloud, you must add it to Veeam Backup & Replication as described in the Veeam Backup & Replication User Guide, section [Google Cloud Platform Service Account](#).

- A service account whose permissions are used to perform data protection and disaster recovery operations with Google Cloud resources.
 - When you deploy a new backup appliance, the default service account is automatically created on this appliance and is assigned all the required permissions.
 - When you connect to an existing backup appliance, Google Cloud Plug-in for Veeam Backup & Replication uses a service account with a set of predefined permissions that has already been created on this appliance.

Virtualization Servers and Hosts Service Account Permissions

If you plan to copy backups to on-premises 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 VM instances to Microsoft Azure must have permissions described in the Veeam Backup & Replication User Guide, section [Permissions](#).

AWS IAM User Permissions

An IAM user whose one-time access keys you plan to use to perform restore of VM instances to Amazon EC2 must have permissions described in the Veeam Backup & Replication User Guide, section [AWS IAM User Permissions](#).

Service Account Permissions

Google Cloud Identity and Access Management (IAM) roles that Veeam Backup for Google Cloud uses to perform data protection and disaster recovery operations must have permissions to access Google Cloud services and resources.

Default Permissions

Veeam Backup for Google Cloud requires a service account in each Google Cloud project where data protection and disaster recovery tasks will be performed. To allow Veeam Backup for Google Cloud to access Google Cloud services and resources that you want to protect, service accounts used by Veeam Backup for Google Cloud must have the following minimal set of permissions:

```
compute.disks.addResourcePolicies
compute.disks.get
compute.instances.get
compute.resourcePolicies.create
compute.resourcePolicies.get
compute.resourcePolicies.use
compute.zones.get
serviceusage.services.list
compute.projects.get
resourcemanager.projects.get
```

Repository Permissions

To allow Veeam Backup for Google Cloud to create a backup repository in a Google Cloud storage bucket and to access the repository when performing backup and restore operations, the service account associated with the Google Cloud project in which this bucket resides must have the following permissions:

```
storage.buckets.list
storage.buckets.get
storage.objects.create
storage.objects.delete
storage.objects.list
storage.objects.get
storage.hmacKeys.create
storage.hmacKeys.list
storage.hmacKeys.get
resourcemanager.projects.get
serviceusage.services.list
storage.buckets.getIamPolicy
storage.buckets.setIamPolicy *
compute.projects.get
storage.multipartUploads.create
storage.multipartUploads.abort
```

* Veeam Backup for Google Cloud will use the `storage.buckets.setIamPolicy` permission only to grant access to repositories while performing SQL backup operations.

Worker Permissions

To allow Veeam Backup for Google Cloud to create a worker instance in a Google Cloud project and to access the instance when performing backup and restore operations, the service account associated with the project must have the following permissions:

VM Backup and Restore Permissions

```
compute.regions.list
compute.disks.list
compute.instances.get
compute.instances.list
compute.snapshots.get
compute.snapshots.list
compute.zones.get
compute.zones.list
compute.globalOperations.get
compute.zoneOperations.get
compute.regionOperations.get
resourcemanager.projects.get
compute.projects.get
compute.firewalls.list
compute.snapshots.getIamPolicy
compute.networks.list
compute.subnetworks.list
resourcemanager.projects.getIamPolicy
resourcemanager.projects.setIamPolicy *
iam.serviceAccounts.actAs
compute.disks.create
compute.disks.createSnapshot
compute.disks.delete
compute.disks.setLabels
compute.instances.attachDisk
compute.instances.create
compute.instances.delete
compute.instances.detachDisk
compute.instances.setMetadata
compute.instances.setServiceAccount
compute.instances.setLabels
compute.instances.setTags
compute.routes.list
compute.regions.get
compute.snapshots.create
compute.snapshots.setLabels
compute.snapshots.setIamPolicy
compute.snapshots.delete
pubsub.subscriptions.consume
pubsub.subscriptions.create
pubsub.subscriptions.delete
pubsub.subscriptions.list
pubsub.subscriptions.get
logging.sinks.get
logging.sinks.delete
logging.sinks.list
pubsub.topics.attachSubscription
pubsub.topics.detachSubscription
pubsub.topics.create
pubsub.topics.delete
pubsub.topics.list
pubsub.topics.get
```

```
pubsub.topics.publish
compute.machineTypes.get
compute.machineTypes.list
compute.subnetworks.get
compute.subnetworks.use
compute.subnetworks.useExternalIp
compute.disks.use
pubsub.subscriptions.setIamPolicy
pubsub.subscriptions.getIamPolicy
pubsub.topics.setIamPolicy
pubsub.topics.getIamPolicy
storage.objects.create
storage.objects.delete
storage.objects.list
storage.objects.get
storage.objects.update
storage.buckets.create
serviceusage.services.list
```

IMPORTANT

- To allow Veeam Backup for Google Cloud to perform restore to the original location while source VM instances still exist there, the deletion protection setting must be disabled for the source instance, and the service account whose permissions are used for the operation must be assigned the `compute.instances.setName` permission.
- To allow Veeam Backup for Google Cloud to connect a created worker instance to a Shared VPC network, the service account associated with the Google Cloud project to which the instance belongs must also have either the `compute.networkUser` role for the whole Shared VPC host project, or the `compute.networkViewer` role for the whole host project plus `compute.networkUser` for specific subnets in the host project. To learn how to provide access to Shared VPC networks, see [Google Cloud documentation](#).

Cloud SQL Backup and Restore Permissions

```
compute.regions.list
compute.disks.list
compute.instances.get
compute.instances.list
compute.snapshots.get
compute.snapshots.list
compute.zones.get
compute.zones.list
compute.globalOperations.get
compute.zoneOperations.get
compute.regionOperations.get
resourcemanager.projects.get
compute.projects.get
compute.firewalls.list
compute.snapshots.getIamPolicy
compute.networks.list
compute.subnetworks.list
resourcemanager.projects.getIamPolicy
resourcemanager.projects.setIamPolicy *
iam.serviceAccounts.actAs
compute.disks.create
compute.disks.createSnapshot
compute.disks.delete
compute.disks.setLabels
compute.instances.attachDisk
compute.instances.create
compute.instances.delete
compute.instances.detachDisk
compute.instances.setMetadata
compute.instances.setServiceAccount
compute.instances.setLabels
compute.instances.setTags
compute.routes.list
compute.regions.get
compute.snapshots.create
compute.snapshots.setLabels
compute.snapshots.setIamPolicy
compute.snapshots.delete
pubsub.subscriptions.consume
pubsub.subscriptions.create
pubsub.subscriptions.delete
pubsub.subscriptions.list
pubsub.subscriptions.get
logging.sinks.get
logging.sinks.delete
logging.sinks.list
pubsub.topics.attachSubscription
pubsub.topics.detachSubscription
pubsub.topics.create
pubsub.topics.delete
pubsub.topics.list
pubsub.topics.get
pubsub.topics.publish
```



```
compute.machineTypes.get
compute.machineTypes.list
compute.subnetworks.get
compute.subnetworks.use
compute.subnetworks.useExternalIp
compute.disks.use
serviceusage.services.list
cloudsql.databases.list
cloudsql.instances.create
cloudsql.instances.delete
cloudsql.instances.export
cloudsql.instances.get
cloudsql.instances.list
cloudsql.instances.listServerCas
cloudsql.users.create
cloudsql.users.list
cloudsql.users.update
compute.projects.get
```

* Veeam Backup for Google Cloud will use the `resourcemanager.projects.setIamPolicy` permission only to assign the `cloudsql.instances.get` and `cloudsql.instances.restoreBackup` permissions to service accounts while performing backup operations.

Cloud Spanner Backup and Restore Permissions

```
compute.regions.list
compute.disks.list
compute.instances.get
compute.instances.list
compute.snapshots.get
compute.snapshots.list
compute.zones.get
compute.zones.list
compute.globalOperations.get
compute.zoneOperations.get
compute.regionOperations.get
resourcemanager.projects.get
compute.projects.get
compute.firewalls.list
compute.snapshots.getIamPolicy
compute.networks.list
compute.subnetworks.list
resourcemanager.projects.getIamPolicy
iam.serviceAccounts.actAs
compute.disks.create
compute.disks.createSnapshot
compute.disks.delete
compute.disks.setLabels
compute.instances.attachDisk
compute.instances.create
compute.instances.delete
compute.instances.detachDisk
compute.instances.setMetadata
compute.instances.setServiceAccount
compute.instances.setLabels
compute.instances.setTags
compute.routes.list
compute.regions.get
compute.snapshots.create
compute.snapshots.setLabels
compute.snapshots.setIamPolicy
compute.snapshots.delete
pubsub.subscriptions.consume
pubsub.subscriptions.create
pubsub.subscriptions.delete
pubsub.subscriptions.list
pubsub.subscriptions.get
logging.sinks.get
logging.sinks.delete
logging.sinks.list
pubsub.topics.attachSubscription
pubsub.topics.detachSubscription
pubsub.topics.create
pubsub.topics.delete
pubsub.topics.list
pubsub.topics.get
pubsub.topics.publish
```

```
compute.machineTypes.get  
compute.machineTypes.list  
compute.subnetworks.get  
compute.subnetworks.use  
compute.subnetworks.useExternalIp  
compute.disks.use  
serviceusage.services.list
```

Snapshot Permissions

To allow Veeam Backup for Google Cloud to create and manage cloud-native snapshots of Google Cloud instances, the service account associated with the Google Cloud project managing instances that you want to protect must have the following permissions.

VM Snapshot Permissions

```
compute.addresses.list
compute.firewalls.list
compute.regions.list
compute.disks.list
compute.disks.createSnapshot
compute.disks.get
compute.instances.get
compute.instances.list
compute.networks.list
compute.projects.get
compute.snapshots.create
compute.snapshots.delete
compute.snapshots.get
compute.snapshots.list
compute.snapshots.getIamPolicy
compute.snapshots.setIamPolicy
compute.snapshots.setLabels
compute.subnetworks.list
compute.routes.list
compute.zones.list
compute.globalOperations.list
compute.globalOperations.get
compute.zoneOperations.get
compute.regionOperations.get
resourcemanager.projects.get
logging.sinks.create
logging.sinks.delete
logging.sinks.get
logging.sinks.list
logging.sinks.update
pubsub.subscriptions.create
pubsub.subscriptions.delete
pubsub.subscriptions.get
pubsub.subscriptions.list
pubsub.subscriptions.consume
pubsub.topics.attachSubscription
pubsub.topics.create
pubsub.topics.delete
pubsub.topics.detachSubscription
pubsub.topics.get
pubsub.topics.getIamPolicy
pubsub.topics.list
pubsub.topics.setIamPolicy
pubsub.topics.update
cloudkms.keyRings.list
cloudkms.cryptoKeys.list
serviceusage.services.list
```

Cloud SQL Snapshot Permissions

```
cloudsql.backupRuns.create
cloudsql.backupRuns.delete
cloudsql.backupRuns.get
cloudsql.backupRuns.list
cloudsql.databases.list
cloudsql.instances.get
cloudsql.instances.list
compute.regions.list
compute.zones.list
logging.sinks.create
logging.sinks.delete
logging.sinks.get
logging.sinks.list
pubsub.subscriptions.consume
pubsub.subscriptions.create
pubsub.subscriptions.delete
pubsub.subscriptions.get
pubsub.subscriptions.list
pubsub.topics.attachSubscription
pubsub.topics.create
pubsub.topics.delete
pubsub.topics.detachSubscription
pubsub.topics.get
pubsub.topics.getIamPolicy
pubsub.topics.list
pubsub.topics.setIamPolicy
serviceusage.services.list
cloudkms.keyRings.list
cloudkms.cryptoKeys.list
compute.projects.get
resourcemanager.projects.get
```

Cloud Spanner Snapshot Permissions

```
spanner.backups.copy,  
spanner.backups.create  
spanner.backups.get  
spanner.backups.list  
spanner.backups.delete  
spanner.backupOperations.cancel  
spanner.backupOperations.get  
spanner.backupOperations.list  
spanner.databases.createBackup  
spanner.databases.list  
spanner.instanceConfigs.get  
spanner.instanceConfigs.list  
spanner.instances.get  
spanner.instances.list  
compute.regions.list  
compute.zones.list  
logging.sinks.create  
logging.sinks.delete  
logging.sinks.get  
logging.sinks.list  
pubsub.subscriptions.consume  
pubsub.subscriptions.create  
pubsub.subscriptions.delete  
pubsub.subscriptions.get  
pubsub.subscriptions.list  
pubsub.topics.attachSubscription  
pubsub.topics.create  
pubsub.topics.delete  
pubsub.topics.detachSubscription  
pubsub.topics.get  
pubsub.topics.getIamPolicy  
pubsub.topics.list  
pubsub.topics.setIamPolicy,  
serviceusage.services.list  
cloudkms.keyRings.list  
cloudkms.cryptoKeys.list  
compute.projects.get  
resourcemanager.projects.get
```


Backup Permissions

To allow Veeam Backup for Google Cloud to perform backup operations, the service account associated with the Google Cloud project managing instances that you want to protect must have the following permissions.

VM Backup Permissions

```
compute.addresses.list
compute.regions.list
compute.disks.list
compute.disks.createSnapshot
compute.disks.get
compute.instances.get
compute.instances.list
compute.snapshots.create
compute.snapshots.delete
compute.snapshots.get
compute.snapshots.list
compute.snapshots.getIamPolicy
compute.snapshots.setIamPolicy
compute.snapshots.setLabels
compute.subnetworks.list
compute.routes.list
compute.machineTypes.get
compute.zones.list
compute.globalOperations.list
compute.globalOperations.get
compute.zoneOperations.get
compute.regionOperations.get
compute.projects.get
compute.regions.get
compute.networks.list
compute.firewalls.list
resourcemanager.projects.get
resourcemanager.projects.getIamPolicy
logging.sinks.create
logging.sinks.delete
logging.sinks.get
logging.sinks.list
logging.sinks.update
pubsub.subscriptions.create
pubsub.subscriptions.delete
pubsub.subscriptions.get
pubsub.subscriptions.list
pubsub.subscriptions.consume
pubsub.topics.attachSubscription
pubsub.topics.create
pubsub.topics.delete
pubsub.topics.detachSubscription
pubsub.topics.get
pubsub.topics.getIamPolicy
pubsub.topics.list
pubsub.topics.setIamPolicy
pubsub.topics.update
cloudkms.keyRings.list
cloudkms.cryptoKeys.list
cloudkms.cryptoKeys.setIamPolicy
cloudkms.cryptoKeys.getIamPolicy
serviceusage.services.list
```

IMPORTANT

To allow Veeam Backup for Google Cloud to back up a VM instance connected to a Shared VPC network, the service account associated with the project to which the instance belongs must also have either the `compute.networkUser` role for the whole Shared VPC host project, or the `compute.networkViewer` role for the whole host project plus `compute.networkUser` for specific subnets in the host project.

To learn how to provide access to Shared VPC networks, see [Google Cloud documentation](#).

Cloud SQL Backup Permissions

```
cloudsql.backupRuns.create
cloudsql.backupRuns.delete
cloudsql.backupRuns.get
cloudsql.backupRuns.list
cloudsql.databases.list
cloudsql.instances.export
cloudsql.instances.get
cloudsql.instances.list
cloudsql.instances.listServerCas
cloudsql.instances.update
cloudsql.users.list
compute.regions.list
compute.zones.list
logging.sinks.create
logging.sinks.delete
logging.sinks.get
logging.sinks.list
pubsub.subscriptions.consume
pubsub.subscriptions.create
pubsub.subscriptions.delete
pubsub.subscriptions.get
pubsub.subscriptions.list
pubsub.topics.attachSubscription
pubsub.topics.create
pubsub.topics.delete
pubsub.topics.detachSubscription
pubsub.topics.get
pubsub.topics.getIamPolicy
pubsub.topics.list
pubsub.topics.setIamPolicy
serviceusage.services.list
cloudkms.keyRings.list
cloudkms.cryptoKeys.list
compute.projects.get
resourcemanager.projects.get
```

IMPORTANT

To allow Veeam Backup for Google Cloud to use Cloud IAM credentials while backing up a MySQL instance, the service account associated with the project to which the instance belongs must also have the `cloudsql.instances.login` permission assigned.

Cloud Spanner Backup Permissions

```
spanner.databases.list
spanner.databases.get
spanner.databases.getDdl
spanner.databases.beginReadOnlyTransaction
spanner.databases.partitionQuery
spanner.databases.select
spanner.instanceConfigs.get
spanner.instanceConfigs.list
spanner.instances.get
spanner.instances.list
spanner.sessions.create
spanner.sessions.delete
compute.regions.list
compute.zones.list
logging.sinks.create
logging.sinks.delete
logging.sinks.get
logging.sinks.list
pubsub.subscriptions.consume
pubsub.subscriptions.create
pubsub.subscriptions.delete
pubsub.subscriptions.get
pubsub.subscriptions.list
pubsub.topics.attachSubscription
pubsub.topics.create
pubsub.topics.delete
pubsub.topics.detachSubscription,
pubsub.topics.get
pubsub.topics.getIamPolicy
pubsub.topics.list
pubsub.topics.setIamPolicy
serviceusage.services.list
cloudkms.keyRings.list
cloudkms.cryptoKeys.list
compute.projects.get
monitoring.timeSeries.list
resourcemanager.projects.get
```

Restore Permissions

To allow Veeam Backup for Google Cloud to perform restore operations, the service account associated with the Google Cloud project that will be used to manage the restored instances must have the following permissions.

VM Restore Permissions

```
compute.addresses.list
compute.disks.create
compute.disks.get
compute.disks.setLabels
compute.disks.use
compute.disks.delete
compute.disks.useReadOnly
compute.firewalls.list
compute.globalOperations.list
compute.globalOperations.get
compute.instances.create
compute.instances.delete
compute.instances.get
compute.instances.setLabels
compute.instances.setMachineResources
compute.instances.setMetadata
compute.instances.setMinCpuPlatform
compute.instances.setScheduling
compute.instances.setServiceAccount
compute.instances.setTags
compute.instances.start
compute.instances.stop
compute.instances.updateDisplayDevice
compute.instances.updateNetworkInterface
compute.instances.setDeletionProtection
compute.machineTypes.list
compute.networks.list
compute.projects.get
compute.regionOperations.get
compute.regions.get
compute.regions.list
compute.snapshots.create
compute.snapshots.delete
compute.snapshots.get
compute.snapshots.getIamPolicy
compute.snapshots.list
compute.snapshots.setLabels
compute.snapshots.useReadOnly
compute.subnetworks.list
compute.subnetworks.use
compute.subnetworks.useExternalIp
compute.zoneOperations.get
compute.zones.get
compute.zones.list
iam.serviceAccounts.actAs
iam.serviceAccounts.list
resourcemanager.projects.get
cloudkms.cryptoKeys.list
cloudkms.keyRings.list
compute.addresses.use
compute.addresses.useInternal
compute.disks.list
```

```
compute.instances.list
compute.routes.list
cloudkms.cryptoKeys.setIamPolicy
cloudkms.cryptoKeys.getIamPolicy
serviceusage.services.list
pubsub.subscriptions.setIamPolicy
pubsub.subscriptions.getIamPolicy
pubsub.topics.setIamPolicy
pubsub.topics.getIamPolicy
storage.objects.create
storage.objects.delete
storage.objects.list
storage.objects.get
storage.objects.update
storage.buckets.create
```

IMPORTANT

- To allow Veeam Backup for Google Cloud to perform VM restore to the original location while source VM instances still exist there, you must also add the permission `compute.instances.setName`.
- To allow Veeam Backup for Google Cloud to connect a restored VM instance to a Shared VPC network, the service account associated with the project to which the instance belongs must also have either the `compute.networkUser` role for the whole Shared VPC host project, or the `compute.networkViewer` role for the whole host project plus `compute.networkUser` for specific subnets in the host project.

To allow Veeam Backup for Google Cloud to check the subnet configuration of the Shared VPC network to which the restored VM instance is connected, you must also add the following permissions to the service account associated with the project to which the instance belongs: `compute.firewalls.list`, `compute.networks.get`, `compute.routes.list` and `compute.subnetworks.get` for the whole Shared VPC host project.

To learn how to provide access to Shared VPC networks, see [Google Cloud documentation](#).

Cloud SQL Restore Permissions

```
cloudkms.cryptoKeys.getIamPolicy
cloudkms.cryptoKeys.list
cloudkms.cryptoKeys.setIamPolicy
cloudkms.keyRings.list
cloudsql.backupRuns.get
cloudsql.instances.create
cloudsql.instances.get
cloudsql.instances.import
cloudsql.instances.restoreBackup
cloudsql.instances.update
compute.firewalls.list
compute.networks.list
compute.projects.get
compute.regions.list
compute.routes.list
compute.subnetworks.list
compute.zones.list
resourcemanager.projects.get
cloudsql.backupRuns.list
cloudsql.databases.create
cloudsql.databases.list
cloudsql.instances.list
cloudsql.instances.listServerCas
cloudsql.users.create
cloudsql.users.list
cloudsql.users.update
pubsub.subscriptions.consume
pubsub.subscriptions.create
pubsub.subscriptions.delete
pubsub.subscriptions.get
pubsub.subscriptions.list
pubsub.topics.attachSubscription
pubsub.topics.create
pubsub.topics.delete
pubsub.topics.detachSubscription
pubsub.topics.get
pubsub.topics.list
serviceusage.services.list
cloudsql.backupRuns.create
cloudsql.backupRuns.delete
cloudsql.databases.get
```

IMPORTANT

To allow Veeam Backup for Google Cloud to use Cloud IAM credentials while restoring a MySQL instance, the service account associated with the project to which the instance belongs must also have the `cloudsql.instances.login` permission assigned.

Cloud Spanner Restore Permissions

```
spanner.backupOperations.get
spanner.backups.get
spanner.backups.restoreDatabase
spanner.backups.delete
spanner.databaseOperations.get
spanner.databases.create
spanner.databases.list
spanner.databases.update
spanner.instanceConfigOperations.get
spanner.instanceConfigs.create
spanner.instanceConfigs.delete
spanner.instanceConfigs.get
spanner.instanceConfigs.list
spanner.instanceOperations.get
spanner.instances.create
spanner.instances.delete
spanner.instances.get
spanner.instances.list
cloudkms.cryptoKeys.getIamPolicy
cloudkms.cryptoKeys.list
cloudkms.cryptoKeys.setIamPolicy
cloudkms.keyRings.list
compute.projects.get
monitoring.timeSeries.list
resourcemanager.projects.get
spanner.databases.get
spanner.databases.updateDdl
spanner.databases.beginOrRollbackReadWriteTransaction
spanner.databases.beginReadOnlyTransaction
spanner.databases.write
spanner.databases.select
spanner.sessions.create
spanner.sessions.delete
pubsub.subscriptions.consume
pubsub.subscriptions.create
pubsub.subscriptions.delete
pubsub.subscriptions.get
pubsub.subscriptions.list
pubsub.topics.attachSubscription
pubsub.topics.create
pubsub.topics.delete
pubsub.topics.detachSubscription
pubsub.topics.get
pubsub.topics.list
serviceusage.services.list
resourcemanager.projects.get
resourcemanager.projects.getIamPolicy
resourcemanager.projects.setIamPolicy
```

Permissions Changelog

This section describes the latest changes in service account permissions required for Veeam Backup for Google Cloud to perform operations.

When you update Veeam Backup for Google Cloud version 4.0 to version 5.0, consider that additional permissions must be granted to the service accounts used to perform the following operations.

Repository Creation

```
storage.multipartUploads.create  
storage.multipartUploads.abort
```

File-Level Restore to Original Location

```
pubsub.subscriptions.setIamPolicy  
pubsub.subscriptions.getIamPolicy  
pubsub.topics.setIamPolicy  
pubsub.topics.getIamPolicy  
storage.objects.create  
storage.objects.delete  
storage.objects.list  
storage.objects.get  
storage.objects.update  
storage.buckets.create
```

Cloud Spanner Backup and Restore

```
compute.regions.list
compute.disks.list
compute.instances.get
compute.instances.list
compute.snapshots.get
compute.snapshots.list
compute.zones.get
compute.zones.list
compute.globalOperations.get
compute.zoneOperations.get
compute.regionOperations.get
resourcemanager.projects.get
compute.projects.get
compute.firewalls.list
compute.snapshots.getIamPolicy
compute.networks.list
compute.subnetworks.list
resourcemanager.projects.getIamPolicy
iam.serviceAccounts.actAs
compute.disks.create
compute.disks.createSnapshot
compute.disks.delete
compute.disks.setLabels
compute.instances.attachDisk
compute.instances.create
compute.instances.delete
compute.instances.detachDisk
compute.instances.setMetadata
compute.instances.setServiceAccount
compute.instances.setLabels
compute.instances.setTags
compute.routes.list
compute.regions.get
compute.snapshots.create
compute.snapshots.setLabels
compute.snapshots.setIamPolicy
compute.snapshots.delete
pubsub.subscriptions.consume
pubsub.subscriptions.create
pubsub.subscriptions.delete
pubsub.subscriptions.list
pubsub.subscriptions.get
logging.sinks.get
logging.sinks.delete
logging.sinks.list
pubsub.topics.attachSubscription
pubsub.topics.detachSubscription
pubsub.topics.create
pubsub.topics.delete
pubsub.topics.list
pubsub.topics.get
```

```
pubsub.topics.publish  
compute.machineTypes.get  
compute.machineTypes.list  
compute.subnetworks.get  
compute.subnetworks.use  
compute.subnetworks.useExternalIp  
compute.disks.use  
serviceusage.services.list
```

Google Cloud APIs

The backup appliance and worker instances must have outbound internet access to the following Google Cloud APIs:

- [Compute Engine API](#)
- [Service Usage API](#)
- [IAM Service Account Credentials API](#)
- [Identity and Access Management \(IAM\) API](#)
- [Cloud Resource Manager API](#)
- [Cloud Billing API](#)
- [Pub/Sub API](#)
- [Cloud Key Management Service API](#)
- [Cloud SQL Admin API](#)
- [Cloud Logging API](#)
- [Cloud Spanner API](#)
- [Cloud Deployment Manager v2 API](#)

Considerations and Limitations

When you plan to deploy and configure Veeam Backup for Google Cloud, keep in mind the following limitations and considerations.

Licensing

If the license file is not installed, Veeam Backup for Google Cloud will operate in the *Free* edition allowing you to protect up to 10 instances free of charge.

Software

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

Security Certificates

Veeam Backup for Google Cloud supports certificates only in the PFX and P12 formats.

Backup Repositories

When managing backup repositories, consider the following:

- The *Coldline* storage class is not supported. For more information on storage classes offered by Cloud Storage, see [Google Cloud documentation](#).
- You cannot change Google Cloud storage buckets, subdirectories and storage classes for backup repositories already added to Veeam Backup for Google Cloud.
- Customer-supplied encryption keys (CSEKs) are not supported for repository encryption.
- 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 Repositories](#).
- A backup repository must not be managed by multiple backup appliances simultaneously. Retention sessions running on different appliances may corrupt backups stored in the repository, which may result in unpredictable data loss.

Worker Instances

When managing worker instances, consider the following:

- For Veeam Backup for Google Cloud to be able to deploy the number of worker instances required for a backup or restore process, you must have enough resource quotas allocated between your projects. To learn how to check your quotas, see [Google Cloud documentation](#).

- To allow Veeam Backup for Google Cloud to connect a created worker instance to a Shared VPC network, the service account associated with the Google Cloud project to which the instance belongs must have either the `compute.networkUser` role for the whole Shared VPC host project, or the `compute.networkViewer` role for the whole host project plus `compute.networkUser` for specific subnets in the host project.

To learn how to provide access to Shared VPC networks, see [Google Cloud documentation](#).

Backup

When protecting Google Cloud resources, consider the following:

- Veeam Backup for Google Cloud allows you to protect MySQL and PostgreSQL instances. SQL Server instances are not supported. For more information on types of Cloud SQL instances, see [Google Cloud documentation](#).
- To allow Veeam Backup for Google Cloud to back up a VM instance connected to a Shared VPC network, the service account associated with the project to which the instance belongs must have either the `compute.networkUser` role for the whole Shared VPC host project, or the `compute.networkViewer` role for the whole host project plus `compute.networkUser` for specific subnets in the host project.

To learn how to provide access to Shared VPC networks, see [Google Cloud documentation](#).

- Backup of the default PostgreSQL databases (*template0*, *template1* and *postgres*) is not supported.
- When backing up Cloud Spanner instances, Veeam Backup for Google Cloud does not process their internal settings — except for the *version_retention_period* setting. For more information on internal settings of Cloud Spanner instances, see [Google Cloud documentation](#).

Restore

When restoring Google Cloud resources, consider the following:

- When restoring a VM instance, Veeam Backup for Google Cloud recovers data from all zonal and regional persistent disks (standard, balanced, extreme and SSD) attached to the instance. However, due to [technical reasons](#), when it comes to local SSDs (SCSI and NVMe), Veeam Backup for Google Cloud is able to recover only the configuration of these disks, which means that any data stored on the disks is lost during the restore process.
- To allow Veeam Backup for Google Cloud to connect a restored VM instance to a Shared VPC network, the service account associated with the project to which the instance belongs must have either the `compute.networkUser` role for the whole Shared VPC host project, or the `compute.networkViewer` role for the whole host project plus `compute.networkUser` for specific subnets in the host project.

To learn how to provide access to Shared VPC networks, see [Google Cloud documentation](#).

- Restore of local SSDs (SCSI and NVMe) is not supported due to [technical reasons](#).
- File-level recovery is supported for FAT, FAT32, NTFS, ext2, ext3, ext4, XFS and Btrfs file systems only. However, attributes of files and folders stored in FAT and FAT32 file systems cannot be restored to the original location.
- Restore of NTFS links (hard links, junction points, symbolic links) to the original location is not supported.
- Due to [Google Cloud limitations](#), Veeam Backup for Google Cloud does not support restore to the original location if the source Cloud SQL instance is still present in Google Cloud, if it has been recently deleted (less than a week ago), or if its name is reserved.

- Restore of PostgreSQL instances to Cloud SQL instances of the *db-f1-micro* and *db-g1-small* machine types is not supported. If you want to restore a PostgreSQL instance to one of the specified machine types, you must first manually create a Cloud SQL instance of the necessary type in the Google Cloud console as described in [Google Cloud documentation](#), and then restore the backed-up databases to the created instance as described in section [Performing Database Restore](#).
- Restore of the default PostgreSQL databases (*template0*, *template1* and *postgres*) is not supported.
- Restore of encrypted files to their original locations is not supported.
- When restoring encrypted folders to the original locations, folder encryption attributes will not be restored.
- When restoring root folders to their original locations while the folders no longer exist in these locations, Veeam Backup for Google Cloud restores all the folder attributes in the *Overwrite* mode.
- Due to Google Cloud limitations, Veeam Backup for Google Cloud does not support data encryption of Cloud SQL instances with multi-regional keys. For more information, see [Cloud SQL for MySQL documentation](#) and [Cloud SQL for PostgreSQL documentation](#).
- Due to [Google Cloud limitations](#), Veeam Backup for Google Cloud does not support database restore to the original location if the source database is still located on the server.
- When restoring Cloud SQL instances, Veeam Backup for Google Cloud turns off the point in time recovery setting, and it is turned on automatically only as soon as the restore process completes, which means that all the historical data is lost.

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 Google Cloud 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 that you understand the implications of the recommended practices, or request assistance. If in doubt, reach out to Veeam professionals on [Veeam R&D Forums](#).

Backup Appliance

You can choose the machine type of the VM instance running Veeam Backup for Google Cloud during the deployment, or later as the environment grows.

General Recommendations

The following recommendations and examples apply to the latest Veeam Backup for Google Cloud builds (4.0.0.1072 or later).

By default, the backup appliance can process 25 workloads per policy and up to 15 sessions simultaneously, including running policies, restore, rescan and retention activities.

Appliance Size*	Recommended Maximum Number of Protected Workloads
e2-highmem-2 (2 vCPU, 8 GB RAM)	100
e2-standard-4 (4 vCPU, 16 GB RAM)	200
e2-standard-8 (8 vCPU, 32 GB RAM)	2,500
e2-standard-16 (16 vCPU, 64 GB RAM)	5,000

*It is recommended to add 8 vCPU and 32 GB RAM per each additional 2,500 workloads.

Veeam Backup & Replication Integration

When you connect a backup appliance to the backup infrastructure, its policy and retention data is imported into the Veeam Backup & Replication database.

You can connect multiple backup appliances to a single Veeam Backup & Replication server. However, when working in a Google service account with cross-region data transfer, it is recommended to use one Veeam

Backup & Replication server per region, to help you avoid latency issues and meet potential data residency regulations.

Time Consumption

When you connect an existing backup appliance to the backup infrastructure, the integration process includes the following steps:

- Retrieving data from the appliance.
- Saving the retrieved data to the Veeam Backup & Replication database.

Protected Workloads	Snapshots	Backups	Policy Sessions	Workload Processing Sessions	Time Consumption
100	10,000	10,000	1	100	0:02:26
200	20,000	20,000	2	200	0:03:44
2,500	250,000	250,000	25	2,500	1:34:52
5,000	500,000	500,000	50	5,000	3:31:20

NOTE

The process of synchronizing data between the backup appliance and Veeam Backup & Replication database runs every 2 minutes after you add the appliance to the backup infrastructure. Creating new backup policies and updating policy settings may also trigger the synchronization process.

Object Storage

Veeam Backup for Google Cloud 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 Google Cloud saves different objects to Google Cloud storage buckets.

Object Type	Block Size
Backup data (Standard)	1 MB (compressed to ~512 KB)
Backup data (Archive)	512 MB
Metadata	4 KB (per 1 GB of VM source data)

Storage Bucket Placement

To achieve best performance, create backup repositories in regional storage buckets and place them in the same region as source instances. A situation where a storage bucket is located far from a source instance may cause slow network throughput between regions.

Cost Estimation

Veeam Backup for Google Cloud comes with a built-in cost calculator that allows you to estimate your Google Cloud expenses. It uses publicly available Google Cloud 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. As a result, you will have several small policies instead of a big one.

The default limit for simultaneously processed policies is 25 instances per one policy, with 15 sessions running in parallel. These values can be changed in the configuration file `/opt/veeam/gcpbackup/ServiceSettings.json`.

```
{
  "JobProcess": {
    "MaxParallelJobProcesses": 15
  },
  "Snapshot": {
    "MaxConcurrentSnapshots": 25
  },
  "Backup": {
    "MaxConcurrentInstanceBackups": 25
  }
}
```

Where:

- `MaxParallelJobProcesses` – the maximum number of simultaneously processed sessions (including policies, restore instance sessions, FLR sessions, and so on).
- `MaxConcurrentSnapshots`, `MaxConcurrentInstanceBackups` – the maximum number of simultaneously processed VM, Cloud SQL or Cloud Spanner instances per one policy.

Keep in mind that changing these values may induce additional monitoring of the backup appliance resource usage since it may require the machine type of the appliance to be changed to a larger one.

IMPORTANT

It is not recommended to manually change the default limit for policies processed simultaneously – to adjust the limit, open a [support case](#).

Worker Instances

If you want initial full backups to be processed quickly, it is recommended to use a larger worker profile, and then change it to a smaller profile for incremental backup. You can change worker profile settings on a regional basis, so make sure that the selected profile is appropriate to process the largest workload within the required time.

Each worker instance is deployed as an Ubuntu image and is removed once the task that it performs completes. Machine types of worker instances depend on the regional quota.

Profile	Machine Type	Case	Backup Speed
Primary	e2-highcpu-8	Processing resources while sufficient disk quota is available	Up to 420 MBps (NTFS disks up to 540 MBps)
Secondary	e2-highcpu-2	Processing resources while running out of disk quota	Up to 210 MBps
Archiving	e2-standard-4	Transferring data to archive repositories	Up to 420 MBps

For details on Google Cloud pricing, see [Google Cloud documentation](#).

VM Instance Backup

The default configuration (e2-highcpu-8) is universal and available in all Google Cloud regions. For this configuration, the backup speed is up to 420 MBps if the sum of the source disk sizes is less than 5 TB. For NTFS disks, the backup speed is up to 540 MBps. For better speed consistency and overall performance at the same price level, it is recommended to use n2d-highcpu-8 as the primary worker profile. By changing the profile to e2-highcpu-16, it is possible to achieve the backup speed up to 800 MBps; however, this will require adjusting the performance disk size of worker instances by changing the value in the configuration file `/opt/veeam/gcpbackup/ServiceSettings.json`.

```
"Backup": {  
  "TotalHddDisksSize": 10000  
}
```

If the sum of source disk sizes is more than 10 TB, changing the primary worker profile to e2-highcpu-16 (n2d-highcpu-16) allows you to achieve the backup speed up to 800 MBps. However, this will increase the total monthly infrastructure costs.

VM Instance Archive

The default configuration (e2-standard-4) allows you to archive data up to 420 MBps. By changing the archiving worker profile to e2-standard-8, it is possible to achieve the archiving speed up to 500 MBps; however, this will increase the total monthly infrastructure costs.

VM Instance Restore

The default configuration (e2-highcpu-4) allows you to restore data up to 170 MBps. For better speed consistency and overall performance at the same price level, you can use n2d-highcpu-4 as the worker profile.

By default, a worker instance deployed for the entire VM instance restore operation will use an additional performance disk up to 1500 GB; the disk size can be changed in the configuration file `/opt/veeam/gcpbackup/ServiceSettings.json` to increase the restore speed. To restore VM instances with a size of 24 TB or more, you can change the worker profile to e2-highcpu-8.

```
"HardwareSettings": {
  "Restore": "e2-highcpu-8"
}
```

To achieve a higher restore speed, you can change the worker profile to e2-highcpu-8 with the additional 2048 GB performance disk. However, this will increase the total monthly infrastructure costs.

```
"HardwareSettings": {
  "Restore": "e2-highcpu-8"
}
"Restore": {
  "TotalHddDisksSize": 2048
}
```

File-Level Recovery

To avoid prolonged execution time of file-level recovery operations, it is recommended to change the worker profile to e2-highmem-4 if the processed VM instances have a lot of disks. The profile can be changed in the configuration file `/opt/veeam/gcpbackup/ServiceSettings.json`.

```
"HardwareSettings": {
  "Flr": "e2-highmem-4"
}
```

IMPORTANT

It is not recommended to manually change the default worker profile that is used to launch worker instances performing restore or file-level recovery operations – to customize the profile, open a [support case](#).

Retention

By default, all retention processes run on the backup appliance. However, it is possible to execute these processes on a worker instance (except for the deletion of an entire backup chain). To do that, it is recommended to change the backup retention threshold to O in the configuration file `/opt/veeam/gcpbackup/ServiceSettings.json`. You can also adjust the worker profile to fit the size of the largest processed source instance.

```
{
  "BackupRetention": {
    "CreateWorkerRestorePointsThresholdGb": 150
  }
  "HardwareSettings": {
    "Retention": "e2-highcpu-8"
  }
}
```

Worker Profile Recommendations for Retention

Largest Source Instance	Worker Profile
Less than 8 TB	e2-highcpu-8
Between 8 and 16 TB	e2-highcpu-16
Larger than 16 TB	e2-highcpu-32

IMPORTANT

It is not recommended to manually change the default worker profile that is used to launch worker instances performing retention operations – to customize the profile, open a [support case](#).

Deployment

To deploy Veeam Backup for Google Cloud, 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 Google Cloud Plug-in for Veeam Backup & Replication [system requirements](#).

2. [Install Google Cloud Plug-in for Veeam Backup & Replication on the backup server](#).

This step applies only to Veeam Backup & Replication versions prior to 12.0. Version 12.0 (and later) comes pre-packed with Google Cloud Plug-in for Veeam Backup & Replication.

3. [Deploy a backup appliance](#).

Deploying Plug-In

If your installation package of Veeam Backup & Replication does not provide features that allow you to protect Google Cloud resources, you must install Google Cloud Plug-in for Veeam Backup & Replication on the backup server to be able to add your backup appliances to the backup infrastructure.

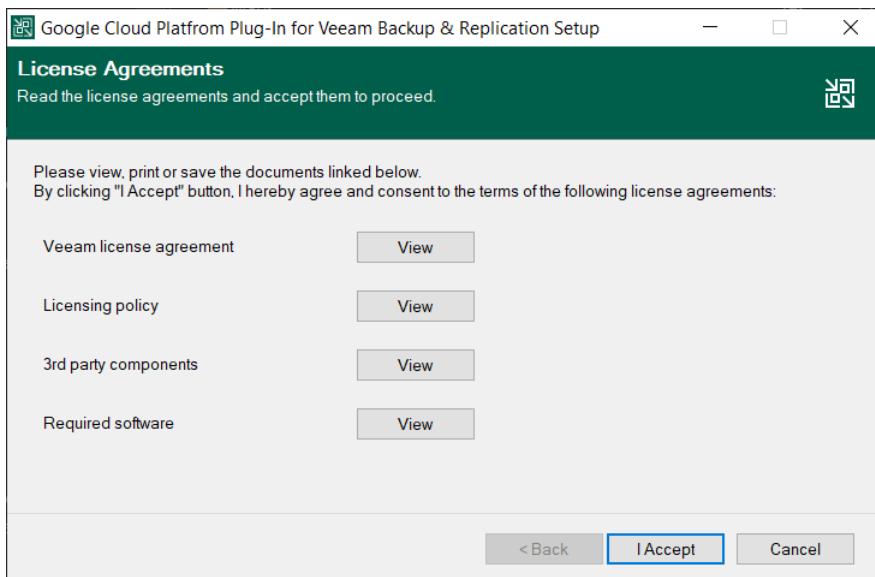
NOTE

Before you install Google Cloud Plug-in for Veeam Backup & Replication, stop all running backup policies, disable all jobs, and close the Veeam Backup & Replication console.

To install Google Cloud 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 Google Cloud Plug-in for Veeam Backup & Replication.
3. Open the downloaded `GCPPlugin_12.5.0.1257.zip` file and launch the `GCPPlugin_12.5.0.1257.exe` installation file.
4. Complete the **Google Cloud Plug-in for Veeam Backup & Replication** wizard:
 - a. At the **License Agreements** step, read and accept both the Veeam license agreement, licensing policy, the 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 license agreements, click **View**.
 - b. At the **Installation Path** step, 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 Plug-In in Unattended Mode

You can install and uninstall Google Cloud Plug-in for Veeam Backup & Replication in the unattended mode using the command line interface. The unattended installation 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 the unattended installation mode to automate the Google Cloud Plug-in for Veeam Backup & Replication installation process in large-scale environments.

To install Google Cloud Plug-in for Veeam Backup & Replication in the unattended mode, use either of the following options:

- If Google Cloud 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 Google Cloud Plug-in for Veeam Backup & Replication is delivered as a separate .EXE file, follow the instructions provided in this section.

Before You Begin

Before you start unattended installation, do the following:

1. Download the Google Cloud Plug-in for Veeam Backup & Replication .EXE file as described in section [Installing Plug-In](#) (steps 1–4).
2. Check compatibility of the Google Cloud 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: "GCPPlugin_12.5.0.1257.exe /silent /accepteula /acceptthirdpartylicenses /uninstall"
/repair	No	Replaces missing files and firewall rules. Example: "GCPPlugin_12.5.0.1257.exe /silent /accepteula /acceptthirdpartylicenses /repair"

Upgrading Plug-In

To upgrade Google Cloud Plug-in for Veeam Backup & Replication, do the following:

1. Install a new version of Google Cloud 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](#).

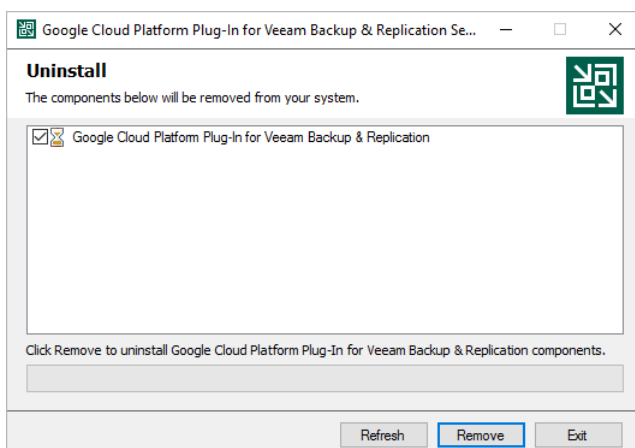
Uninstalling Plug-In

Before you uninstall Google Cloud Plug-in for Veeam Backup & Replication, it is recommended to [remove all connected backup appliances](#) from the backup infrastructure. If you keep the appliances in the backup infrastructure, the following will happen:

- You will be able to see information on snapshots of VM instances, Cloud SQL instances and Cloud Spanner Instances in the Veeam Backup & Replication console. However, you will not be able to perform any operations with these snapshots.
- You will be able to see information on image-level backups of Cloud SQL instances in the Veeam Backup & Replication console. However, you will not be able to perform any operations with these backups.
- You will be able to see information on image-level backups of VM instances and perform data recovery operations using these backups. However, restore of entire VM instances to Google Cloud will start working as described in the Veeam Backup & Replication User Guide, section [How Restore to Google Compute Engine Works](#).
- You will be able to see information on image-level backups of Cloud Spanner instances in the Veeam Backup & Replication console. However, you will not be able to perform any operations with these backups.
- 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 Google Cloud 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. Open the **Start** menu, navigate to **Control Panel > Programs > Programs and Features**.
3. In the program list, click **Google Cloud Plug-in for Veeam Backup & Replication** and click **Uninstall**.
4. In the opened window, click **Remove**.



NOTE

After you uninstall Google Cloud Plug-in for Veeam Backup & Replication, you will be no longer able to add backup appliances and cloud repositories to the backup infrastructure.

Deploying Backup Appliance

After you install Google Cloud Plug-in for Veeam Backup & Replication, you must add backup appliances to the backup infrastructure. To do that, use either of the following options:

- [Deploying Backup Appliance from Console](#)
- [Deploying Backup Appliance from Google Cloud Marketplace](#)

Deploying Backup Appliance from Console

A backup appliance comes as an image of a Linux-based VM that you can deploy from Veeam Backup & Replication console. Veeam Backup for Google Cloud is installed on a VM instance that is created in Google Cloud during the product installation.

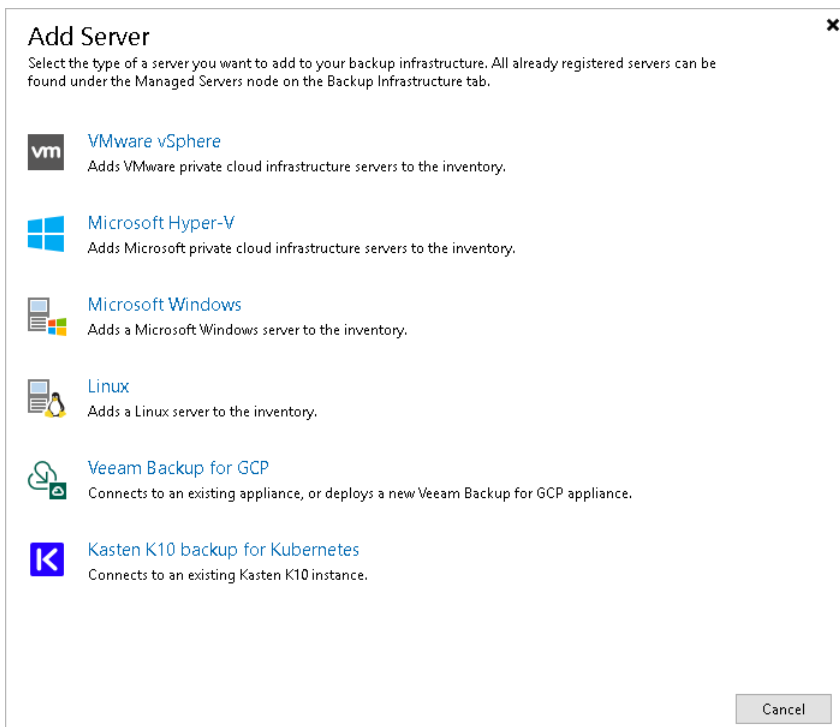
To deploy a new backup appliance from the Veeam Backup & Replication console, do the following:

1. [Launch the New Veeam Backup for Google Cloud Appliance wizard.](#)
2. [Choose a deployment mode.](#)
3. [Specify a Veeam Backup for Google Cloud account in which the appliance will be deployed.](#)
4. [Specify a name and description for the appliance.](#)
5. [Specify network settings for the appliance.](#)
6. [Specify IP address settings.](#)
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 GCP Appliance Wizard

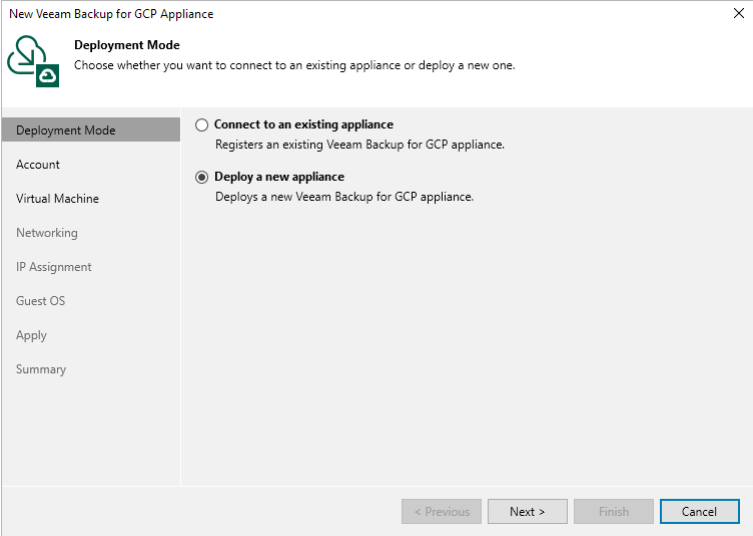
To launch the **New Veeam Backup for GCP Appliance** wizard, do the following:

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 GCP**.



Step 2. Choose Deployment Mode

At the **Deployment Mode** step of the wizard, select the **Deploy a new appliance** option.



Step 3. Specify Service Account Settings

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

1. From the **GCP service account** drop-down list, select a service account whose permissions will be used to deploy the new backup appliance. Note that the specified service account will further be used by Veeam Backup & Replication to connect to this appliance.

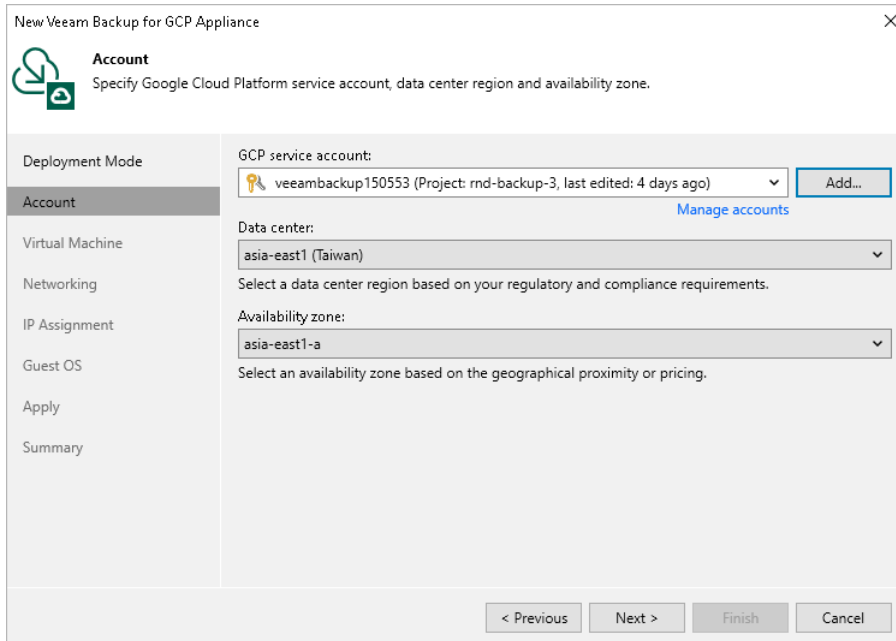
For a service account to be displayed in the **GCP service account** drop-down list, it must be created in Google Cloud and added to the Cloud Credentials Manager. You can either create a new service account or add an existing service account to the Cloud Credentials Manager without closing the wizard. To do that, follow the instructions provided in the Veeam Backup & Replication User Guide, section [Google Cloud Platform Service Accounts](#).

NOTE

When you create a service account using the Veeam Backup & Replication console, the service account is automatically assigned the **Owner IAM role** with a wide scope of permissions and capabilities. If you want the service account to be assigned a limited list of permissions, create a service account **manually in Google Cloud** beforehand and then add it to the Cloud Credentials Manager. For more information on required permissions that must be assigned to the service account, see [Plug-In Permissions](#).

2. From the **Data center** drop-down list, select a Google Cloud region in which the backup appliance will reside.
3. From the **Availability zone** drop-down list, select a location within a Google Cloud region where you want to deploy the backup appliance.

For more information on regions and zones in Google Cloud, see [Google Cloud documentation](#).



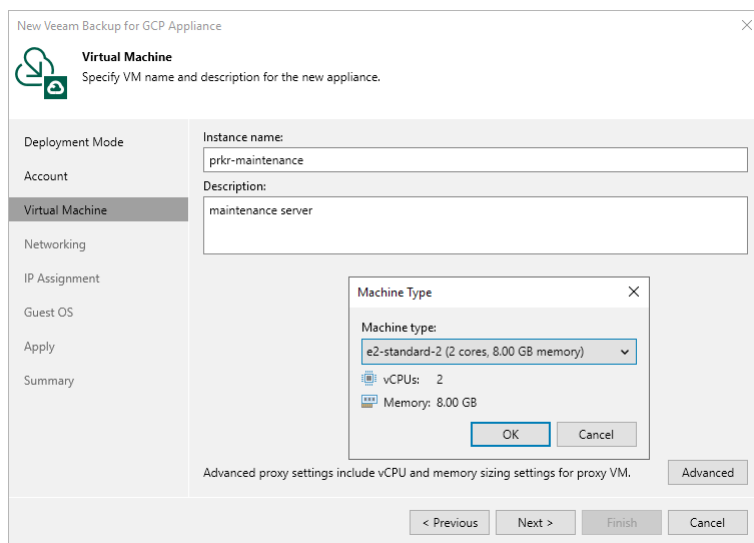
Step 4. Specify VM Instance Name and Description

At the **Virtual Machine** step of the wizard, specify a name and description for the VM instance where Veeam Backup for Google Cloud will be deployed. Note that the name must meet the [naming convention for Compute Engine resources](#).

TIP

By default, Veeam Backup & Replication uses the recommended *e2-standard-2* machine type for the backup appliance. If you want to define a specific machine type for the VM instance, click **Advanced** and select the necessary type in the **Machine Type** window.

For the list of all existing machine types, see [Sizing and Scalability Guidelines](#).



Step 5. Specify Network Settings

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

1. Choose a virtual private cloud (VPC) network to which the backup appliance will be connected.

You can create a new VPC network or specify an existing one:

- To create a new VPC network, select the **(create new)** option from the **VPC** drop-down list. Veeam Backup & Replication will automatically create a network with a set of predefined firewall rules.
- To specify an existing VPC network, select it from the **VPC** drop-down list. For a VPC network to be displayed in the list of available networks, it must be created in the Google Cloud for the region specified at [step 3](#) of the wizard, as described in [Google Cloud documentation](#).

2. Choose a subnet to which the backup appliance will be connected.

You can create a new subnet or specify an existing one:

- 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 network.
- 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 network as described in [Google Cloud documentation](#).

3. Choose a network tag that will be assigned to the backup appliance.

You can create a new tag or specify an existing one:

- To create a new tag, select the **(create new)** option from the **Network tag** drop-down list. Veeam Backup & Replication will automatically create a tag with the appliance name.

If you have chosen to connect the backup appliance to a shared VPC network, Veeam Backup & Replication will not be able to create a new network tag with required firewall rules automatically while deploying the appliance. That is why you must either specify an existing network tag, or configure firewall rules associated with the selected VPC manually.

- To specify an existing tag, select it from the **Network tag** drop-down list. For a tag to be displayed in the list of available tags, it must be created in Google Cloud as described in [Google Cloud documentation](#).

IMPORTANT

If you specify an existing network tag, consider that the following firewall rules must apply to the tag:

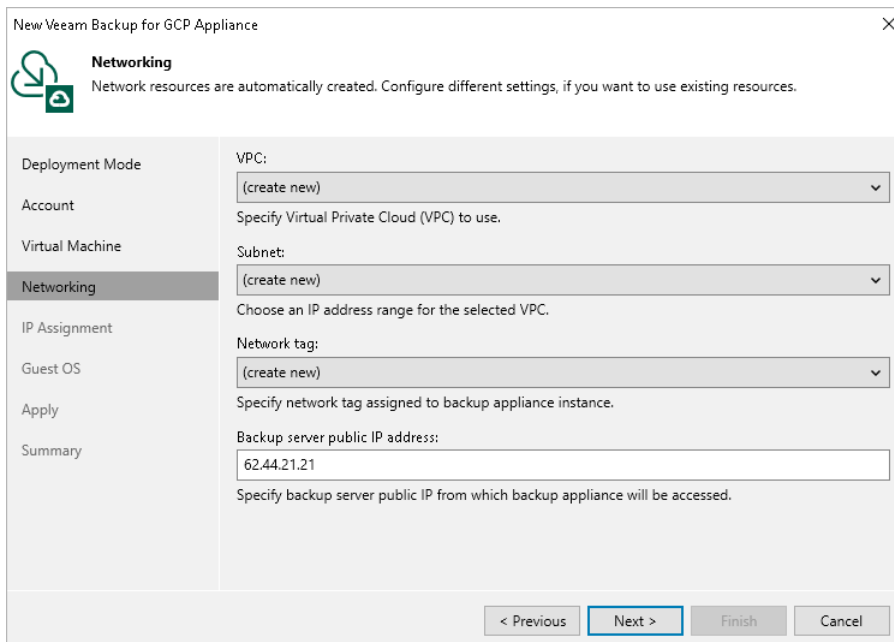
- A rule that allows outbound internet access from the backup appliance to Google Cloud APIs listed in section [Planning and Preparation](#).
- A rule that allows inbound internet access to the backup appliance from both the backup server and a local machine that you plan to use to work with Veeam Backup for Google Cloud.
- A rule that allows ingress traffic from the Google IAP to the backup appliance through the SSH protocol (IP range 35.235.240.0/20) to perform automatic updates of the TLS certificates installed on the appliance. For more information on the Google IAP, see [Google Cloud documentation](#).

To learn how to create firewall rules, see [Google Cloud documentation](#).

4. [Applies only if you have chosen to create a new network tag] In the **Backup server public IP address** field, specify an IP address or a scope of IP addresses that will be allowed to access the backup appliance. Veeam Backup & Replication will create a firewall rule for the specified IP addresses. Note that the IP address of the backup server must fall into the specified IP address range.

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 scopes of IP addresses, use a comma-separated list.



The screenshot shows a wizard window titled "New Veeam Backup for GCP Appliance" with a close button (X) in the top right corner. The "Networking" step is active, indicated by a green checkmark icon and a bolded title. Below the title, a note states: "Network resources are automatically created. Configure different settings, if you want to use existing resources." The wizard is divided into two main sections: a left-hand navigation pane and a right-hand configuration area. The navigation pane lists steps: Deployment Mode, Account, Virtual Machine, Networking (highlighted), IP Assignment, Guest OS, Apply, and Summary. The configuration area contains the following fields and instructions:

- VPC:** A dropdown menu with "(create new)" selected. Below it, the instruction reads: "Specify Virtual Private Cloud (VPC) to use."
- Subnet:** A dropdown menu with "(create new)" selected. Below it, the instruction reads: "Choose an IP address range for the selected VPC."
- Network tag:** A dropdown menu with "(create new)" selected. Below it, the instruction reads: "Specify network tag assigned to backup appliance instance."
- Backup server public IP address:** A text input field containing "62.44.21.21". Below it, the instruction reads: "Specify backup server public IP from which backup appliance will be accessed."

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

Step 6. Specify IP Address Settings

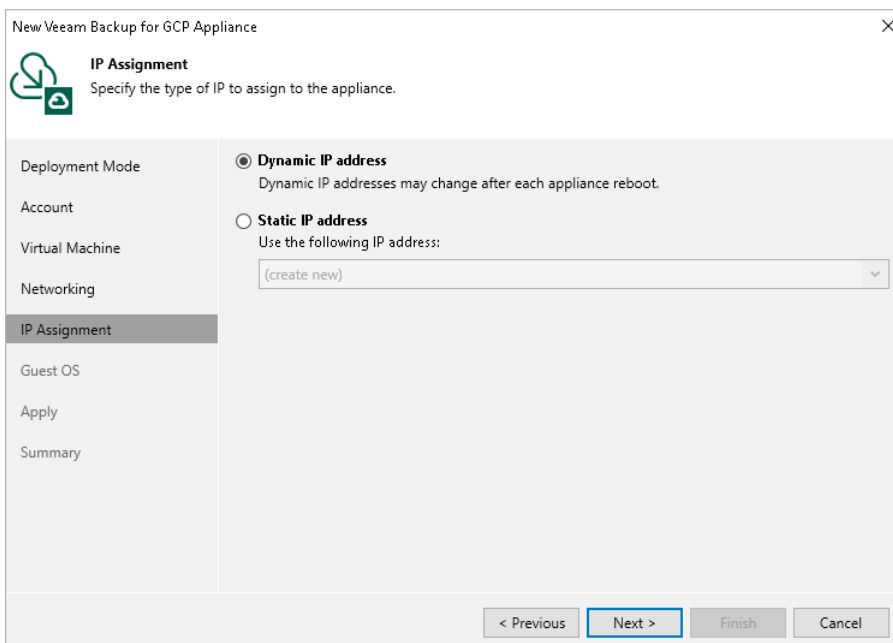
At the **IP Assignment** step of the wizard, choose whether you want to assign a dynamic or a static IP address to the backup appliance.

To assign a static 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 static IP addresses, it must be reserved in Google Cloud as described in [Google Cloud documentation](#).

NOTE

You can use only IPv4 regional IP address as static external IP addresses for backup appliances.

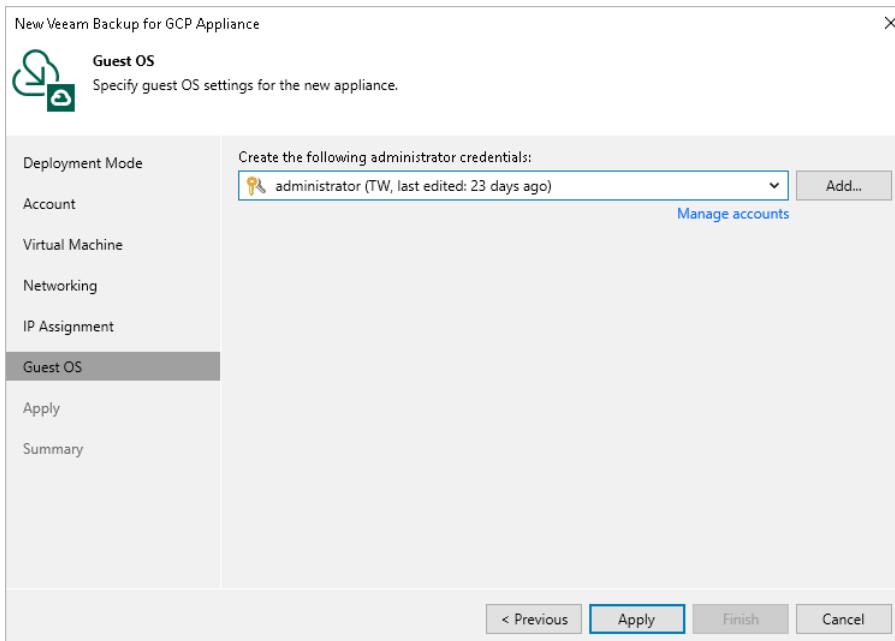


The screenshot shows the 'IP Assignment' step of the 'New Veeam Backup for GCP Appliance' wizard. The window title is 'New Veeam Backup for GCP Appliance' with a close button (X) in the top right corner. The main heading is 'IP Assignment' with a sub-heading 'Specify the type of IP to assign to the appliance.' Below this, there is a sidebar on the left with navigation options: Deployment Mode, Account, Virtual Machine, Networking, IP Assignment (highlighted), Guest OS, Apply, and Summary. The main content area has two radio button options: 'Dynamic IP address' (selected) and 'Static IP address'. Under 'Dynamic IP address', it says 'Dynamic IP addresses may change after each appliance reboot.' Under 'Static IP address', it says 'Use the following IP address:' followed by a dropdown menu currently showing '(create new)'. At the bottom of the window, there are four buttons: '< Previous', 'Next >' (highlighted with a blue border), 'Finish', and 'Cancel'.

Step 7. Specify User Credentials

At the **Guest OS** step of the wizard, specify a user whose credentials Veeam Backup & Replication will use to create the Default Administrator account on the backup appliance.

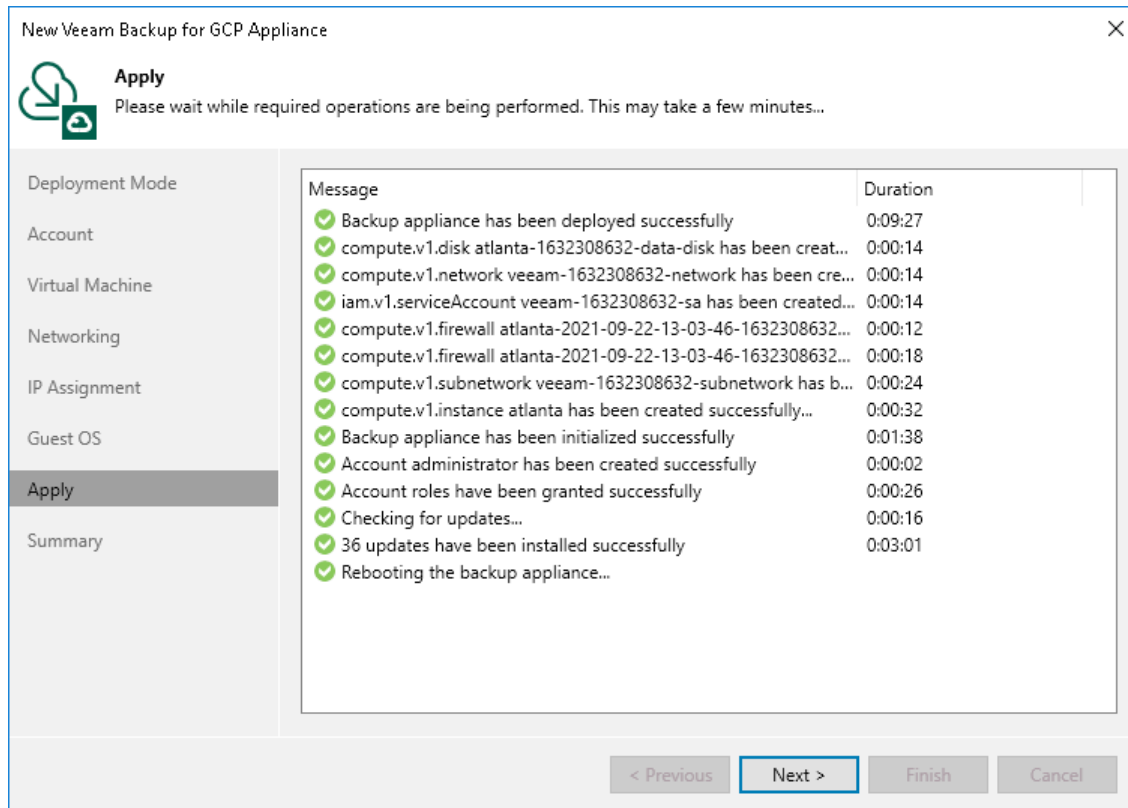
For a user to be displayed in the **Create the following administrator credentials** drop-down list, it must be added to the Credentials Manager. If you have not added a user to the Credential Manager beforehand, you can do it without closing the **New Veeam Backup for GCP Appliance** wizard. To add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



The screenshot shows the 'New Veeam Backup for GCP Appliance' wizard window. The title bar reads 'New Veeam Backup for GCP Appliance' with a close button (X) on the right. Below the title bar is a logo and the text 'Guest OS Specify guest OS settings for the new appliance.' A sidebar on the left contains a list of steps: Deployment Mode, Account, Virtual Machine, Networking, IP Assignment, Guest OS (highlighted), Apply, and Summary. The main area is titled 'Create the following administrator credentials:' and features a dropdown menu with the selected user 'administrator (TW, last edited: 23 days ago)'. To the right of the dropdown is an 'Add...' button and a blue link labeled 'Manage accounts'. At the bottom of the window, there are four buttons: '< Previous', 'Apply' (highlighted with a blue border), 'Finish', and 'Cancel'.

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**.



New Veeam Backup for GCP Appliance

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

Deployment Mode	Message	Duration
Account	Backup appliance has been deployed successfully	0:09:27
Virtual Machine	compute.v1.disk atlanta-1632308632-data-disk has been creat...	0:00:14
Networking	compute.v1.network veeam-1632308632-network has been cre...	0:00:14
IP Assignment	iam.v1.serviceAccount veeam-1632308632-sa has been created...	0:00:14
Guest OS	compute.v1.firewall atlanta-2021-09-22-13-03-46-1632308632...	0:00:12
Apply	compute.v1.firewall atlanta-2021-09-22-13-03-46-1632308632...	0:00:18
Summary	compute.v1.subnetwork veeam-1632308632-subnetwork has b...	0:00:24
	compute.v1.instance atlanta has been created successfully...	0:00:32
	Backup appliance has been initialized successfully	0:01:38
	Account administrator has been created successfully	0:00:02
	Account roles have been granted successfully	0:00:26
	Checking for updates...	0:00:16
	36 updates have been installed successfully	0:03:01
	Rebooting the backup appliance...	

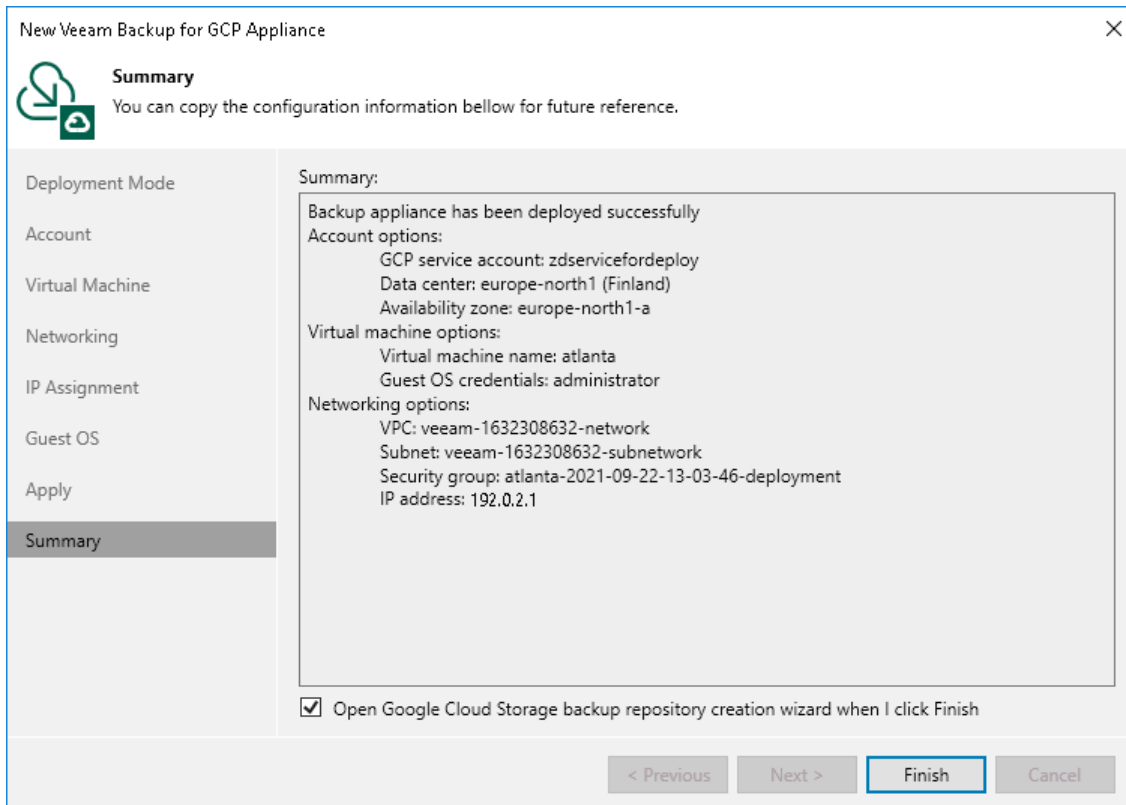
< Previous **Next >** Finish Cancel

Step 9. Finish Working with Wizard

At the **Summary** step of the wizard, review summary information and click **Finish**. After the backup appliance is deployed, you will be able to configure its settings in the Veeam Backup for Google Cloud Web UI.

TIP

If you want to configure repositories immediately after the backup appliance is deployed, select the **Open Google Cloud Storage backup repository creation wizard when I click Finish** check box and follow the instructions provided in section [Adding Repositories](#).



The screenshot shows the 'New Veeam Backup for GCP Appliance' wizard at the 'Summary' step. The window title is 'New Veeam Backup for GCP 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, Virtual Machine, Networking, IP Assignment, Guest OS, Apply, and Summary (which is currently selected and highlighted). The main area contains the following text:

Summary
You can copy the configuration information below for future reference.

Summary:

- Backup appliance has been deployed successfully
- Account options:
 - GCP service account: zdservicefordeploy
 - Data center: europe-north1 (Finland)
 - Availability zone: europe-north1-a
- Virtual machine options:
 - Virtual machine name: atlanta
 - Guest OS credentials: administrator
- Networking options:
 - VPC: veeam-1632308632-network
 - Subnet: veeam-1632308632-subnetwork
 - Security group: atlanta-2021-09-22-13-03-46-deployment
 - IP address: 192.0.2.1

At the bottom of the main area, there is a checkbox labeled 'Open Google Cloud Storage backup repository creation wizard when I click Finish' which is checked.

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

Deploying Backup Appliance from Google Cloud Marketplace

A backup appliance comes as an image of a Linux-based VM that you can deploy from Google Cloud Marketplace. Veeam Backup for Google Cloud is installed on a VM instance that is created in Google Cloud during the product installation.

To deploy a backup appliance, do the following:

1. Log in to [Google Cloud Marketplace](#) using credentials of a Google account that has the Editor role granted.
To learn how to manage user roles in the Google Cloud console, see [Google Cloud documentation](#).
2. Click **Explore the marketplace**.
3. In the search field, enter *Veeam Backup for Google Cloud* and press [Enter] on the keyboard.
4. In the list of search results, click *Veeam Backup for Google Cloud* to open the product overview page.
5. Click **Launch**.

The screenshot shows the product overview page for Veeam Backup for Google Cloud. At the top, there is a Veeam logo and the product name. Below that, a description states: "Google-native backup and recovery, built for simplicity, scale, savings and security". There are three buttons: "LAUNCH" (highlighted in blue), "VIEW DEPLOYMENTS", and "CONTACT SALES". A navigation bar includes links for "OVERVIEW", "PRICING", "DOCUMENTATION", "SUPPORT", and "RELATED PRODUCTS". The "Overview" section contains text about the product's integration with Google Cloud and its benefits. The "Additional details" section lists technical specifications: "Runs on: Google Compute Engine", "Type: Virtual machines, Single VM", "Architecture: X86_64", "Last product update: 12/6/23", "Category: Security, Storage", "Version: 5.0", "Operating System: Ubuntu 20.04", "Package contents: Veeam Backup for Google Cloud 5.0", and "Add to Service Catalog: Deployment .zip file".

6. On the **New Veeam Backup for Google Cloud deployment** page, configure the following installation settings:
 - a. Select a project to which the VM instance running Veeam Backup for Google Cloud will belong.

IMPORTANT

Make sure that Google Cloud APIs listed in the [Planning and Preparation](#) section are enabled for the selected project. Otherwise, Veeam Backup for Google Cloud deployment may fail or cause unexpected errors. To learn how to enable APIs for Google Cloud projects, see [Google Cloud documentation](#).

- b. In the **Deployment name** field, enter a name for the new Veeam Backup for Google Cloud deployment.

The deployment will include the VM instance running Veeam Backup for Google Cloud, the Google Cloud service account used by the VM instance to access Google Cloud APIs, firewall rules defined to allow traffic to and from the VM instance, and other configuration details specified during installation.

- c. From the **Zone** drop-down list, select an availability zone within a Google Cloud region in which the VM instance running Veeam Backup for Google Cloud will reside.

To learn how to configure availability and redundancy settings for Google Cloud resources, see [Google Cloud documentation](#).

- d. In the **Machine type** section, specify the number of vCPUs and the amount of memory on Compute Engine that will be allocated to the VM instance running Veeam Backup for Google Cloud.

The recommended hardware requirement for a VM instance running Veeam Backup for Google Cloud is an *e2-standard-2* instance with 2 vCPUs and 8 GB RAM.

- e. In the **Disks** section, specify the size of a boot disk that will be attached to the VM instance running Veeam Backup for Google Cloud, and the size of an additional data disk where the application database and logs will be stored.

New Veeam Backup for Google Cloud deployment

Deployment name *
veeam-backup-for-google-cloud-2

Zone
us-west3-a

Machine type

General purpose Compute optimized

Machine types for common workloads, optimized for cost and flexibility

Series
E2

CPU platform selection based on availability

Machine type
e2-standard-2 (2 vCPU, 1 core, 8 GB memory)

	vCPU	Memory
	2	8 GB

Disks

Boot disk size in GB
10

Data disk size in GB
20

Additional information

Veeam Backup for Google Cloud overview
Product provided by Veeam Software

Veeam Backup for Google Cloud Platform Usage Fee CHF 0.00/mo
Veeam Software does not charge a usage fee.

All products are priced in USD and charged in the currency (CHF) specified by your Billing Account. The price for this month is calculated with an exchange rate of 1 USD = 0.88 CHF

created or consumed by this product (or the fees charged for such consumption). Veeam Software may be able to provide a more accurate estimate of monthly GCP IaaS consumption.

Software

Operating System	Ubuntu(20.04)
Software	Veeam Backup for Google Cloud(5.0)

- f. In the **Networking** section, specify a VPC network and a subnet to which the VM instance running Veeam Backup for Google Cloud will be connected.

For a VPC network and a subnet to be displayed in the lists of available networks, they must be created in the Google Cloud console for the region specified at step 5b, as described in [Google Cloud documentation](#).

IMPORTANT

- The specified VPC network and subnet must have the outbound internet access to Google Cloud APIs listed in the [Planning and Preparation](#) section. Otherwise, Veeam Backup for Google Cloud will not work properly.
- The specified VPC network and subnet must allow the inbound internet access from a local machine that you plan to use to work with Veeam Backup for Google Cloud.

To learn how to enable internet access for VPC networks and subnets, see [Google Cloud documentation](#).

If there are no firewall rules that allow inbound HTTPS traffic in the specified network, you must select the **Allow HTTPS traffic from the internet** and specify the allowed IP address ranges explicitly.

If you plan connect to Veeam Backup for Google Cloud using the [Veeam Backup for Google Cloud REST API](#), you must select the **Allow public API traffic from the internet** check box and specify the allowed IP address ranges. You can also allow SSH connections to the backup appliance, which may be required for debugging and troubleshooting purposes.

TIP

The IPv4 address ranges must be specified in the CIDR notation (for example, 12.23.34.0/24). To let all IPv4 addresses access the VM instance running Veeam Backup for Google Cloud, you can enter 0.0.0.0/0. However, note that allowing access from all IPv4 addresses is unsafe and thus not recommended in production environments.

- g. Click **Deploy** to begin installation.

After installation completes, the **Suggested next steps** section will display a link to the Veeam Backup for Google Cloud Web UI. Click the link to proceed to the [initial configuration](#) required to start working with Veeam Backup for Google Cloud.

After You Install

To start working with Veeam Backup for Google Cloud, you must perform the initial configuration of the backup appliance:

1. In a web browser, navigate to the Veeam Backup for Google Cloud 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

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

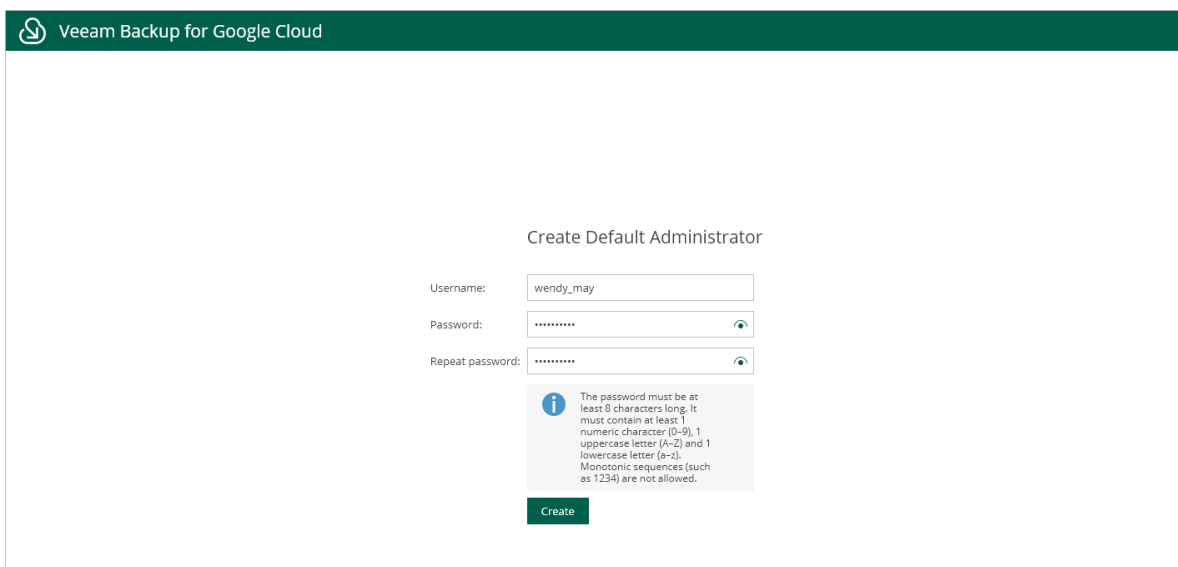
2. Read and accept the Veeam license agreement, Veeam licensing policy, license agreements of the 3rd party components and license agreements of the required 3rd party software. If you reject the terms and conditions, you will not be able to continue installation.
3. In the **Instance ID** field, specify the unique numeric identifier of the VM instance running Veeam Backup for Google Cloud to prove that you are the owner of this VM instance.

To obtain the ID assigned to the VM instance upon creation, you can either look it up on the **Instances** page in the Google Cloud console, or send a query to the metadata server API using the gcloud command-line tool. To learn how to retrieve instance metadata, see [Google Cloud documentation](#).

4. Create the Default Administrator account whose credentials you will use for your first login to Veeam Backup for Google Cloud.

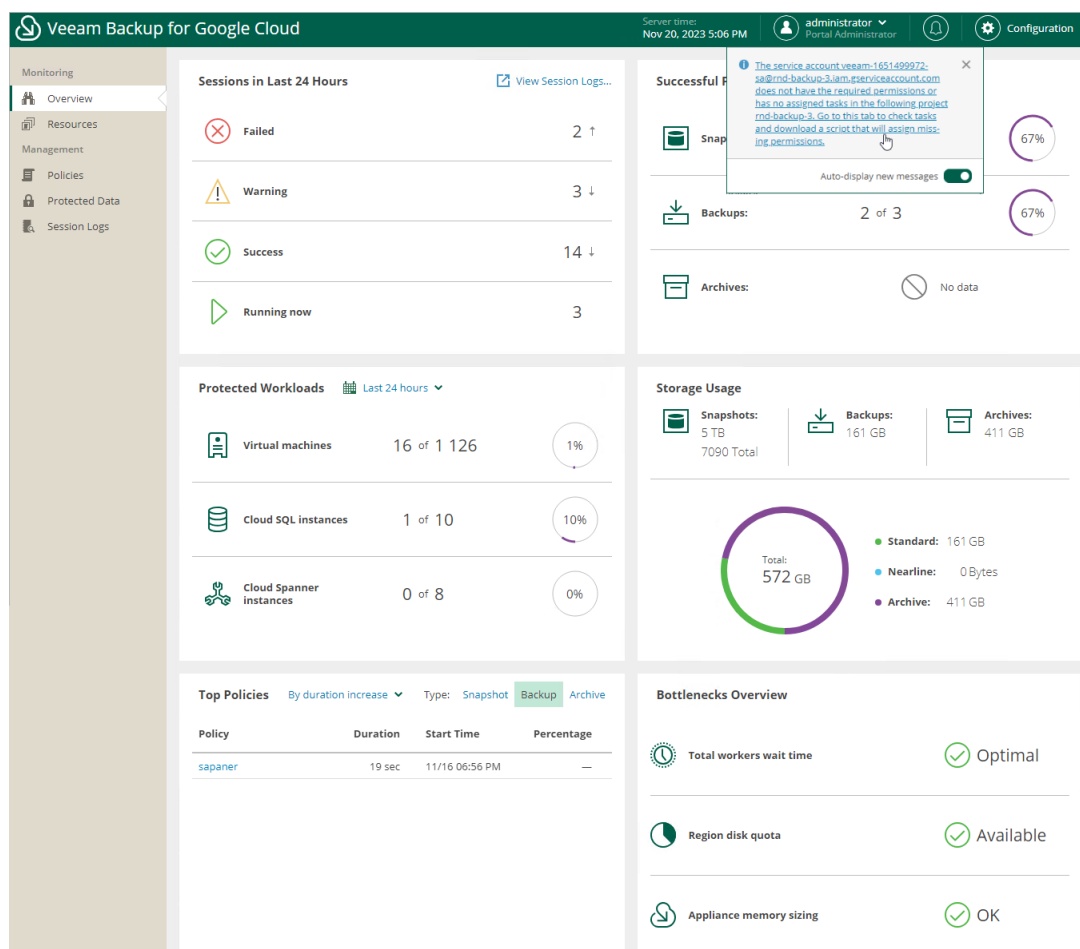
NOTE

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



5. Log in to Veeam Backup for Google Cloud with the credentials of the Default Administrator account as described in section [Accessing Veeam Backup for Google Cloud](#).

You will receive a warning in the notification area notifying that the service account created during product installation does not have the permissions required to perform data protection tasks for the project to which the backup appliance belongs. You can grant the missing permissions to the service account later when configuring Veeam Backup for Google Cloud as described in section [Managing Projects and Folders](#).



Uninstalling Veeam Backup for Google Cloud

Veeam Backup for Google Cloud creates a number of resources while operating in Google Cloud, and these resources are not removed from Google Cloud automatically when you uninstall the solution. That is why you must perform the following steps to uninstall Veeam Backup for Google Cloud:

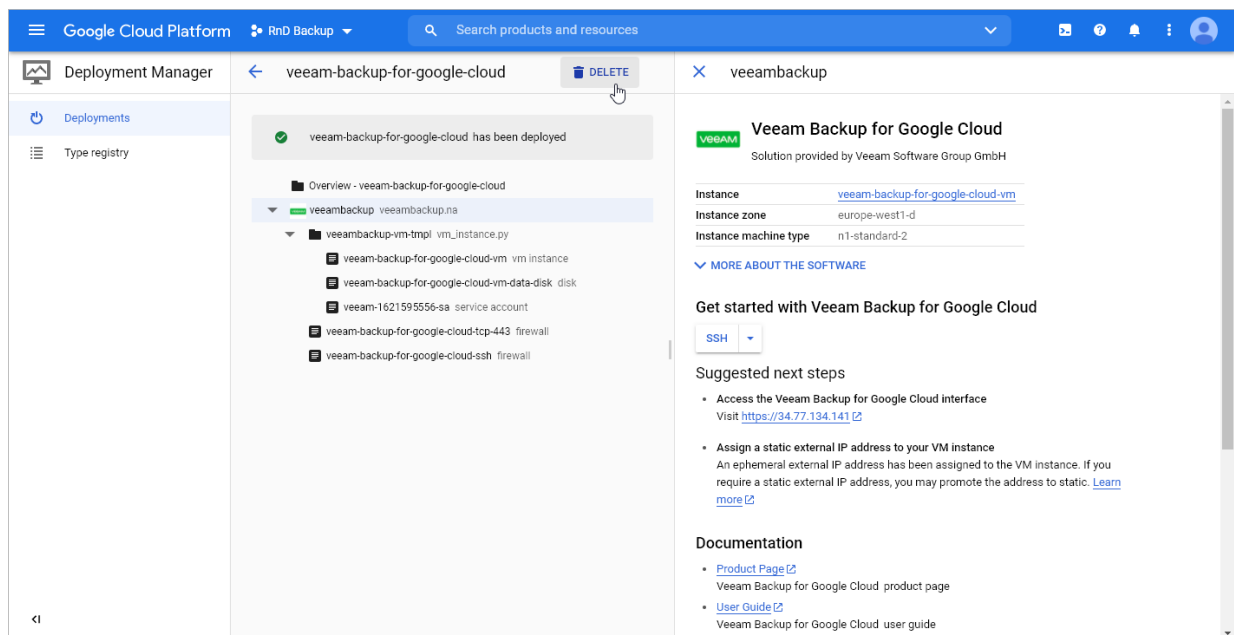
1. Locate and save the unique numeric identifier of the VM instance running Veeam Backup for Google Cloud — you will need it later.

To obtain the ID, you can either look it up on the **Instances** page in the Google Cloud console, or send a query to the metadata server API using the `gcloud` command-line tool. To learn how to retrieve instance metadata, see [Google Cloud documentation](#).

2. Save the names of Google Cloud projects that have ever been added to Veeam Backup for Google Cloud — you will need it later.

To obtain the names, you can look them up on the **Infrastructure > Projects and Folders** tab in the Veeam Backup for Google Cloud UI.

3. Log in to [Google Cloud Marketplace](#) using credentials of the Google account that you used to install Veeam Backup for Google Cloud.
4. Navigate to **Your products**.
5. Click *Veeam Backup for Google Cloud* to open the product overview page.
6. Click **Delete**.



7. Wait until Veeam Backup for Google Cloud is removed from your organization domain.

8. Navigate to **IAM & Admins > IAM**.

In the list of permissions, locate the `deleted:serviceAccount:veeam` member, and then unassign all existing roles from this member.

9. Navigate to **IAM & Admins > Roles**.

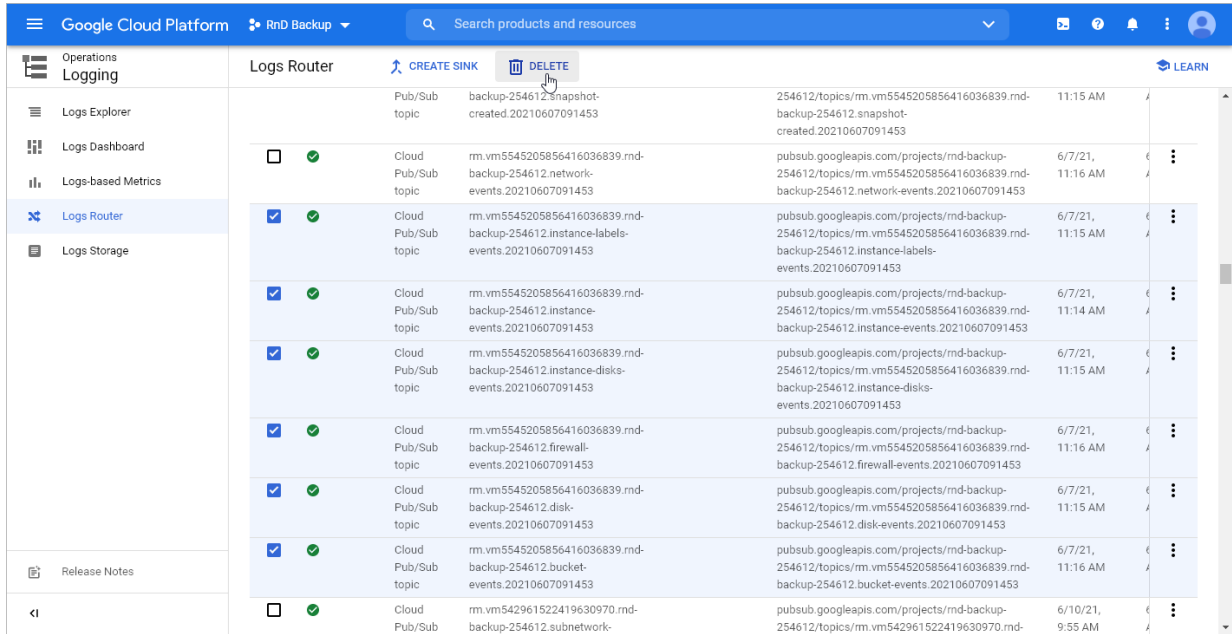
In the list of roles, locate the role either with the name that starts with *Veeam.VB.** (for Veeam Backup for Google Cloud version 4.0 or later) or with the name that contains the ID of the VM instance that was running Veeam Backup for Google Cloud (for the previous versions), and then delete this role.

NOTE

It may take up to one week for the role to be deleted.

10. Navigate to **Logging > Logs Router**.

In the list of logs router sinks, locate all sinks with the *Cloud Pub/Sub topic* type created by Veeam Backup for Google Cloud (the names of these sinks will contain the ID of the VM instance that was running Veeam Backup for Google Cloud), and then delete these sinks.



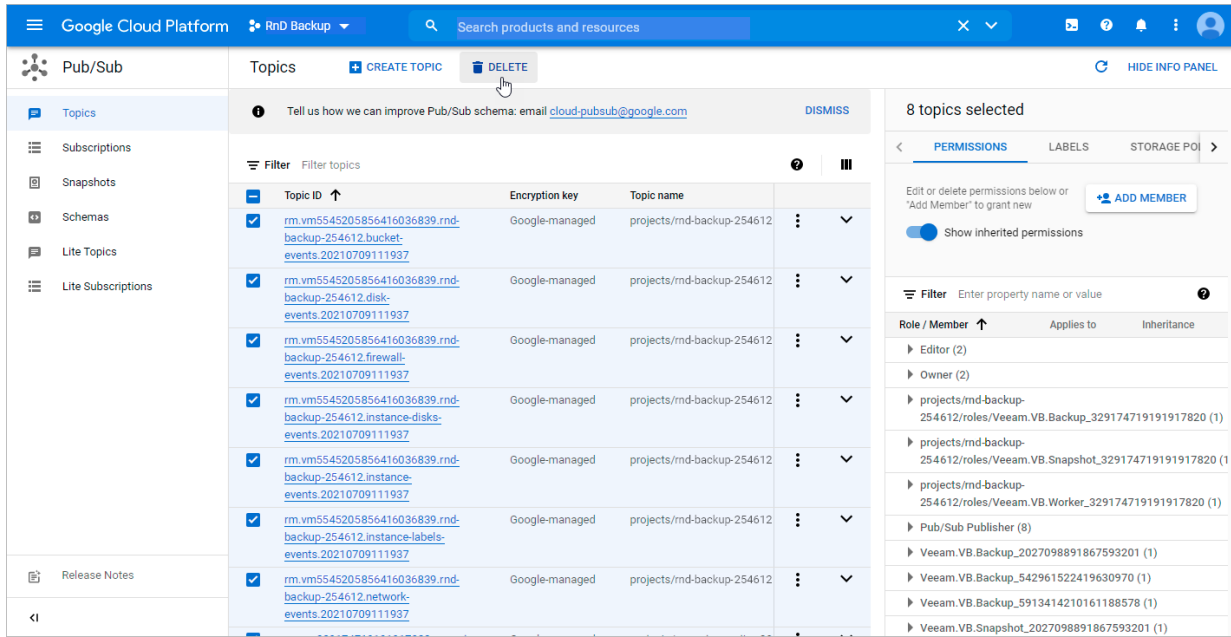
Operations Logging	Logs Router	CREATE SINK	DELETE	LEARN
Logs Explorer	Pub/Sub topic	backup-254612.snapshot-created.20210607091453	254612/topics/rm.vm5545205856416036839.md-backup-254612.snapshot-created.20210607091453	11:15 AM
Logs Dashboard	Cloud Pub/Sub topic	mm.vm5545205856416036839.md-backup-254612.network-events.20210607091453	pubsub.googleapis.com/projects/md-backup-254612/topics/rm.vm5545205856416036839.md-backup-254612.network-events.20210607091453	6/7/21, 11:16 AM
Logs-based Metrics	Cloud Pub/Sub topic	mm.vm5545205856416036839.md-backup-254612.instance-labels-events.20210607091453	pubsub.googleapis.com/projects/md-backup-254612/topics/rm.vm5545205856416036839.md-backup-254612.instance-labels-events.20210607091453	6/7/21, 11:15 AM
Logs Router	Cloud Pub/Sub topic	mm.vm5545205856416036839.md-backup-254612.instance-disks-events.20210607091453	pubsub.googleapis.com/projects/md-backup-254612/topics/rm.vm5545205856416036839.md-backup-254612.instance-disks-events.20210607091453	6/7/21, 11:14 AM
Logs Storage	Cloud Pub/Sub topic	mm.vm5545205856416036839.md-backup-254612.instance-disks-events.20210607091453	pubsub.googleapis.com/projects/md-backup-254612/topics/rm.vm5545205856416036839.md-backup-254612.instance-disks-events.20210607091453	6/7/21, 11:15 AM
Release Notes	Cloud Pub/Sub topic	mm.vm5545205856416036839.md-backup-254612.firewall-events.20210607091453	pubsub.googleapis.com/projects/md-backup-254612/topics/rm.vm5545205856416036839.md-backup-254612.firewall-events.20210607091453	6/7/21, 11:16 AM
	Cloud Pub/Sub topic	mm.vm5545205856416036839.md-backup-254612.disk-events.20210607091453	pubsub.googleapis.com/projects/md-backup-254612/topics/rm.vm5545205856416036839.md-backup-254612.disk-events.20210607091453	6/7/21, 11:15 AM
	Cloud Pub/Sub topic	mm.vm5545205856416036839.md-backup-254612.bucket-events.20210607091453	pubsub.googleapis.com/projects/md-backup-254612/topics/rm.vm5545205856416036839.md-backup-254612.bucket-events.20210607091453	6/7/21, 11:16 AM
	Cloud Pub/Sub topic	mm.vm542961522419630970.md-backup-254612.subnetwork-	pubsub.googleapis.com/projects/md-backup-254612/topics/rm.vm542961522419630970.md-	6/10/21, 9:55 AM

11. Navigate to **Pub/Sub > Subscriptions**.

In the list of subscriptions, locate all subscriptions created by Veeam Backup for Google Cloud (the names of these subscriptions will contain the ID of the VM instance that was running Veeam Backup for Google Cloud), and then delete these subscriptions.

12. Navigate to **Pub/Sub > Topics**.

In the list of topics, locate all topics created by Veeam Backup for Google Cloud (the names of these topics will contain the ID of the VM instance that was running Veeam Backup for Google Cloud), and then delete these topics.



13. Repeat steps 8–12 for each project that has ever been added to Veeam Backup for Google Cloud.

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](#)
- [Backup Appliance Licensing](#)

Limitations

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 Google Cloud 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](#).

Scenarios

An instance is considered to be protected if it has a restore point (backup or snapshot) created by a [backup policy](#) during the past 31 days. The number of license units that a protected instance consumes depends on the instance type and product edition. For more information, see [Veeam Licensing Policy](#).

NOTE

If an instance has only snapshots created manually, it does not consume any license units. To learn how to create cloud-native snapshots of VM, Cloud SQL and Cloud Spanner instances manually, see [Performing VM Backup](#), [Performing SQL Backup](#) and [Performing Spanner Backup](#).

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 appliance from the backup infrastructure, Veeam Backup & Replication stops counting backed-up instances. Veeam Backup for Google Cloud continues using the license that had been used before you added the appliance to the backup infrastructure.

- If you [deploy a new backup appliance](#) from the Veeam Backup & Replication console, instances start consuming license units from the license installed on the backup server after you create and run backup policies.

When you remove the appliance from the backup infrastructure, Veeam Backup & Replication stops counting backed-up instances and Veeam Backup for Google Cloud switches to the *Free* edition that allows you to protect up to 10 instances free of charge. To back up more than 10 instances, you must install a *BYOL* license on the appliance. To learn how to install a new *BYOL* license, see [Installing and Removing License](#).

Licensing When Connection to Veeam Backup & Replication is Lost

Veeam Backup for Google Cloud 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.

Backup Appliance Licensing

Veeam Backup for Google Cloud is licensed by the number of protected instances. An instance is defined as a single Google Cloud resource — a VM, Cloud SQL or Cloud Spanner instance. 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 one license unit from the license scope. However, if an instance has only snapshots created manually, it does not consume any license units.

NOTE

If an instance has not been backed up within the past 31 days, Veeam Backup for Google Cloud 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 Google Cloud is available in 2 editions:

- **Free**

By default, Veeam Backup for Google Cloud operates in the *Free* edition that allows you to protect up to 10 instances free of charge.

TIP

If you earlier deployed the *Free* edition of the product for evaluation and testing purposes on one instance, and now want to switch to the *BYOL* edition running on another instance without reconfiguring Veeam Backup for Google Cloud, follow the instructions provided in [this Veeam KB article](#).

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

The *BYOL* (Bring Your Own License) edition allows you to protect the number of instances equivalent to the number of units specified in your license. Veeam Backup for Google Cloud *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](#).

When the license expires, Veeam Backup for Google Cloud 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 30 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 Google Cloud stops processing all instances and disables all scheduled backup policies. You must update your license before the end of the grace period.

To learn how to install and update the license, see [Installing and Removing License](#).

NOTE

Veeam Backup & Replication licensing is applied to backup appliances managed by standalone Veeam Backup & Replication servers. For more information, see [Scenarios](#).

Installing and Removing Backup Appliance License

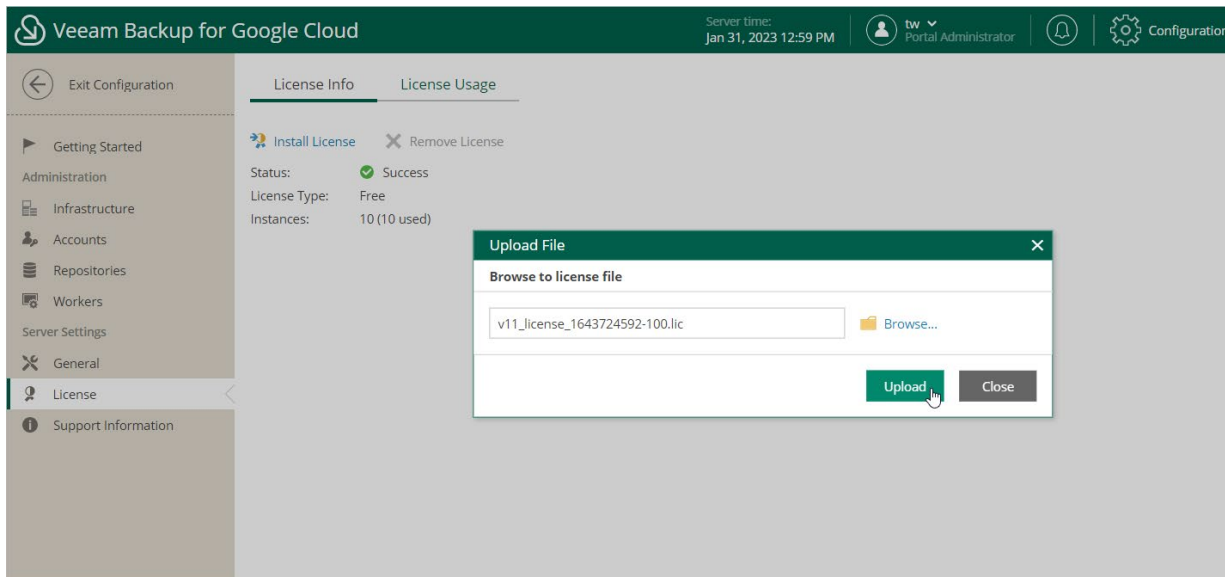
NOTE

This section applies only to the *BYOL* edition of Veeam Backup for Google Cloud.

Installing License

To install or update a license installed on the backup appliance, do the following:

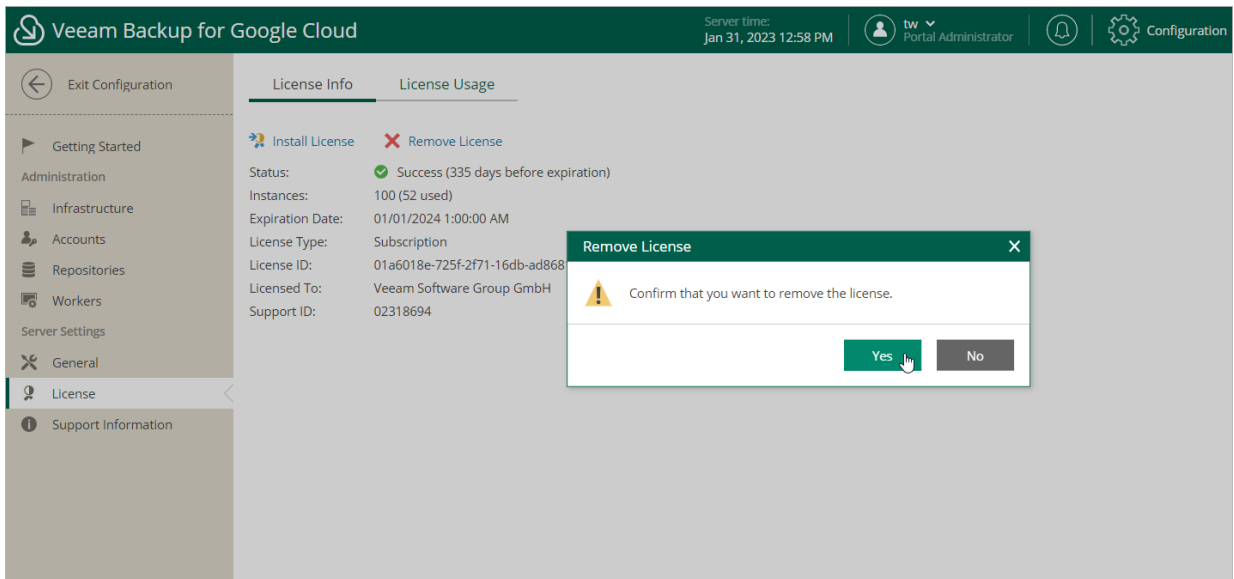
1. Switch to the **Configuration** page.
2. Navigate to **License > License Info**.
3. Click **Install license**.
4. In the **Upload File** window, click **Browse** to browse to a license file, and then click **Upload**.



Removing License

To remove a license installed on the backup appliance if you no longer need it, do the following:

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 Google Cloud 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 Using Veeam Backup & Replication Console

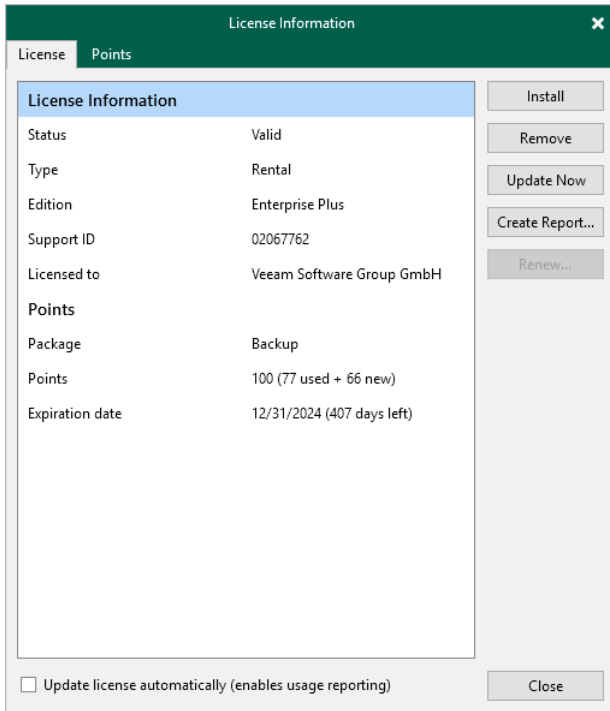
To view Google Cloud Plug-in for Veeam Backup & Replication license details in the Veeam Backup & Replication console, do the following:

1. Open the main menu.
2. Select **License**.

The **License** tab of the **License Information** window provides general information on the currently installed Google Cloud 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 instances.
 - **Virtual Machines** – protected VM instances.
 - **Cloud VMs** – protected VM instances.
 - **Cloud Databases** – protected Cloud SQL and Cloud Spanner instances.
- **Count** – the number of protected instances.
- **Multiplier** – the number of license units that one protected instance consumes.

- **Instances** – the total number of the consumed license units.

Type	Count	Multiplier	Instances
Cloud VMs	7	1	7
Cloud Databases	9	1	9

Viewing License Details Using Veeam Backup for Google Cloud Web UI

To view details on the license that is currently installed on the backup appliance in the Veeam Backup for Google Cloud Web UI, do the following:

1. Switch to the **Configuration** page.
2. Navigate to **License > License Info**.

The **License Info** tab provides general information on the Veeam Backup for Google Cloud 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 instances.

Each instance that has a restore point created in the past 31 days is considered to be protected and consumes one 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, Subscription*).

NOTE

Subscription is the name of the BYOL license in Veeam Backup for Google Cloud.

- **License ID** – the unique identification number of the provided license file (required for contacting the Veeam Customer Support Team).
- **Licensed To** – the name of an organization to which the license was issued.

- **Support ID** – the unique identification number of the support contract (required for contacting the Veeam Customer Support Team).

The screenshot displays the Veeam Backup for Google Cloud configuration interface. The top header shows the product name, server time (Jan 31, 2023 1:00 PM), and the user (tw Portal Administrator). The left sidebar contains navigation options: Exit Configuration, Getting Started, Administration, Infrastructure, Accounts, Repositories, Workers, Server Settings, General, License (selected), and Support Information. The main content area is titled 'License Info' and 'License Usage'. It features two buttons: 'Install License' and 'Remove License'. Below these buttons, the license status is shown as 'Success (335 days before expiration)'. Other details include 100 instances (52 used), an expiration date of 01/01/2024 1:00:00 AM, a subscription license type, license ID 01a6018e-725f-2f71-16db-ad8681407806, licensed to Veeam Software Group GmbH, and a support ID of 02318694.

Property	Value
Status:	Success (335 days before expiration)
Instances:	100 (52 used)
Expiration Date:	01/01/2024 1:00:00 AM
License Type:	Subscription
License ID:	01a6018e-725f-2f71-16db-ad8681407806
Licensed To:	Veeam Software Group GmbH
Support ID:	02318694

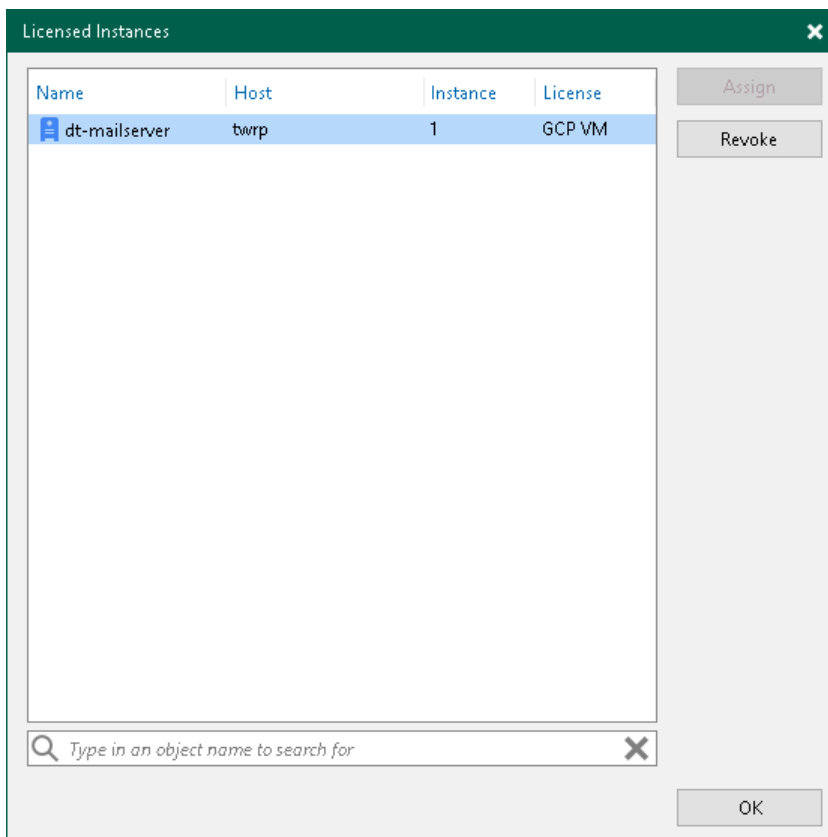
Revoking License Units

By default, Veeam Backup for Google Cloud 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 instance and click **Revoke**. Veeam Backup & Replication will revoke a license unit from the selected instance.

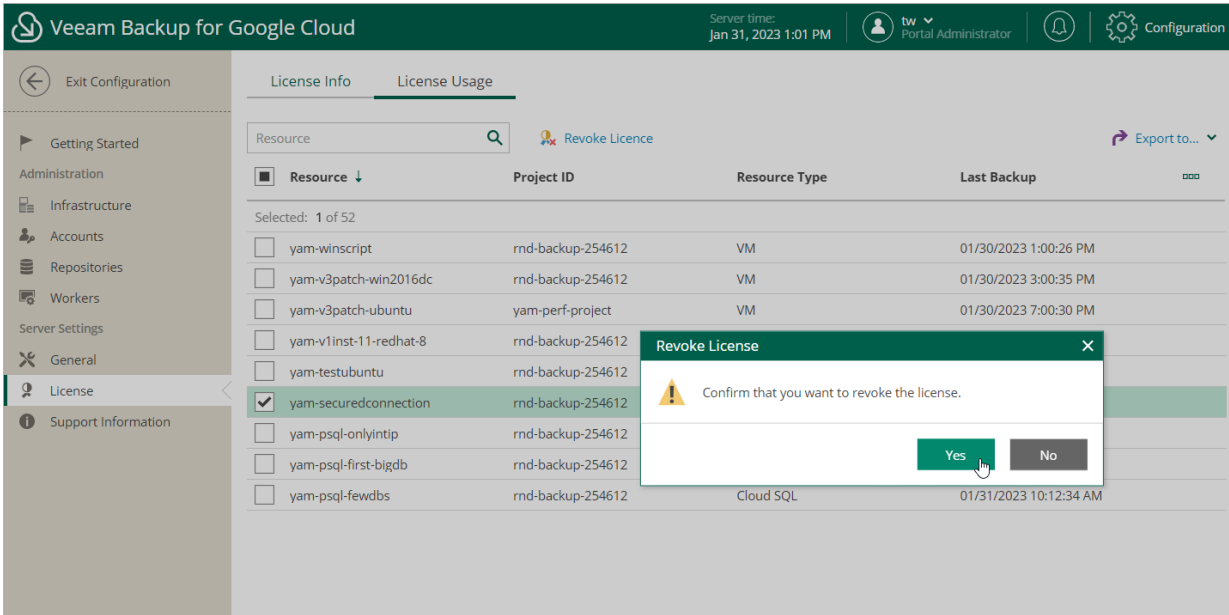


Revoking License Units Using Veeam Backup for Google Cloud Web UI

To revoke a license unit from a protected instance in the Veeam Backup for Google Cloud Web UI, do the following:

1. Switch to the **Configuration** page.
2. Navigate to **License > 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.



Accessing Veeam Backup for Google Cloud

After you install Veeam Backup for Google Cloud and [add backup appliances](#) to the backup infrastructure, you will be able to back up and restore Google Cloud resources using both the Veeam Backup & Replication console and the Veeam Backup for Google Cloud 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](#).

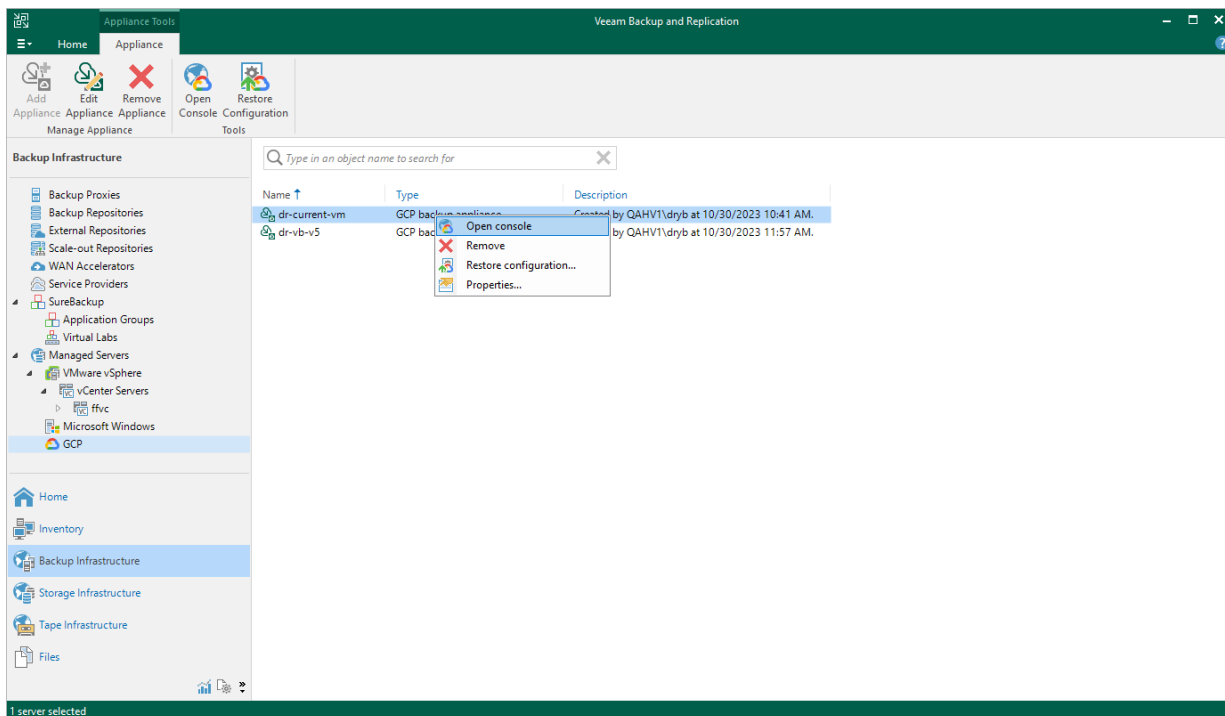
Accessing Web UI from Console

To access the Veeam Backup for Google Cloud 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 Google Cloud Web UI in your default web browser.



Accessing Web UI from Workstation

To access the Veeam Backup for Google Cloud Web UI from a workstation, do the following:

1. In a web browser, navigate to the Veeam Backup for Google Cloud web address.

IMPORTANT

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

The address consists of a public IPv4 address or DNS hostname of the backup appliance. Note that the website is available over HTTPS only.

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 Web Certificates](#).

2. In the **Username** and **Password** fields, specify credentials of an authorized user account.

If you log in for the first time, use credentials of the Default Administrator account that was created [after the product installation](#). In future, you can add other user accounts to grant access to Veeam Backup for Google Cloud. For more information, see [Managing User Accounts](#).

TIP

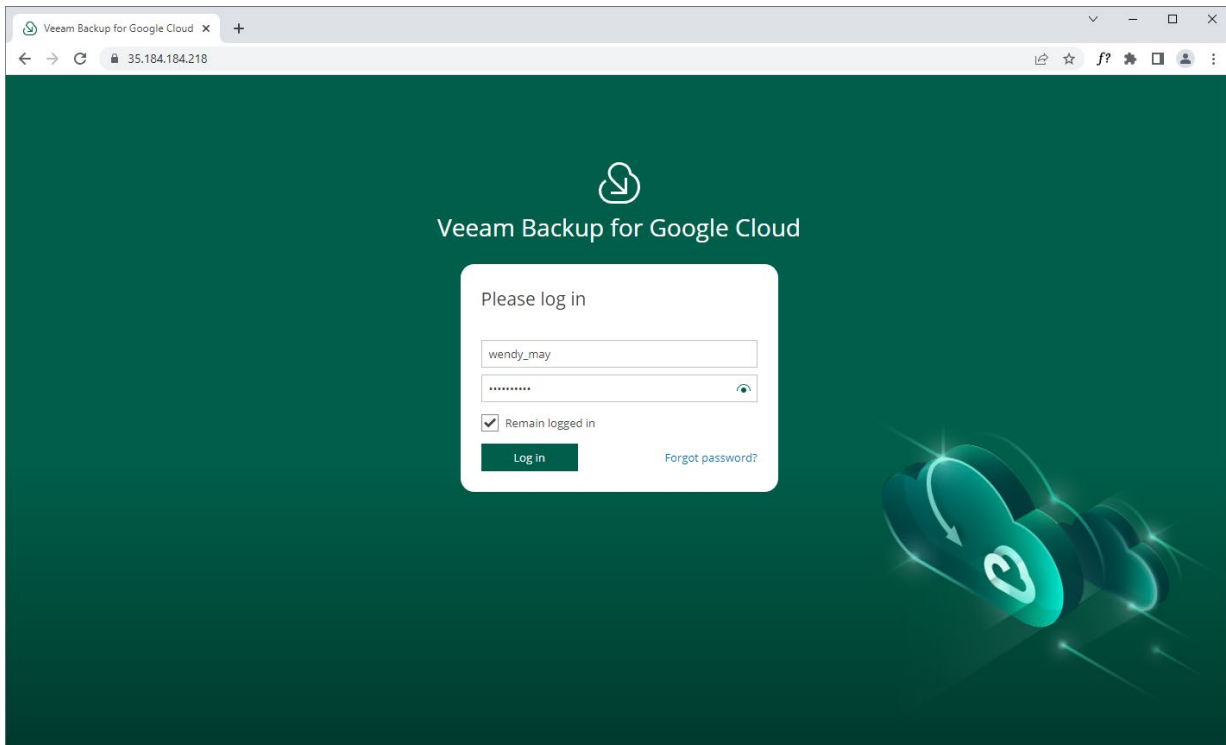
If you do not remember the password, you can reset it. To do that, click the **Forgot password?** link and follow the instructions provided in [this Veeam KB article](#).

3. Select the **Remain logged in** check box to save the specified credentials in a persistent browser cookie so that your session does not expire after 60 minutes of inactivity.

If you select this check box, you will be logged in for 7 days and will not have to provide credentials every time you access the Veeam Backup for Google Cloud Web UI in a new browser session.

4. Click **Log in**.

If **multi-factor authentication (MFA) is enabled** for the user, Veeam Backup for Google Cloud 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 Out

To log out, at the top right corner of the Veeam Backup for Google Cloud window, click the user name and then click **Log out**.

Configuring Veeam Backup for Google Cloud

To start working with Veeam Backup for Google Cloud, 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.](#)
3. Configure the added backup appliances:
 - a. [Add service accounts to authorize requests to Google Cloud APIs.](#)
 - b. [Add projects and folders to get access to Google Cloud resources that you want to protect.](#)
 - c. [\[Optional\] Add users to control access to Veeam Backup for Google Cloud.](#)
 - d. [Create worker configurations.](#)
 - e. [\[Optional\] Configure global retention, email notification and Google authentication settings.](#)

NOTE

Even after you add projects that manage your Google Cloud resources and configure all the necessary settings, Veeam Backup for Google Cloud will not populate the lists of VM, Cloud SQL and Cloud Spanner instances on the [Resources tab](#) — unless you create backup policies and specify regions where the instances belong, as described in sections [Performing VM Backup](#), [Performing SQL Backup](#) and [Performing Spanner Backup](#).

Managing Backup Appliances

Google Cloud 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 Google Cloud 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 Google Cloud appliances](#) from the Veeam Backup & Replication console.
- [Connect to existing Veeam Backup for Google Cloud appliances](#) if you have already deployed them.

NOTE

One backup appliance can be managed by one backup server only. If you add the appliance to the backup infrastructure of another backup server, the synchronization between the appliance and the previous backup server will be terminated, and appliance will be displayed as unavailable.

Connecting to Existing Appliances


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

1. [Launch the New Veeam Backup for GCP Appliance wizard.](#)
2. [Choose a deployment mode.](#)
3. [Specify a service account that will be used to connect the appliance.](#)
4. [Select 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 GCP Appliance Wizard

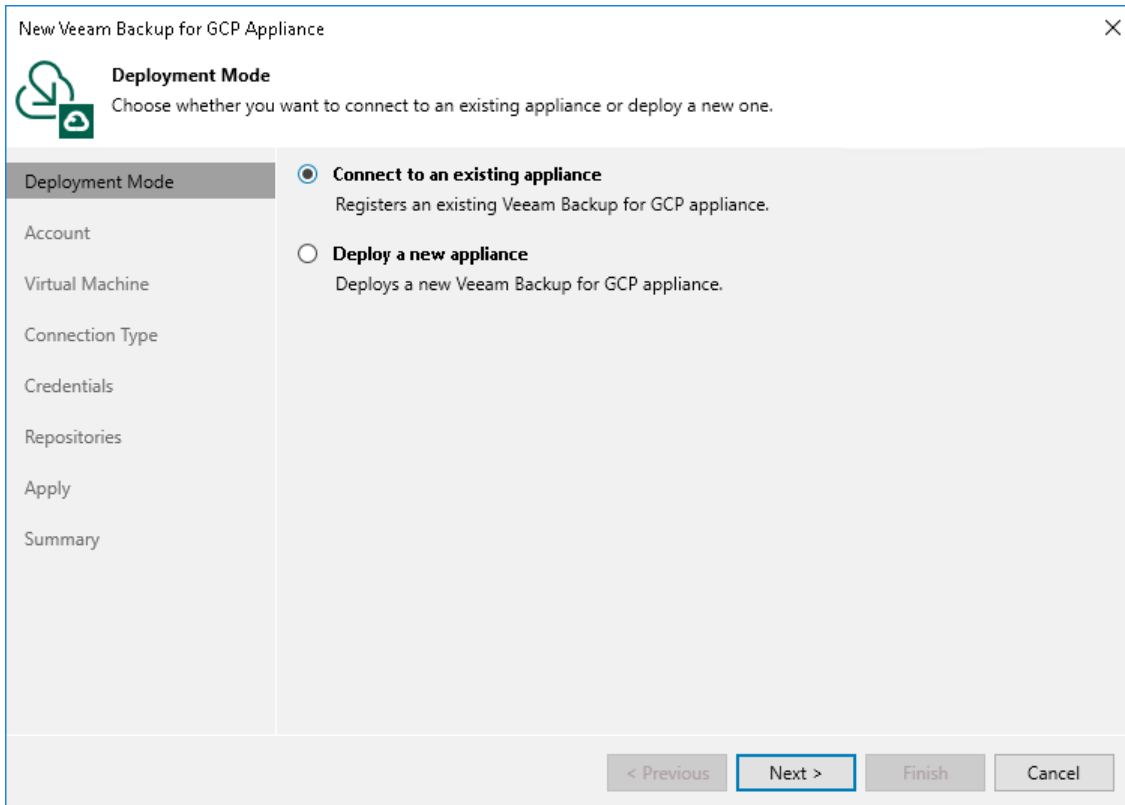
To launch the **New Veeam Backup for GCP Appliance** wizard, do the following:

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 several cloud plug-ins are installed] Click **Veeam cloud-native backup appliance**.
 - b. Choose **Veeam Backup for GCP**.



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 GCP Appliance" with a close button (X) in the top right corner. The main heading is "Deployment Mode" with a sub-instruction: "Choose whether you want to connect to an existing appliance or deploy a new one." Below this, there are two radio button options:

- Connect to an existing appliance**
Registers an existing Veeam Backup for GCP appliance.
- Deploy a new appliance**
Deploys a new Veeam Backup for GCP appliance.

A vertical sidebar on the left lists the wizard steps: "Deployment Mode" (highlighted), "Account", "Virtual Machine", "Connection Type", "Credentials", "Repositories", "Apply", and "Summary". At the bottom right, there are four buttons: "< Previous" (disabled), "Next >" (active/highlighted), "Finish" (disabled), and "Cancel" (disabled).

Step 3. Specify Service Account Settings

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

1. From the **GCP service account** drop-down list, select a service account whose permissions will be used to connect the backup appliance.

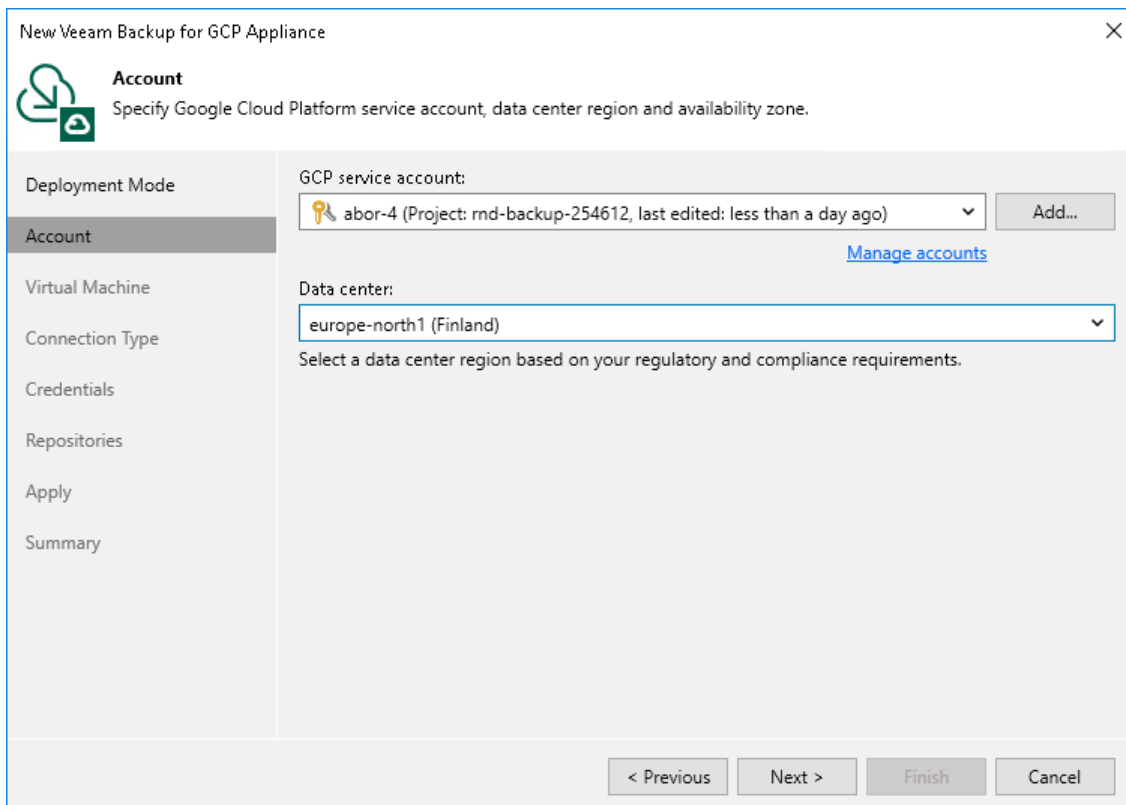
For a service account to be displayed in the **GCP service account** drop-down list, it must be created in Google Cloud and added to the Cloud Credentials Manager. You can create a new service account or add an existing service account to the Cloud Credentials Manager without closing the wizard. To do that, follow the instructions provided in the Veeam Backup & Replication User Guide, section [Google Cloud Platform Service Accounts](#).

NOTE

When you create a service account using the Veeam Backup & Replication console, the service account is automatically assigned the Owner IAM role with a wide scope of permissions and capabilities. If you want the service account to be assigned a limited list of permissions, create a service account manually in Google Cloud beforehand and then add it to the Cloud Credentials Manager. For more information on required permissions that must be assigned to the service account, see [Plug-In Permissions](#).

2. From the **Data center** drop-down list, select the Google Cloud region in which the backup appliance resides.

For more information on regions and zones in Google Cloud, see [Google Cloud documentation](#).

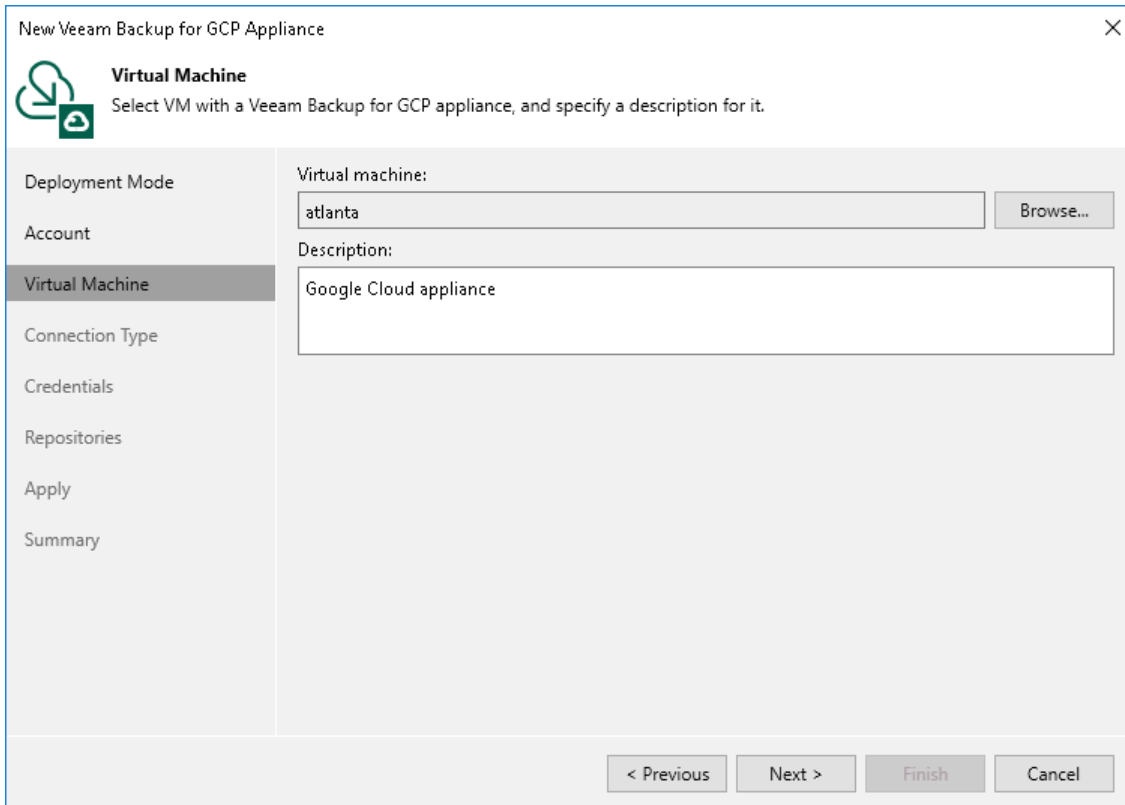


The screenshot shows the 'New Veeam Backup for GCP Appliance' wizard window, specifically the 'Account' step. The window title is 'New Veeam Backup for GCP Appliance' with a close button (X) in the top right corner. Below the title bar, there is a Veeam logo and the text 'Account' followed by the instruction 'Specify Google Cloud Platform service account, data center region and availability zone.' A left-hand navigation pane lists several steps: 'Deployment Mode', 'Account' (which is highlighted), 'Virtual Machine', 'Connection Type', 'Credentials', 'Repositories', 'Apply', and 'Summary'. The main content area contains two dropdown menus. The first is labeled 'GCP service account:' and has a search icon on the left and an 'Add...' button on the right. The selected value is 'abor-4 (Project: rnd-backup-254612, last edited: less than a day ago)'. Below this dropdown is a blue link that says 'Manage accounts'. The second dropdown is labeled 'Data center:' and has the selected value 'europe-north1 (Finland)'. Below this dropdown is the text 'Select a data center region based on your regulatory and compliance requirements.' At the bottom of the window, there are four buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'.

Step 4. Select Appliance

At the **Virtual Machine** step of the wizard, select the backup appliance that you want to add to the backup infrastructure:

1. Click **Browse**.
2. In the **Select Virtual Machine** 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 GCP Appliance" with a close button (X) in the top right corner. The window has a sidebar on the left with the following steps: Deployment Mode, Account, **Virtual Machine** (highlighted), Connection Type, Credentials, Repositories, Apply, and Summary. The main area is titled "Virtual Machine" and contains the instruction: "Select VM with a Veeam Backup for GCP appliance, and specify a description for it." Below this, there are two input fields: "Virtual machine:" with the text "atlanta" and a "Browse..." button, and "Description:" with the text "Google Cloud appliance". At the bottom of the window, there are four buttons: "< Previous", "Next >", "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 network with the inbound internet access allowed and you want the backup server to connect to this appliance over the internet. In this case, Veeam Backup & Replication will detect the public IP address of the appliance automatically.
- Select the **Private network** option if the backup appliance and the backup server are connected to the same private network, or you want the backup server to connect to this appliance over VPN. In this case, you must specify the private IP address or the DNS hostname of the appliance in the **Specify the IP address or DNS name of the appliance** field.

New Veeam Backup for GCP Appliance

Connection Type
Specify if the Veeam Backup for Veeam Backup for GCP Appliance is connected to the Internet.

Deployment Mode

Account

Virtual Machine

Connection Type

Credentials

Repositories

Apply

Summary

Direct connection
The backup server will identify the IP address automatically.

Private network
Specify the IP address or DNS name of the appliance:

< Previous Next > Finish Cancel

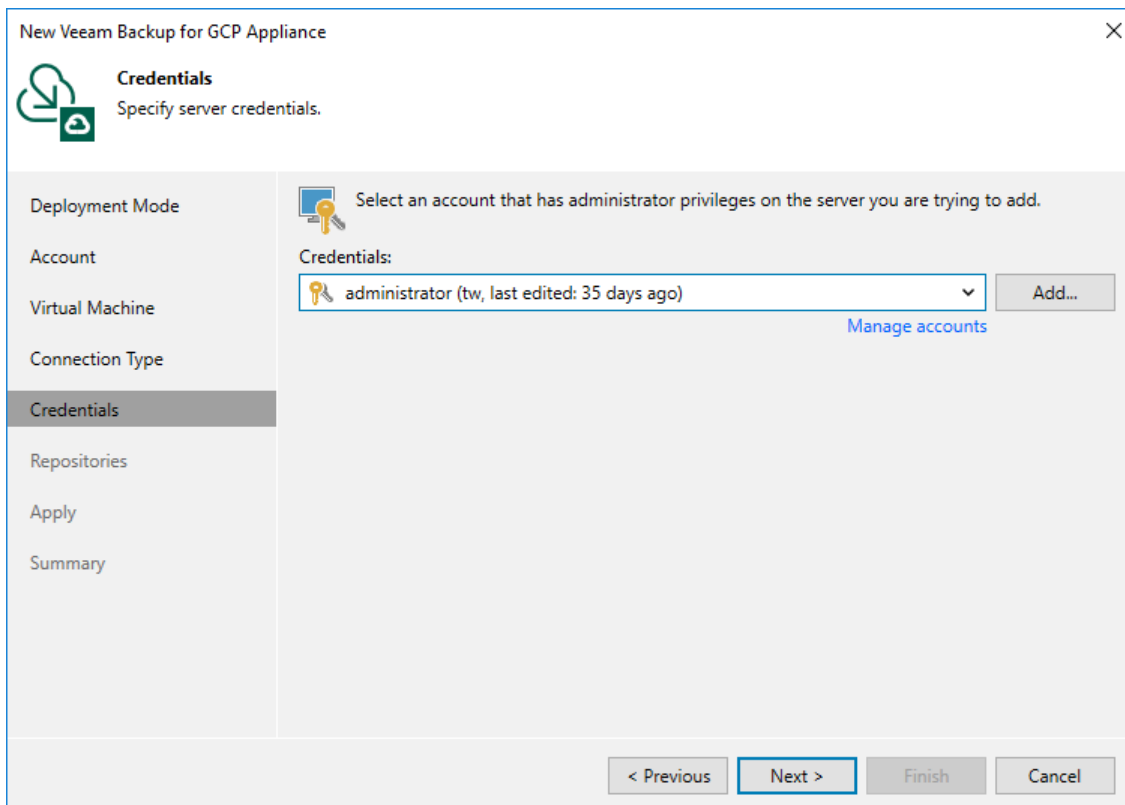
Step 6. Specify User Credentials

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 GCP 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 user must have multi-factor authentication (MFA) disabled and the Portal Administrator role assigned.



The screenshot shows the 'New Veeam Backup for GCP Appliance' wizard at the 'Credentials' step. The window title is 'New Veeam Backup for GCP Appliance' with a close button (X) in the top right corner. The main heading is 'Credentials' with the subtitle 'Specify server credentials.' Below this, there is a navigation pane on the left with options: Deployment Mode, Account, Virtual Machine, Connection Type, Credentials (selected), Repositories, Apply, and Summary. The main content area has a key icon and the instruction 'Select an account that has administrator privileges on the server you are trying to add.' Underneath, there is a 'Credentials:' label followed by a dropdown menu showing 'administrator (tw, last edited: 35 days ago)' and a small 'Add...' button to its right. Below the dropdown is a blue link labeled 'Manage accounts'. At the bottom of the window, there are four buttons: '< Previous', 'Next >' (highlighted with a blue border), 'Finish', and 'Cancel'.

NOTE

As soon as you click **Next**, Veeam Backup & Replication will verify connection to the specified backup appliance. If the version of the appliance is not compatible with the Veeam Backup & Replication version or if the TLS certificate used to connect to the Veeam Backup for Google Cloud 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 whose version is not compatible with the Veeam Backup & Replication version, 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.

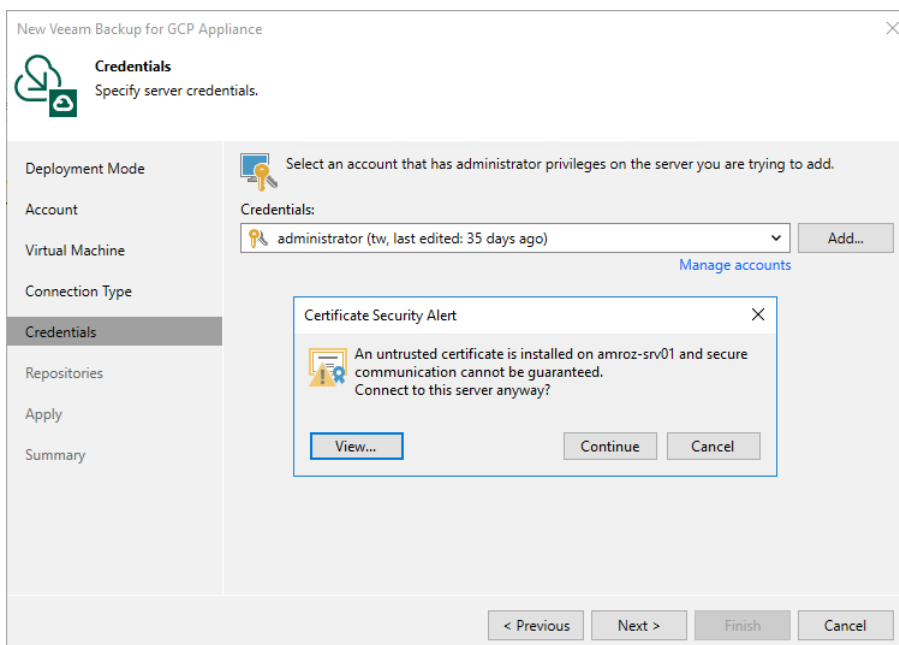
Certificate Security Alert

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. To learn how to manage 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 replace a TLS certificate installed on a backup appliance, this appliance becomes unavailable in the Veeam Backup & Replication console. To make the appliance available again, [modify the appliance settings](#) to acknowledge the new certificate. For Veeam Backup & Replication to be able further to automatically update the TLS certificate in the Veeam Backup & Replication configuration database, make sure that ingress traffic [is allowed from the Google IAP](#) through the SSH protocol (IP range 35.235.240.0/20) on the appliance.



Step 7. Configure Repository Settings

At the **Repositories** step of the wizard, a list of all standard and archive 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 will be able to use the Veeam Backup & Replication console to perform [entire VM instance restore](#), [entire SQL instance restore](#) and [entire Spanner instance restore](#) only — unless you 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 repositories of the *Standard* and *Nearline* storage classes. For repositories of the *Archive* storage class, there is no possibility to specify any configuration settings.

1. In the **Repositories** list, select the necessary standard repository and click **Edit**.
2. In the **Repository** window:
 - a. From the **Credentials** drop-down list, select a Hash-based Message Authentication Code (HMAC) key associated with the service account that will be used to access the repository.

For an HMAC key to be displayed in the **Credentials** list, it must be added to the [Cloud Credentials Manager](#). If you have not added the key 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 HMAC key access ID and secret in the **Credentials** window.
 - 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 [Architecture Overview](#).
 - c. If encryption is enabled for the repository, 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 passwords** link or the **Add** button, and specify a hint and the password in the **Password** window.


If you do not specify a password for a standard repository with encryption enabled, you will have to decrypt data stored in this repository manually as described in section [Decrypting Backups](#).

After you finish working with the wizard, all the 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](#).

New Veeam Backup for GCP Appliance ×

 **Repositories**
The following repositories are available on the specified Veeam Backup for GCP appliance.

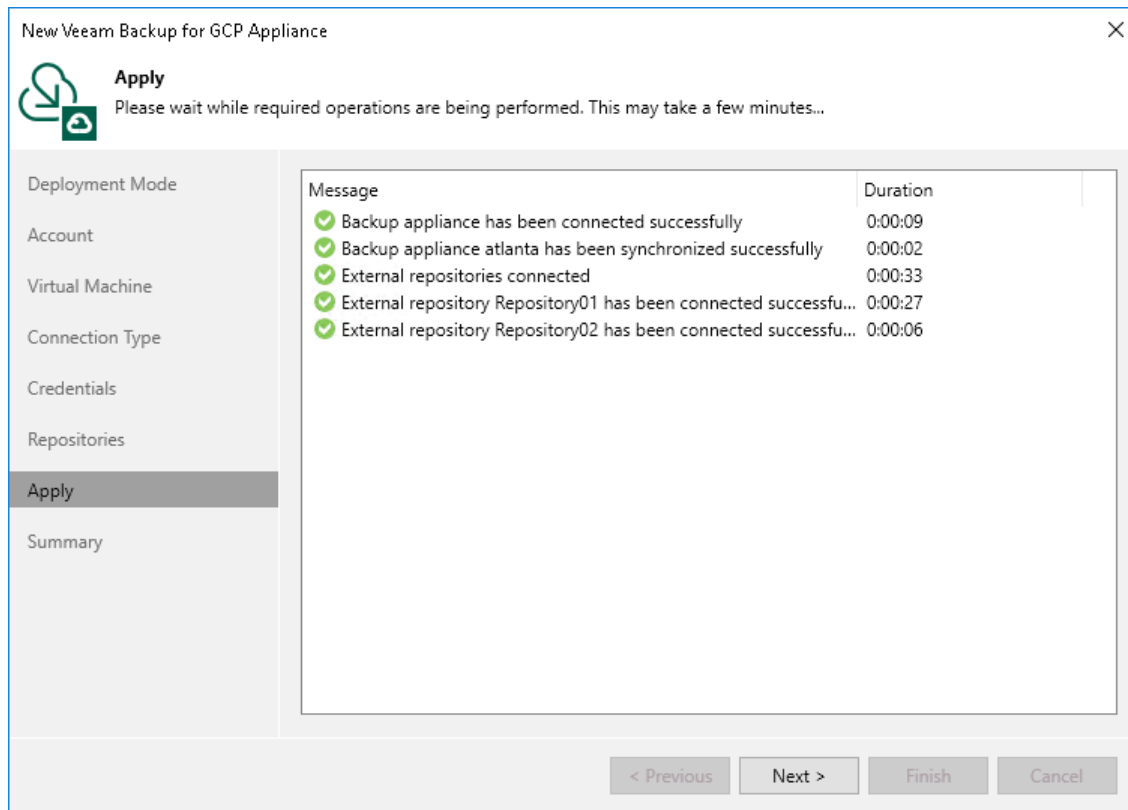
Deployment Mode
Account
Virtual Machine
Connection Type
Credentials
Repositories
Apply
Summary

Repositories:

Repository	Type	Credentials	Encryption password	Edit...
 Repository01	Standard	GOOG1EXDVJ4IDB...	standard (Last edite...	<input type="button" value="Edit..."/>
 Repository02	Archive	N/A	N/A	

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 GCP Appliance

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


Message	Duration
✓ Backup appliance has been connected successfully	0:00:09
✓ Backup appliance atlanta has been synchronized successfully	0:00:02
✓ External repositories connected	0:00:33
✓ External repository Repository01 has been connected successfu...	0:00:27
✓ External repository Repository02 has been connected successfu...	0:00:06

< 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 will be able configure its settings in the Veeam Backup for Google Cloud Web UI.

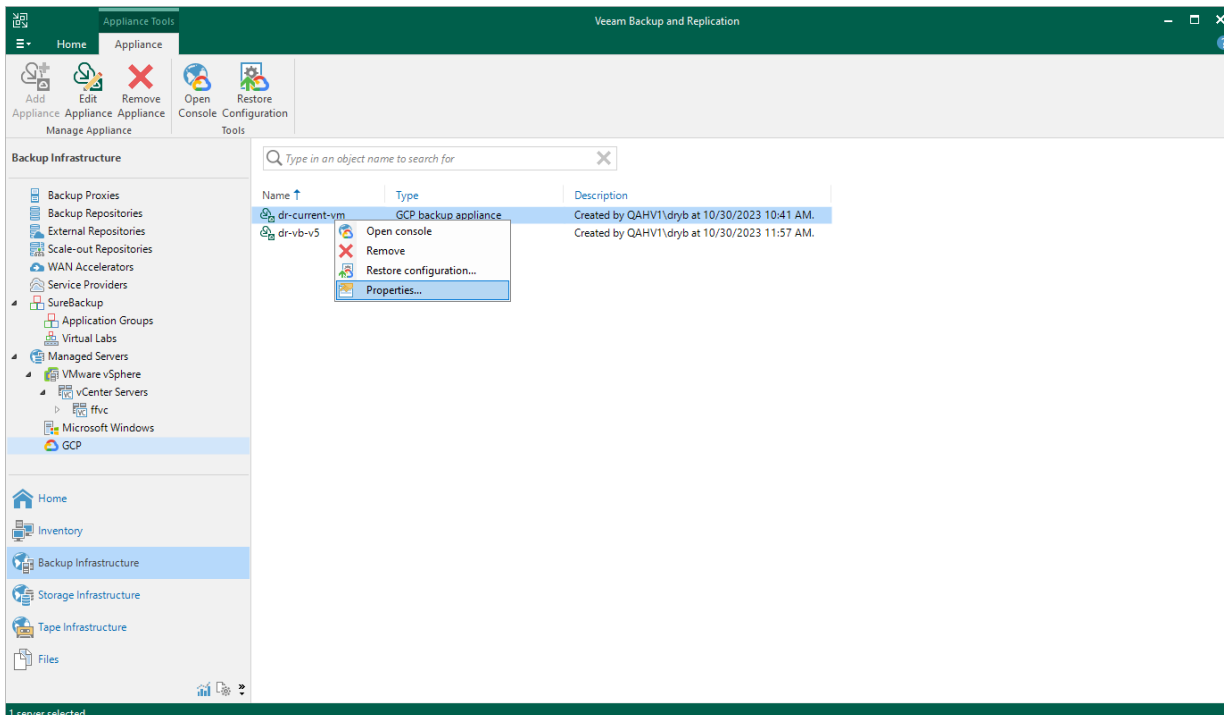
If you want Veeam Backup & Replication to open the Web UI of the added appliance immediately, click the **backup appliance console** link.



Editing Appliance Settings

For each backup appliance managed by the backup server, you can modify the settings configured while adding the appliance to the backup infrastructure:

1. In the Veeam Backup & Replication console, open the **Backup Infrastructure** view.
2. Navigate to **Managed Servers**.
3. Select the necessary backup appliance and click **Edit Appliance** on the ribbon.
Alternatively, you can right-click the appliance and select **Properties**.
4. Complete the **Edit Veeam Backup for GCP Appliance** wizard:
 - a. To change the service account that is used to connect to the appliance, follow the instructions provided in section [Connecting to Existing Appliances](#) (step 1).
 - b. To provide a new description for the appliance, follow the instructions provided in section [Connecting to Existing Appliances](#) (step 4).
 - c. To change the way Veeam Backup & Replication connects to the appliance, follow the instructions provided in section [Connecting to Existing Appliances](#) (step 5).
 - d. To change the user whose credentials Veeam Backup & Replication uses to connect to the appliance, follow the instructions provided in section [Connecting to Existing Appliances](#) (step 6).
 - e. To edit settings of the appliance repositories added to the backup infrastructure, follow the instructions provided in section [Connecting to Existing 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** to confirm the changes.



Rescanning Appliances

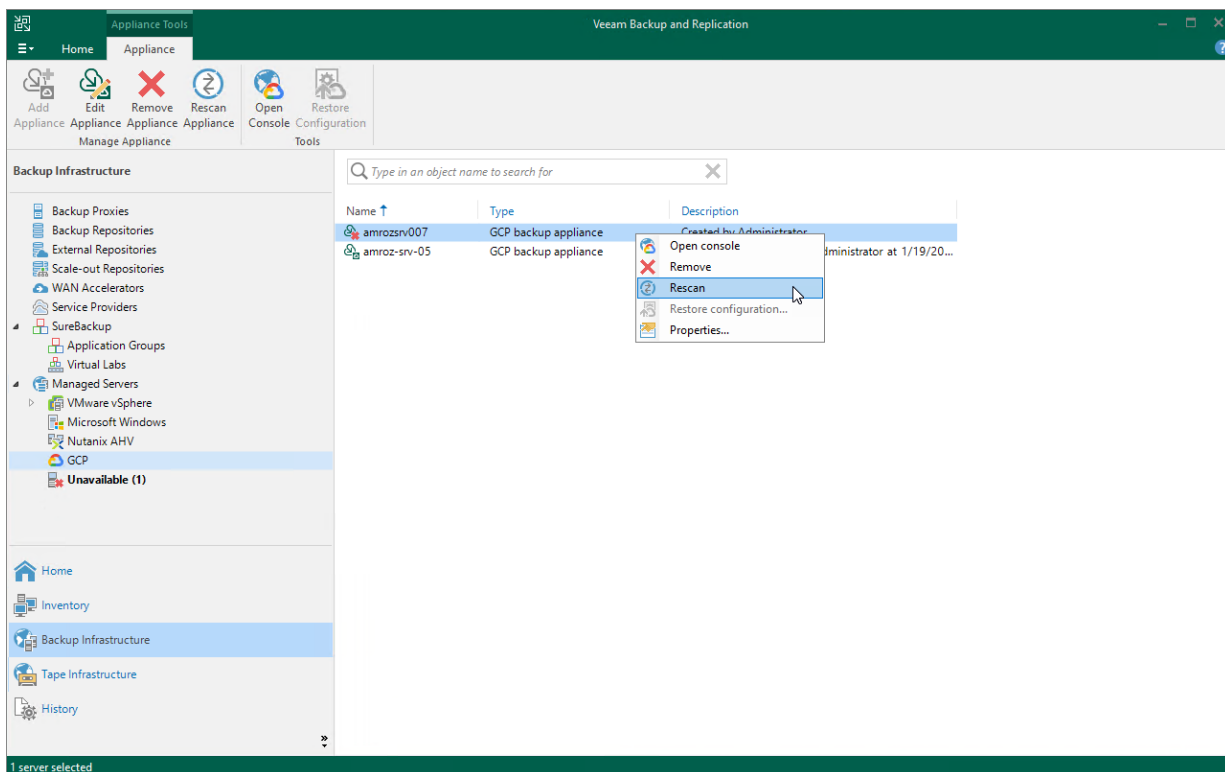
If a backup appliance becomes unavailable, for example, due to connectivity problems, you can rescan the appliance:

1. In the Veeam Backup & Replication console, open the **Backup Infrastructure** view.
2. Navigate to **Managed Servers**.
3. Select the necessary backup appliance and click **Rescan appliance** on the ribbon.
Alternatively, you can right-click the appliance and select **Rescan**.
4. In the opened window, click **Yes**.

Veeam Backup & Replication will remove all data collected from the appliance configuration database. Then, Veeam Backup & Replication will recollect session results for the past 24 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](#).



Removing Appliances

Google Cloud 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 HMAC keys will be removed automatically from the backup infrastructure.
- Repositories for which you have specified HMAC keys 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 VM, Cloud SQL and Cloud Spanner instances from snapshots.
- Restore to Google Cloud from image-level backups will start working as described in the Veeam Backup & Replication User Guide, section [How Restore to Google Compute Engine 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 Google Cloud 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.

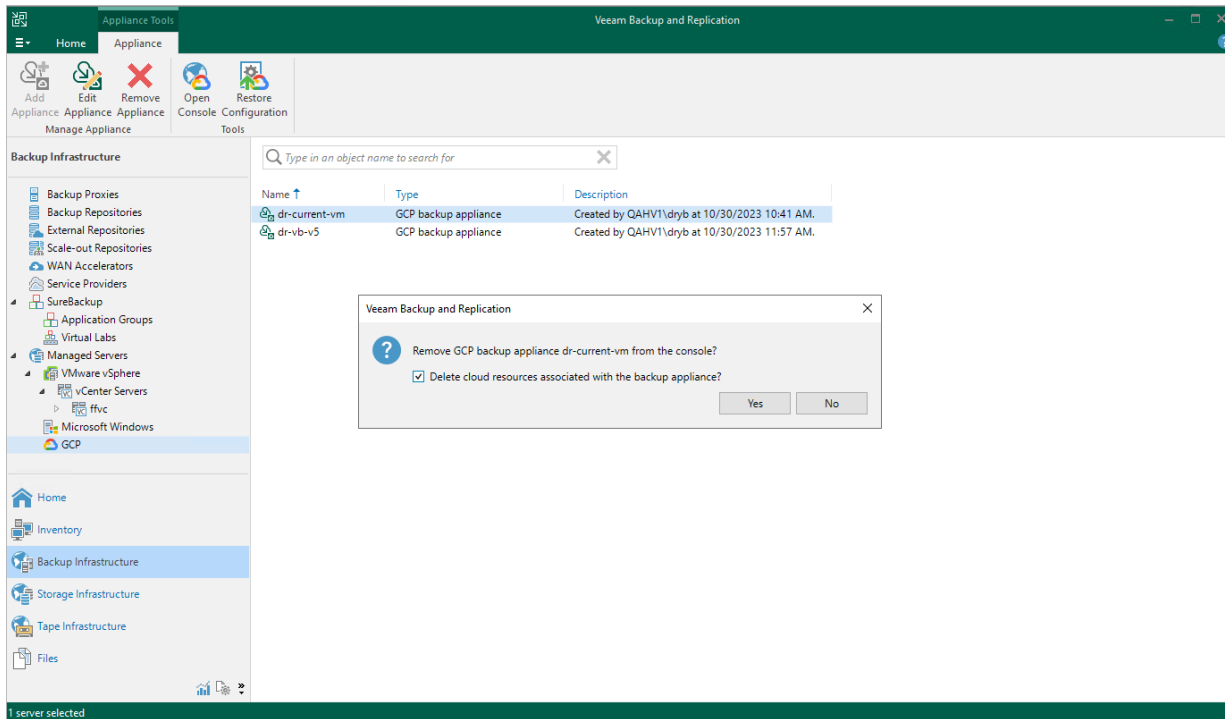
Alternatively, you can right-click the appliance and select **Remove**.

4. In the **Veeam Backup & Replication** window, do either of the following:
 - If you want to remove the appliance from the backup infrastructure but leave it in Google Cloud, do not select the **Delete cloud resources associated with the backup appliance?** check box. Click **Yes**.

In this case, the appliance will continue creating restore points in its repositories according to configured backup policy settings, and you will still be able to use these restore points to perform restore from the Veeam Backup & Replication console. However, you will have to rescan the repositories every time you want to collect information on all newly created restore points, or to update the list of restore points that were removed manually or by retention.

- [Applies only to backup appliances version 3.0 or later] If you want to remove the appliance from both the backup infrastructure and the Google Cloud environment, select the **Delete cloud resources associated with the backup appliance?** check box. Then, click **Yes**.

In this case, Veeam Backup for Google Cloud will remove all resources associated with this appliance in Google Cloud.



Managing Backup Repositories

Veeam Backup for Google Cloud uses Google Cloud storage buckets as target locations for image-level backups of VM, Cloud SQL and Cloud Spanner instances, and for backups of the configuration database. To store backups in storage buckets, configure backup repositories. A repository is a specific subdirectory created by Veeam Backup for Google Cloud in a storage bucket.

IMPORTANT

A backup repository must not be managed by multiple backup appliances simultaneously. Retention sessions running on different appliances may corrupt backups stored in the repository, which may result in unpredictable data loss.

Adding Backup Repositories Using Console

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

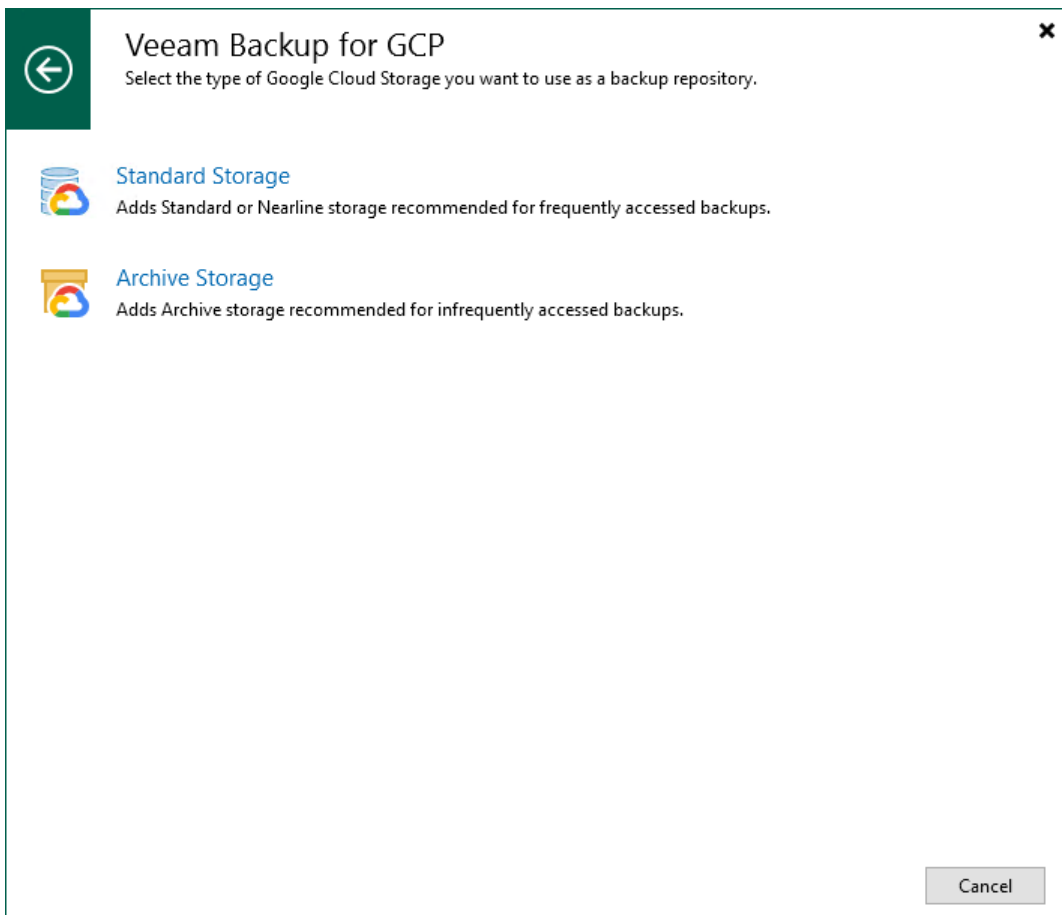
To add a new repository, do the following:

1. [Launch the Add External Repository wizard.](#)
2. [Specify an appliance, and provide a repository name and description.](#)
3. [Specify a project for the repository.](#)
4. [Specify a service account that will be used to access the project.](#)
5. [Configure repository settings.](#)
6. [Wait for the repository to be added to the backup infrastructure.](#)
7. [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 GCP**.
 - b. Choose whether you want to create a standard or an archive repository:
 - § Select the **Standard Storage** option if you want to create a repository with the *Standard* or *Nearline* storage class assigned.
 - § Select the **Archive Storage** option if you want to create a repository with the *Archive* storage class assigned.



Step 2. Specify Repository Details

At the **Veeam Backup for GCP** 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 [Adding Appliances](#).
2. Use the **Repository name** and **Description** fields to enter a name for the new repository and to provide a description for future reference. The maximum length of the name is 127 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 [step 5](#) of the wizard. This folder will be used to store backed-up data.

Add External Repository

Veeam Backup for GCP
Specify the Veeam Backup for Google Cloud Platform appliance to create the backup repository for.

Veeam Backup for GCP

Project

Service Account

Bucket

Apply

Summary

Appliance: amroz-srv01

Repository name: Repository03

Description: Secondary repository

< Previous **Next >** Finish Cancel

Step 3. Specify Project

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

1. From the **Project** drop-down list, select a project where the new repository will belong.
For a project to be displayed in the **Project** list, it must be added to the backup appliance as described in section [Adding Projects and Folders](#).
2. [Applies only if you have chosen to create a standard 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 [Architecture Overview](#).

Add External Repository

Project
Specify a Google Cloud project where the repository will be created.

Veeam Backup for GCP

Project

Service Account

Bucket

Apply

Summary

Project:
RnD Backup

Specify a Google Cloud project.

Gateway server:
srv12win16.tech.local (Backup server)

Select a gateway server to proxy access to Google Cloud Storage bucket with backup files. The server will store a cache of backup metadata for enhanced performance.

< Previous **Next >** Finish Cancel

Step 4. Specify Service Account

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

1. From the **Service account** drop-down list, select a service account whose permissions will be used to access the project specified at [step 3](#) of the wizard.

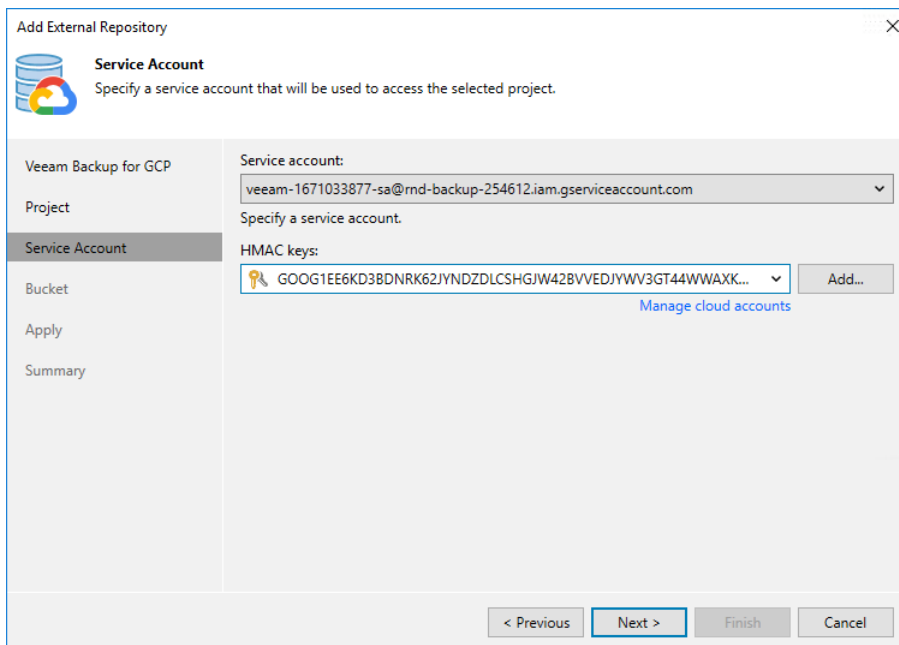
For a service account to be displayed in the list of available accounts, it must be added to the selected backup appliance as described in section [Adding Service Accounts](#).

2. From the **HMAC keys** drop-down list, select a Hash-based Message Authentication Code (HMAC) key that will be used by Veeam Backup & Replication to authenticate requests to the repository. The specified HMAC key can belong to any service account that has permissions to access the project specified at [step 3](#) of the wizard.

For an HMAC key to be displayed in the **Credentials** list, it must be added to the [Cloud Credentials Manager](#). If you have not added the key to the Cloud Credentials Manager beforehand, you can do it without closing the **Add External Repository** wizard. To add an HMAC key, click either the **Manage accounts link** or the **Add** button, and specify the HMAC key access ID and secret in the **Credentials** window.

NOTE

As the backup appliance also needs an HMAC key to authenticate its requests to the repository, Veeam Backup & Replication will verify whether the selected service account has an associated HMAC key saved in the configuration database of the appliance. If the account has no associated HMAC key or the associated HMAC key is not valid, Veeam Backup & Replication will verify whether the selected HMAC key belongs to the selected service account. If yes, Veeam Backup & Replication will add this key to the configuration database of the backup appliance. If no, Veeam Backup & Replication will generate a new HMAC key for the selected account and add this key to the appliance configuration database.



The screenshot shows the 'Add External Repository' wizard window, specifically the 'Service Account' step. The window title is 'Add External Repository' and it has 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 GCP', 'Project', 'Service Account' (which is selected and highlighted), 'Bucket', 'Apply', and 'Summary'. The main area of the wizard is titled 'Service Account' and contains the following fields and controls:

- A sub-header 'Service Account' with the instruction 'Specify a service account that will be used to access the selected project.'
- A 'Service account:' dropdown menu with the selected value 'veeam-1671033877-sa@rmd-backup-254612.iam.gserviceaccount.com'.
- A sub-header 'Specify a service account.'
- An 'HMAC keys:' dropdown menu with the selected value 'GOOG1EE6KD3BDNRK62JYNDZDLCSHGJW42BVVEDJYWV3GT44WWAXK...'. To the right of this dropdown is an 'Add...' button.
- A link labeled 'Manage cloud accounts' below the HMAC keys dropdown.

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

Step 5. Configure Repository Settings

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

1. Choose whether you want to use an existing bucket or to create a new one as the target location for image-level backups of VM, Cloud SQL and Cloud Spanner instances:

- o To specify an existing bucket, enter the name of a storage 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 the Google Cloud for the project specified at [step 3](#) of the wizard, as described in [Google Cloud documentation](#).

- o [Applies only if you have specified the project to which the backup appliance belongs] 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 same region where the backup appliance resides.

2. [Applies only if you have chosen to create a standard repository] When you create a standard repository, backups are stored in a high-performance, short-term *Standard* storage class by default. To store backups in a cost-effective, high-durable *Nearline* storage that you plan to access infrequently, select the **Use nearline storage class** check box. Note that after the repository is created, you will not be able to change its storage class.


3. Use the **Enable backup file encryption** check box to choose whether you want to encrypt backups stored in the created repository. If you enable encryption, also specify a password that will be used to encrypt data.

For a password to be displayed in the list of available passwords, 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 **Bucket** 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

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 Repositories](#).

Add External Repository ⋮ ×

 **Bucket**
Specify Google Cloud Storage bucket to connect to.

Veeam Backup for GCP

Project

Service Account

Bucket

Apply

Summary

Bucket:

Use nearline storage class (may result in higher costs)
With lower price per GB but higher retrieval and early deletion fees, this storage class is best suited for long-term restore points, such as GFS fulls.

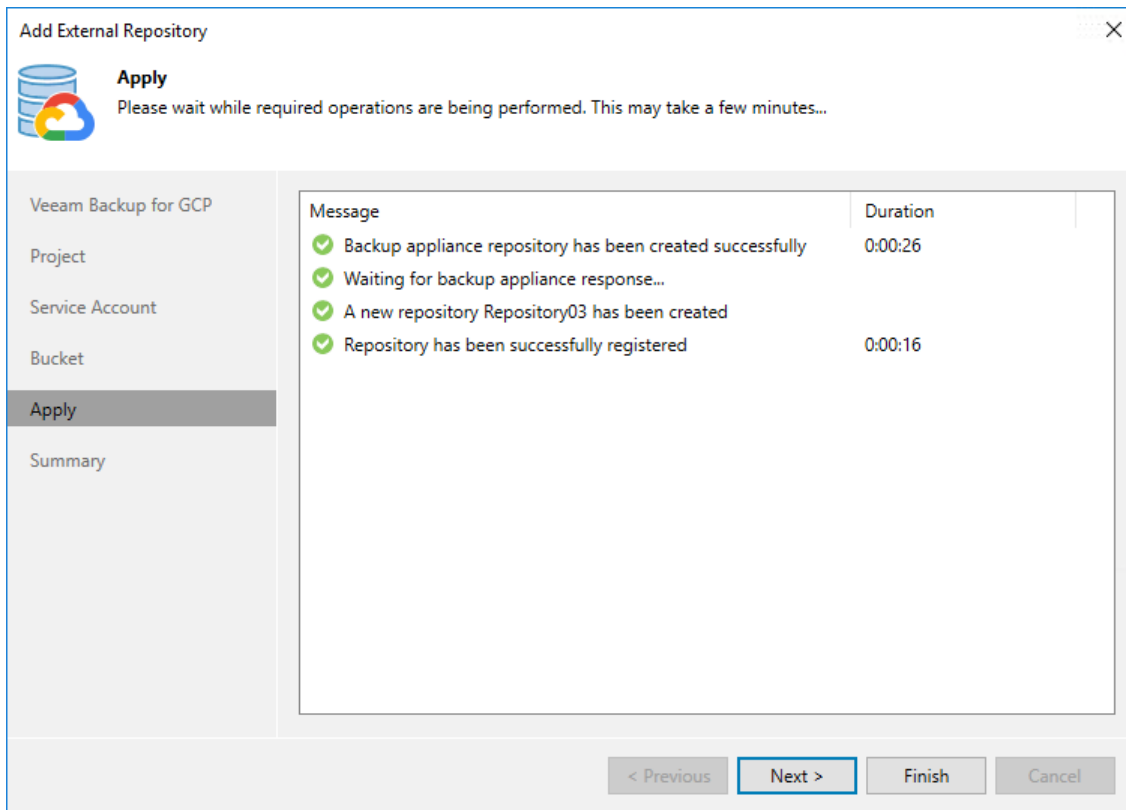
Enable backup file encryption:

[Manage passwords](#)

< Previous Finish Cancel

Step 6. 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**.



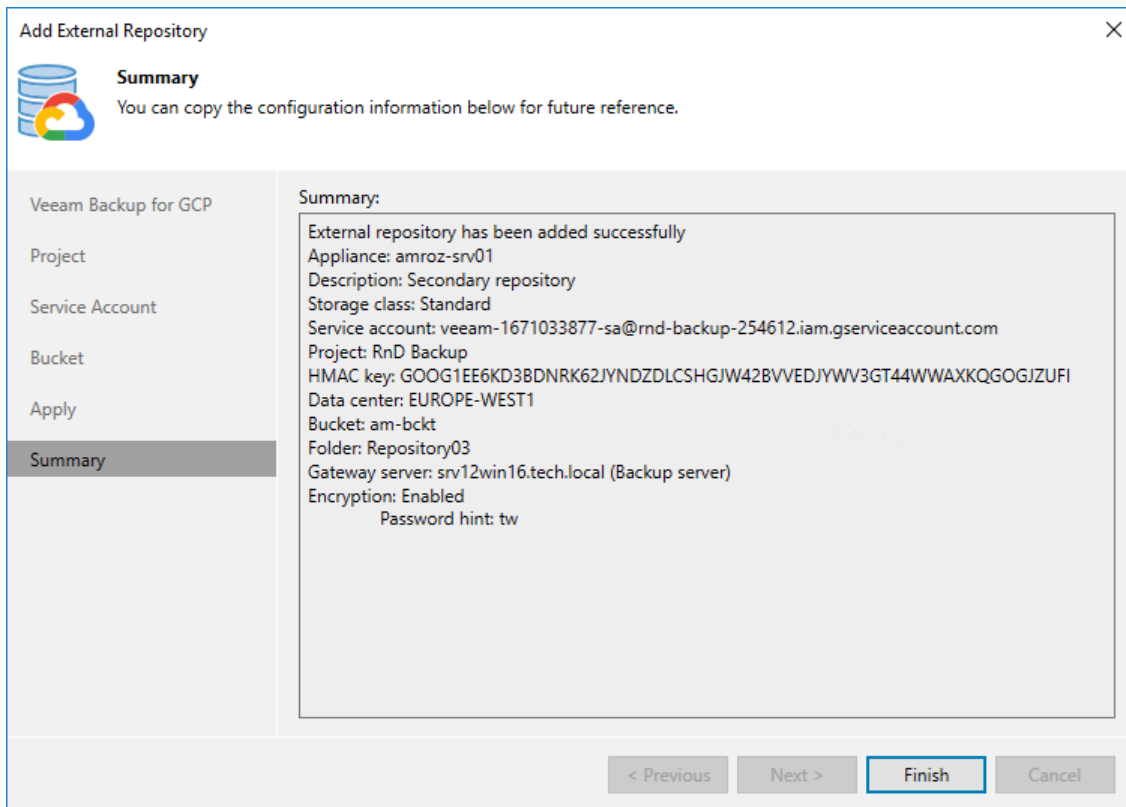
The screenshot shows the 'Add External Repository' wizard window. The title bar reads 'Add External Repository' with a close button. Below the title bar is a 'Veeam Backup for GCP' logo and the word 'Apply' in bold. A message says 'Please wait while required operations are being performed. This may take a few minutes...'. On the left, a navigation pane lists 'Veeam Backup for GCP', 'Project', 'Service Account', 'Bucket', 'Apply' (highlighted), and 'Summary'. The main area contains a table with the following data:

Message	Duration
✔ Backup appliance repository has been created successfully	0:00:26
✔ Waiting for backup appliance response...	
✔ A new repository Repository03 has been created	
✔ Repository has been successfully registered	0:00:16

At the bottom, there are four buttons: '< Previous', 'Next >' (highlighted), 'Finish', and 'Cancel'.

Step 7. 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 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 appliance and select **Properties**.
4. In the **Edit Veeam Backup for GCP Appliance** wizard, do the following:
 - a. Navigate to the **Repositories** step and provide the required information as described in section [Connecting to Existing Appliances](#) (step 7).
 - b. Complete the **Edit Veeam Backup for GCP Appliance** wizard as described in section [Connecting to Existing Appliances](#) (steps 8–9).

Open the **Backup Infrastructure** view to verify that the repository is displayed under the **External Repositories** node.



Adding Backup Repositories Using Web UI

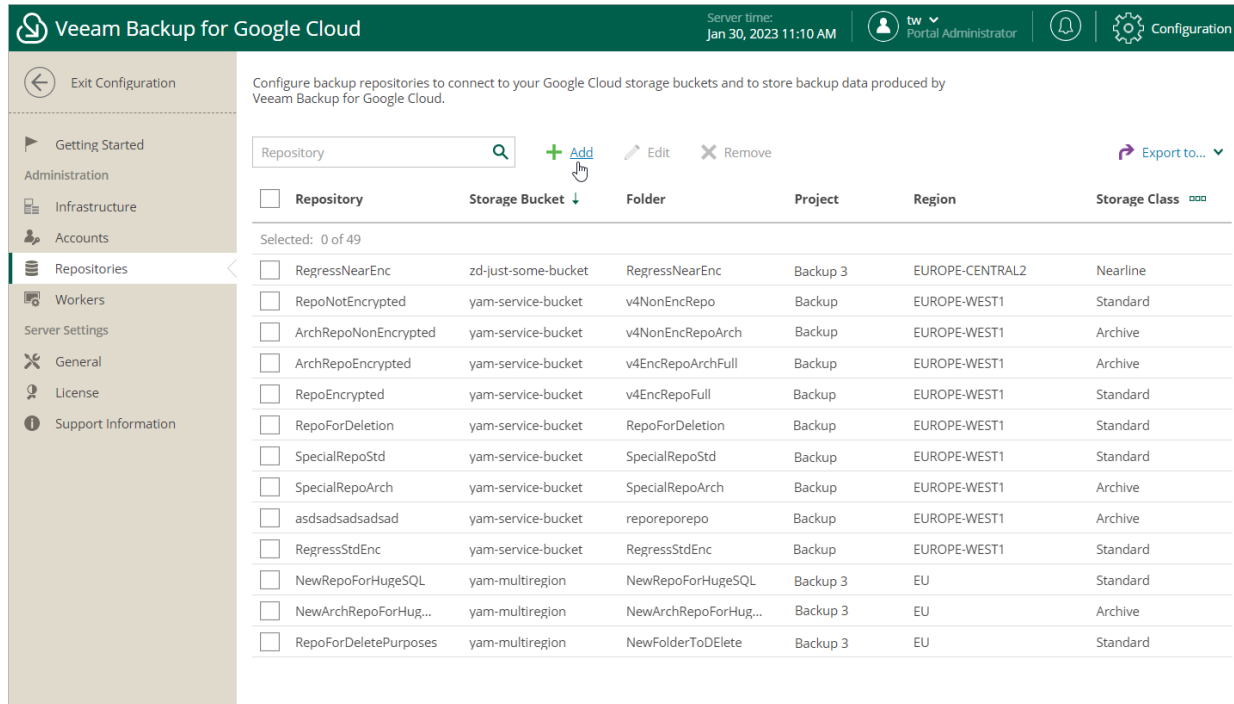
To add a new backup repository, do the following:

1. [Launch the Add Repository wizard.](#)
2. [Specify a repository name and description.](#)
3. [Specify a project for the repository.](#)
4. [Specify a HMAC key for the repository.](#)
5. [Configure repository settings.](#)
6. [Enable encryption for the repository.](#)
7. [Finish working with the wizard.](#)

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**.



The screenshot shows the Veeam Backup for Google Cloud configuration interface. The top navigation bar includes the Veeam logo, the text "Veeam Backup for Google Cloud", the server time "Jan 30, 2023 11:10 AM", the user "tw Portal Administrator", and a "Configuration" gear icon. A left sidebar contains navigation options: "Exit Configuration", "Getting Started", "Administration", "Infrastructure", "Accounts", "Repositories" (highlighted), "Workers", "Server Settings", "General", "License", and "Support Information". The main content area is titled "Configure backup repositories to connect to your Google Cloud storage buckets and to store backup data produced by Veeam Backup for Google Cloud." It features a search bar, an "Add" button (with a mouse cursor over it), "Edit", and "Remove" buttons, and an "Export to..." dropdown. Below this is a table with columns: Repository, Storage Bucket, Folder, Project, Region, and Storage Class. The table lists 15 repositories, including "RegressNearEnc", "RepoNotEncrypted", "ArchRepoNonEncrypted", "ArchRepoEncrypted", "RepoEncrypted", "RepoForDeletion", "SpecialRepoStd", "SpecialRepoArch", "asdsadsadsad", "RegressStdEnc", "NewRepoForHugeSQL", "NewArchRepoForHug...", and "RepoForDeletePurposes".

Repository	Storage Bucket	Folder	Project	Region	Storage Class
<input type="checkbox"/> RegressNearEnc	zd-just-some-bucket	RegressNearEnc	Backup 3	EUROPE-CENTRAL2	Nearline
<input type="checkbox"/> RepoNotEncrypted	yam-service-bucket	v4NonEncRepo	Backup	EUROPE-WEST1	Standard
<input type="checkbox"/> ArchRepoNonEncrypted	yam-service-bucket	v4NonEncRepoArch	Backup	EUROPE-WEST1	Archive
<input type="checkbox"/> ArchRepoEncrypted	yam-service-bucket	v4EncRepoArchFull	Backup	EUROPE-WEST1	Archive
<input type="checkbox"/> RepoEncrypted	yam-service-bucket	v4EncRepoFull	Backup	EUROPE-WEST1	Standard
<input type="checkbox"/> RepoForDeletion	yam-service-bucket	RepoForDeletion	Backup	EUROPE-WEST1	Standard
<input type="checkbox"/> SpecialRepoStd	yam-service-bucket	SpecialRepoStd	Backup	EUROPE-WEST1	Standard
<input type="checkbox"/> SpecialRepoArch	yam-service-bucket	SpecialRepoArch	Backup	EUROPE-WEST1	Archive
<input type="checkbox"/> asdsadsadsad	yam-service-bucket	reporeporepo	Backup	EUROPE-WEST1	Archive
<input type="checkbox"/> RegressStdEnc	yam-service-bucket	RegressStdEnc	Backup	EUROPE-WEST1	Standard
<input type="checkbox"/> NewRepoForHugeSQL	yam-multiregion	NewRepoForHugeSQL	Backup 3	EU	Standard
<input type="checkbox"/> NewArchRepoForHug...	yam-multiregion	NewArchRepoForHug...	Backup 3	EU	Archive
<input type="checkbox"/> RepoForDeletePurposes	yam-multiregion	NewFolderToDElete	Backup 3	EU	Standard

Step 2. Specify Repository Name and Description


At the **Repository Info** step of the wizard, use the **Name** and **Description** fields to enter a name for the new backup repository and to provide a description for future reference. The maximum length of the name is 127 characters; the following characters are not supported: \ / " ' [] : | < > + = ; , ? * @ & _ .

The screenshot shows the 'Add Repository' wizard in Veeam Backup for Google Cloud. 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 Google Cloud', the server time 'Jan 30, 2023 11:19 AM', the user 'tw Portal Administrator', and a 'Configuration' link. The left sidebar contains a list of steps: 'Repository Info' (highlighted), 'Project', 'Service Account', 'Storage Bucket', 'Encryption', and 'Summary'. The main content area is titled 'Specify repository name and description' and contains two input fields: 'Name:' with the value 'backup-repo-01' and 'Description:' with the value 'primary backup repository'. At the bottom right, there are 'Next' and 'Cancel' buttons.

Step 3. Specify Project

At the **Project** step of the wizard, select a project to which the new backup repository will belong.

For a project to be displayed in the **Project** list, it must be added to Veeam Backup for Google Cloud as described in section [Adding Projects and Folders](#). If you have not added the necessary project to Veeam Backup for Google Cloud beforehand, you can do it without closing the **Add Repository** wizard. To add a repository, click **Add** and complete the **Add Projects and Folders** wizard.



Step 4. Specify Service Account

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

1. In the **Service account** section, click **Choose** to select a service account whose permissions Veeam Backup for Google Cloud will use to access the specified project. For more information on the required permissions, see [Service Account Permissions](#).

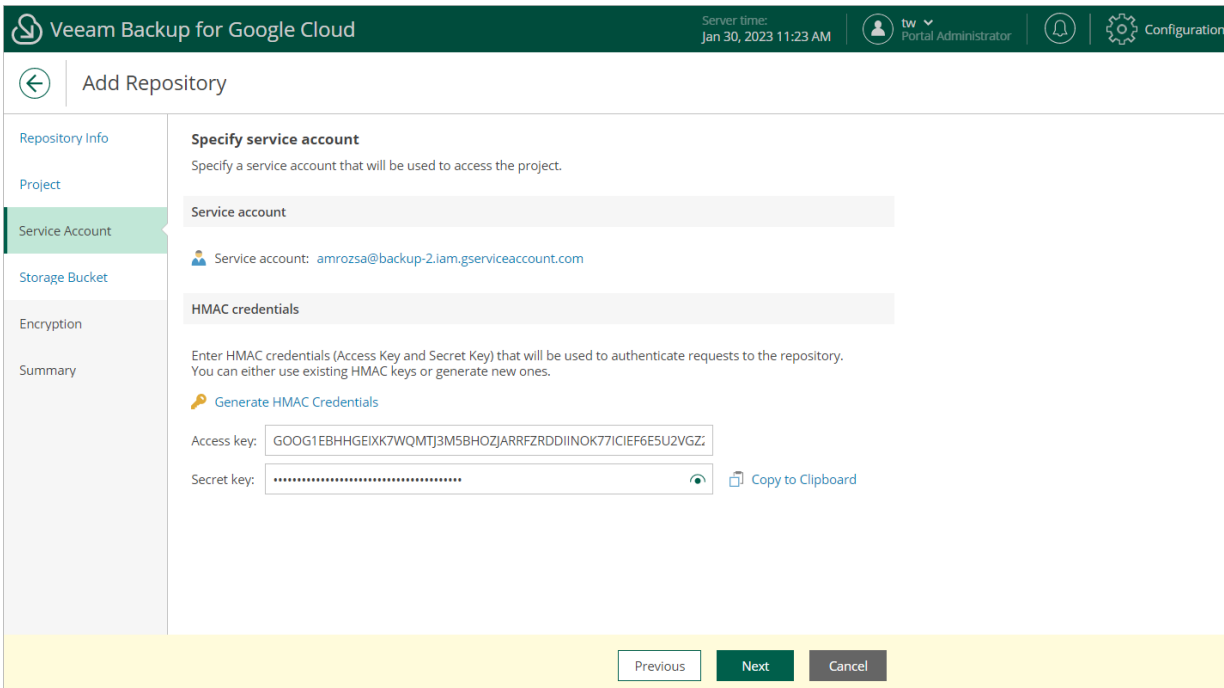
For a service account to be displayed in the **Service Accounts** list, it must be added to Veeam Backup for Google Cloud as described in section [Adding Service Accounts](#), and must be assigned permissions required to access the specified project as described in section [Adding Projects and Folders](#).

IMPORTANT

The selected service account must belong to the same project as that you have specified at [step 3](#) of the wizard.

2. In the **HMAC credentials** section, use the **Access key** and **Secret key** fields to provide a Hash-based Message Authentication Code (HMAC) key associated with the account — Veeam Backup for Google Cloud will use the HMAC key to authenticate requests to the backup repository.

You can create the necessary HMAC key beforehand in the Google Cloud console as described in [Google Cloud documentation](#). Alternatively, you can click **Generate HMAC Credentials** to create a new HMAC key and associate it with the service account without closing the **Add Repository** wizard.



Step 5. Configure Repository Settings

At the **Storage Bucket** step of the wizard, do the following:

1. In the **Storage bucket** section, click **Choose bucket**.

In the **Choose storage bucket** window, select a storage bucket that will be used as a target location for image-level backups of VM, Cloud SQL and Cloud Spanner instances, and click **Apply**.

For a storage bucket to be displayed in the **Available Buckets** list, it must be created for the selected project in the Google Cloud console as described in [Google Cloud documentation](#).

2. In the **Folder** section, choose whether you want to use an existing subdirectory inside the selected storage bucket or to create a new one to group backups stored in the bucket.

- To use an existing subdirectory, select the **Use existing folder** option and click **Choose folder**. In the **Choose folder** window, select the necessary subdirectory and click **Apply**.

For a subdirectory to be displayed in the **Available Folders** list, it must be previously created by a backup appliance in the selected storage bucket.

NOTE

If you select an existing subdirectory for storing backup files, consider the following:

- The created backup repository will have the storage class that has been specified when creating the subdirectory. You cannot change the storage class for the repository.
- If encryption at the repository level was enabled for the selected subdirectory, you must provide the password that was used to encrypt data at [step 6](#) of the wizard.
- If the selected subdirectory already contains backups created by the Veeam backup service, Veeam Backup for Google Cloud 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 Google Cloud applies retention settings saved in the backup metadata to the imported backups. However, if the selected subdirectory 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 Google Cloud will rewrite the saved retention settings and will apply to the imported backups new retention settings configured for that backup policy.

- To create a new subdirectory, select the **Create new folder** option and specify a name for the subdirectory. The maximum length of the name is 127 characters; the following characters are not supported: \ / " ' [] : | < > + = ; , ? * @ & _ .

3. [This step applies only if you have selected the **Create new folder** option] In the **Storage class** section, select a storage class for the backup repository – it can be either the Standard Storage, Nearline Storage or Archive Storage:

- To store backups in a high-performance, short-term storage that you plan to access frequently, select **Standard**.
- To store backups a high-durable, low-cost storage that you plan to access infrequently, select **Nearline**.
- To store backups in a cost-effective, long-term storage that you plan to access less than once a year, select **Archive**.

For the full description of Google Cloud storage classes, see [Google Cloud documentation](#).

IMPORTANT

If you select the **Archive** option, you must also enable backup archiving for any backup policy that will store backups in this repository. For more information, see [Performing VM Backup](#), [Performing SQL Backup](#) and [Performing Spanner Backup](#).

The screenshot shows the Veeam Backup for Google Cloud interface. At the top, there is a header bar with the Veeam logo, the text 'Veeam Backup for Google Cloud', the server time 'Jan 30, 2023 11:30 AM', and user information 'tw Portal Administrator'. Below the header, there is a navigation bar with a back arrow and the text 'Add Repository'. The main content area is divided into three panes:

- Specify repository settings:** This pane contains settings for the repository, including 'Storage bucket' (with a 'Choose bucket...' button), 'Folder' (with radio buttons for 'Use existing folder' and 'Create new folder'), and 'Storage class' (with a 'Class: Select...' dropdown).
- Choose storage bucket:** This pane contains a search bar labeled 'Bucket' and a 'Rescan' button.
- Available Buckets:** This pane contains a table of available buckets and their regions. The 'abor-duo-us' bucket is highlighted in green.

Available Buckets ↑	Region
abaichurin-archive-repo	EUROPE-WEST3
abaichurin-bucket-encrypted	EUROPE-WEST3
abor-asiamulti	ASIA
abor-belgium	EU
abor-dual-eu	EUR4
abor-duo-us	NAM4
abor-netherlands	EUROPE-WEST4
abor-plus-eu	EUROPE-WEST1+EUROPE-WEST4
abor-us-multi	US
agtestbucket	EUROPE-WEST3
ak-eu-west1	EUROPE-WEST1

At the bottom of the dialog, there are two buttons: 'Apply' and 'Cancel'.

Step 6. Enable Encryption

At the **Encryption** step of the wizard, choose whether you want to encrypt backups stored in the selected storage bucket. If you enable encryption, specify a password that will be used to encrypt data.

If you have selected an existing subdirectory at the **Storage Bucket** step of the wizard, you must provide the currently used password to let Veeam Backup for Google Cloud access this subdirectory and add it as a backup repository. You cannot change the encryption settings while adding the repository, but you will be able to [edit the repository settings](#) later.

IMPORTANT

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 Repositories](#).

The screenshot shows the 'Add Repository' wizard in Veeam Backup for Google Cloud. The interface is divided into a left sidebar and a main content area. The sidebar contains a navigation menu with the following items: Repository Info, Project, Service Account, Storage Bucket, Encryption (highlighted in green), and Summary. The main content area is titled 'Specify encryption settings' and contains the following fields:

- Enable encryption
- Password: [password field with eye icon]
- Repeat password: [password field with eye icon]
- Password hint: [text field containing 'standard']

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

Step 7. Finish Working with Wizard

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

As soon as you click **Finish**, Veeam Backup for Google Cloud will start creating the new backup repository. To track the progress, click **Go to Sessions** in the **Session Info** window to proceed to the [Session Logs tab](#).

The screenshot shows the 'Add Repository' wizard in the 'Summary' step. The interface includes a top navigation bar with the Veeam logo, 'Veeam Backup for Google Cloud', server time (Jan 30, 2023 11:33 AM), user 'tw Portal Administrator', and a 'Configuration' gear icon. A left sidebar lists steps: Repository Info, Project, Service Account, Storage Bucket, Encryption, and Summary (highlighted). The main area is titled 'Review configured settings' and contains a 'Copy to Clipboard' icon. The settings are organized into sections: 'Repository' (Name: backup-repo-01, Description: primary backup repository), 'Project' (Service account: amroza@backup-2.iam.gserviceaccount.com, Project: Backup 2), 'Storage settings' (Bucket: am-bckt, Type: Region, Region: EUROPE-WEST1, Folder: am-backup, Storage class: Standard), and 'Encryption' (Status: Enabled). At the bottom, there are three buttons: 'Previous', 'Finish' (with a mouse cursor), and 'Cancel'.

Editing Backup Repositories

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 Google Cloud Web UI.

Editing Backup Repository Settings Using Console

For each standard backup repository, you can modify settings configured while adding the repository to the backup infrastructure:

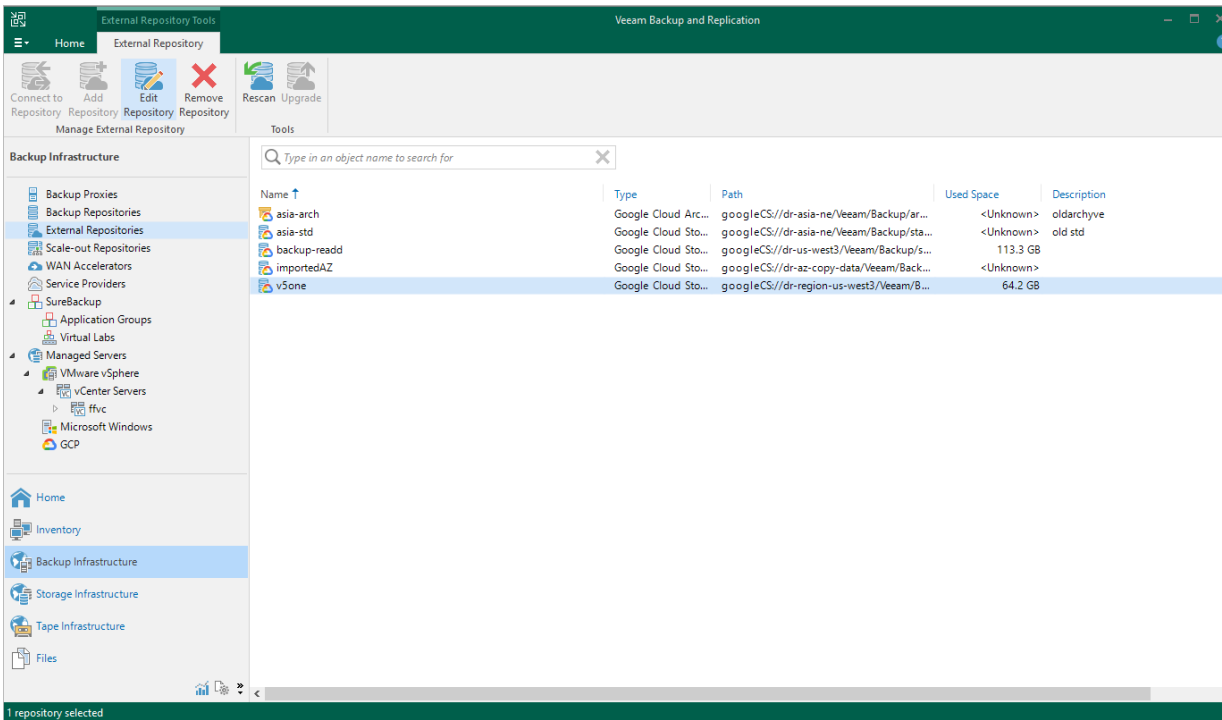
1. In the Veeam Backup & Replication console, open the **Backup Infrastructure** view.
2. Navigate to **External Repositories**.
3. Select the necessary repository and click **Edit Repository** on the ribbon.
Alternatively, you can right-click the repository and select **Properties**.
4. Complete the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

IMPORTANT

If you change the encryption settings of a standard backup repository using 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 section [Editing Backup Repository Settings Using Veeam Backup for Google Cloud Web UI](#).

- d. At the **Apply** step of the wizard, wait for the changes to be applied and click **Next**.

- e. At the **Summary** step of the wizard, review summary information and click **Finish** to confirm the changes.



Editing Backup Repository Settings Using Web UI

For each backup repository, you can modify settings configured while adding the repository to Veeam Backup for Google Cloud:

1. Switch to the **Configuration** page.
2. Navigate to **Repositories**.
3. Select the repository and click **Edit**.
4. Complete the **Edit Repository** wizard:
 - a. To provide a new name and description for the repository, follow the instructions provided in section [Adding Backup Repositories](#) (step 2).
 - b. To change the HMAC key used to authenticate requests to the backup repository, follow the instructions provided in section [Adding Backup Repositories](#) (step 4).
 - c. To enable encryption or change the encryption settings for the repository, follow the instructions provided in section [Adding Backup Repositories](#) (step 6).
 - d. 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 Google Cloud will start modifying the backup repository settings. To track the progress, click **Go to Sessions** in the **Session Info** window to proceed to the [Session Logs tab](#).

The screenshot shows the 'Edit Repository' configuration page for 'backup-repo-01'. The interface includes a top navigation bar with the Veeam logo, server time (Jan 30, 2023 11:38 AM), and user information (tw Portal Administrator). A left sidebar contains navigation tabs: Repository Info, Project, Service Account, Storage Bucket, Encryption, and Summary (which is currently selected). The main content area is titled 'Review configured settings' and contains a 'Copy to Clipboard' button. Below this, the settings are organized into sections: 'Repository' (Name: backup-repo-01, Description: primary backup repository), 'Project' (Service account: amrozs@backup-2.iam.gserviceaccount.com, Project: Backup 2), 'Storage settings' (Bucket: am-bckt, Type: Region, Region: EUROPE-WEST1, Folder: amroz, Storage class: Nearline), and 'Encryption' (Status: Enabled). At the bottom of the page, there are three buttons: 'Previous', 'Finish' (highlighted with a mouse cursor), and 'Cancel'.

Rescanning Backup Repositories

Veeam Backup & Replication periodically rescans standard 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**.

NOTE

Veeam Backup & Replication does not rescan image-level backups of Cloud SQL instances and Cloud Spanner instances stored in repositories.



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 Google Cloud Web UI.

Removing Backup Repositories Using Console

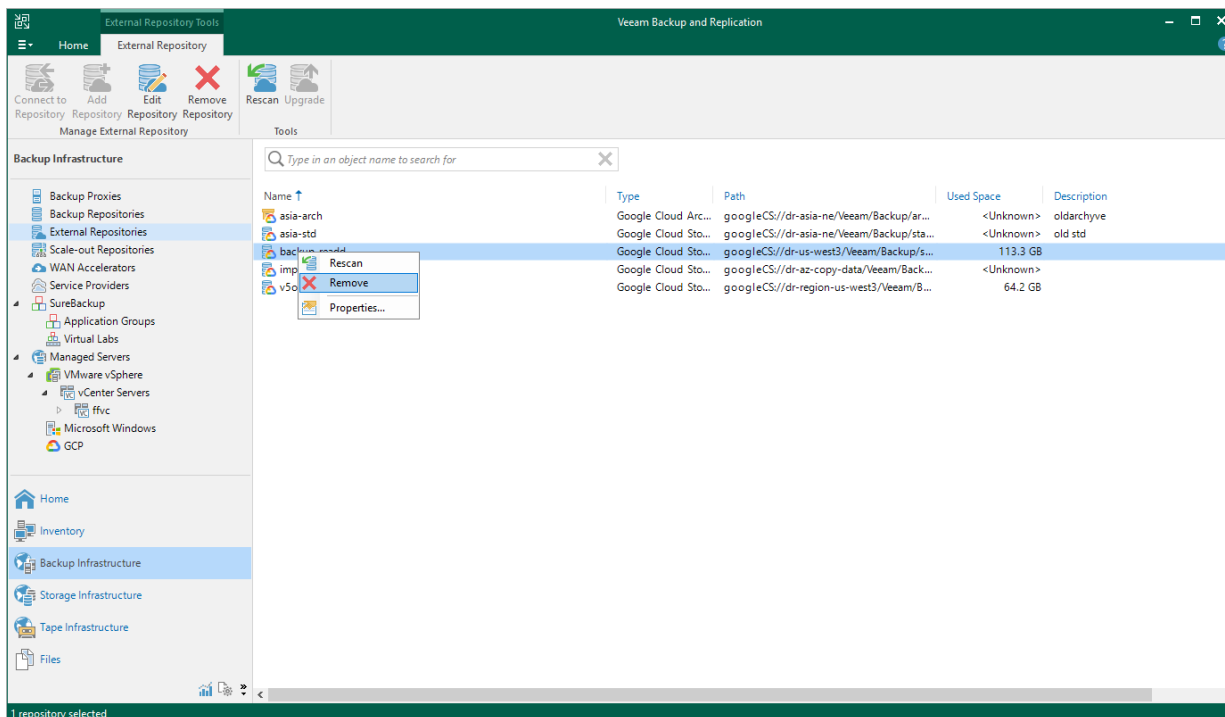
Google Cloud Plug-in for Veeam Backup & Replication allows you to permanently remove repositories from the backup infrastructure:

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**.

IMPORTANT

If you remove a backup repository using the Veeam Backup & Replication console, Veeam Backup & Replication will not propagate these settings to the backup appliance automatically. Consider removing the repository manually as described in section [Removing Backup Repositories Using Veeam Backup for Google Cloud Web UI](#).



Removing Backup Repositories Using Web UI

The Veeam Backup for Google Cloud Web UI allows you to permanently remove backup repositories if you no longer need them. When you remove a backup repository, Veeam Backup for Google Cloud unassigns the

repository role from the target storage bucket subdirectory so that the subdirectory is no longer used as a repository.

NOTE

Even though the storage bucket subdirectory is no longer used as a repository, Veeam Backup for Google Cloud preserves all backups previously stored in the repository and keeps these backups in Google Cloud Storage. You can assign the subdirectory to a new backup repository so that Veeam Backup for Google Cloud 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, you can remove it as described in section [Removing Backups and Snapshots](#).

To remove a backup repository, do the following:

1. Switch to the **Configuration** page.
2. Navigate to **Repositories**.
3. Select the repository and click **Remove**.

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, [change the configuration backup schedule](#) — and then try removing the repository again.

Server time: Nov 20, 2023 9:35 PM administrator Portal Administrator Configuration

Exit Configuration

Configure backup repositories to connect to your Google Cloud storage buckets and to store backup data produced by Veeam Backup for Google Cloud.

Repository

Repository	Storage Bucket	Folder	Project	Region	Storage Class	Encryption
<input type="checkbox"/> arch	tv-g-regress-v5-copy	arch-repo	veeam-rnd-backup...	US-EAST1+US-WEST1	Archive	Enabled
<input type="checkbox"/> st	tv-g-regress-v5-copy	st-repo	veeam-rnd-backup...	US-EAST1+US-WEST1	Standard	Enabled
<input checked="" type="checkbox"/> st-for-vm-b	tv-g-regress-v5-copy	st-for-vm-b-bug	veeam-rnd-backup...	US-EAST1+US-WEST1	Standard	Enabled
<input type="checkbox"/> st-rescanned	tv-g-prj2	st-v5-2	veeam-rnd-backup...	NAM4	Standard	Enabled

Managing Service Accounts

For each data protection and disaster recovery operation performed for a Google Cloud resource, you must specify a service account that has access to the resource and is assigned a set of permissions required to perform the operation.

Particularly, Veeam Backup for Google Cloud uses service accounts to perform the following tasks:

- To access projects and folders that manage Google Cloud resources.
- To synchronize the Google Cloud environment data with the data stored in the configuration database of the backup appliance.
- To create and remove snapshots of VM instances.
- To create and remove snapshots of Cloud SQL instances.
- To create and remove snapshots of Cloud Spanner instances.

During the product installation, the project in which the backup appliance is being deployed is automatically added to the configuration database, and the default service account is created in this project. The account can be further assigned permissions to perform operations within the initial project or any other project (or folder) added to Veeam Backup for Google Cloud. You can also create new and add existing Google Cloud service accounts to use them to access resources for data protection and disaster recovery tasks.

After you create or add a service account, you must grant the account the necessary permissions required to perform operations in a specific project (or folder), as described in section [Managing Projects and Folders](#).

Adding Service Accounts

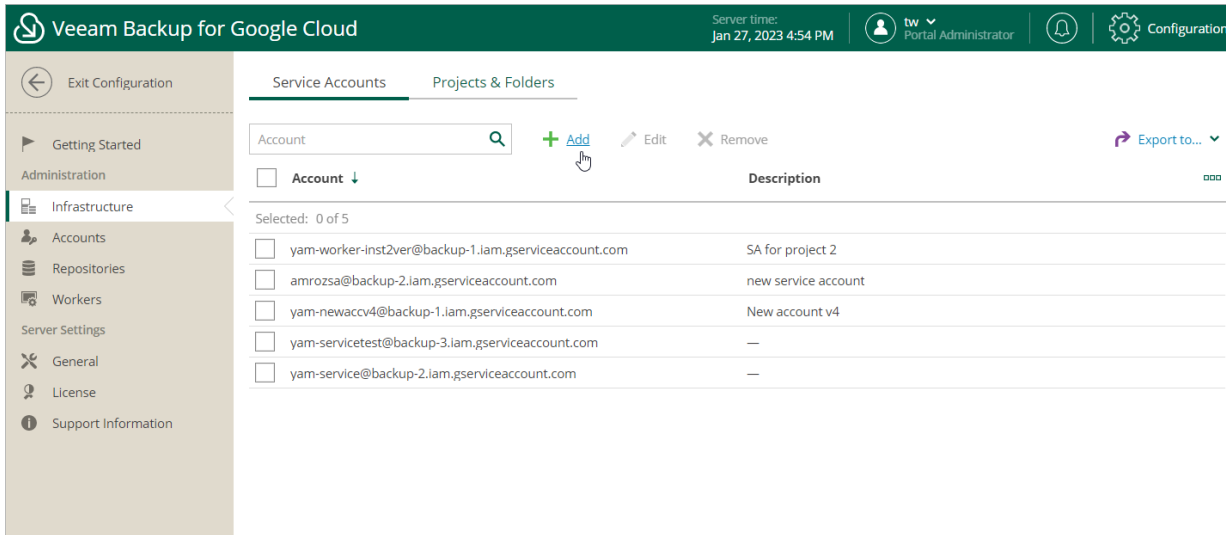
To add a new service account, do the following:

1. [Launch the Add Service Account wizard.](#)
2. [Choose a service account type.](#)
3. [Specify a project for the service account.](#)
4. [Specify service account details.](#)
5. [Track the account creation progress.](#)
6. [Finish working with the wizard.](#)

Step 1. Launch Add Service Account Wizard

To launch the **Add Service Account** wizard, do the following:

1. Switch to the **Configuration** page.
2. Navigate to **Infrastructure > Service Accounts**.
3. Click **Add**.



Step 2. Choose Service Account Type

At the **Service Account** step of the wizard, choose whether you want to add an already existing service account, or to create a new service account and add it to Veeam Backup for Google Cloud.

The screenshot shows the 'Add Service Account' step in the Veeam Backup for Google Cloud wizard. The interface includes a dark green header with the Veeam logo, the text 'Veeam Backup for Google Cloud', and system information: 'Server time: Jan 27, 2023 5:05 PM', a user profile for 'tw Portal Administrator', a notification bell, and a 'Configuration' gear icon. Below the header, a breadcrumb trail shows a back arrow and the text 'Add Service Account'. A left-hand navigation pane contains links for 'Service Account' (highlighted in green), 'Project', 'Account Details', 'Creation Log', and 'Summary'. The main content area is titled 'Specify service account' and contains the instruction 'Choose one of the following options:'. Two radio button options are listed: 'Add existing account' (unselected) and 'Create new account' (selected). At the bottom of the wizard, there are two buttons: 'Next' (highlighted in green) and 'Cancel'.

Step 3. Specify Project

[This step applies only if you have selected the **Create new account** option at the **Service Account** step of the wizard]

At the **Project** step of the wizard, specify the ID of a project in which the new service account will be created. You can find the project ID on the **Dashboard** page in the Google Cloud console. For more information, see [Google Cloud documentation](#).

TIP

If you want Veeam Backup for Google Cloud to automatically create a service account in the specified project, click **Sign in with Google** and specify credentials of a Google account that has [permissions required to create service accounts](#). For Veeam Backup for Google Cloud to be able to authorize in Google Cloud, the OAuth consent screen must be configured as described in section [Registering Applications](#).

Note that Veeam Backup for Google Cloud does not store in the configuration database the Google account credentials provided or access tokens received during authorization.

The screenshot shows the 'Add Service Account' wizard in Veeam Backup for Google Cloud. The interface has a dark green header with the product name and user information. A left sidebar contains navigation options: Service Account, Project (selected), Account Details, Creation Log, and Summary. The main content area is titled 'Specify project ID' and includes instructions, a text input field for the Project ID (containing 'backup-2'), a 'Sign in with Google' button, and a 'Previous' button at the bottom.

Veeam Backup for Google Cloud

Server time: Jan 27, 2023 5:05 PM | tw Portal Administrator | Configuration

← Add Service Account

Service Account

Project

Account Details

Creation Log

Summary

Specify project ID

Enter the ID of the project where the account will be created, and enter credentials with the required permissions. Veeam Backup for Google Cloud will not store these credentials — they will be used only to create the new account.

Project ID:

backup-2

Click 'Sign in with Google' to be redirected to the authentication page.

Sign in with Google

Previous Next Cancel

Step 4. Specify Account Details

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

- If you have selected the **Add existing account** option at the **Service Account** step of the wizard, use the **Email** and **Description** fields to specify an email address generated for the service account upon the account creation and to provide a description for future reference.
- If you have selected the **Create new account** option at the **Service Account** step of the wizard, use the **Account ID** and **Description** fields to specify an ID for the new service account and to provide a description for future reference.

The minimum length of the account ID is 6 characters. The following characters are supported: lowercase Latin letters, numeric characters and hyphens.

NOTE

If you have not signed in to Google Cloud at [step 3](#) of the wizard, Veeam Backup for Google Cloud will try to use the [default service account](#) to create the new service account automatically. If the default service account is missing the necessary permissions required to create service accounts in the specified project, you can generate a gcloud script and run it in the Google Cloud console to create the account manually. To generate the script, click **Download Script**.

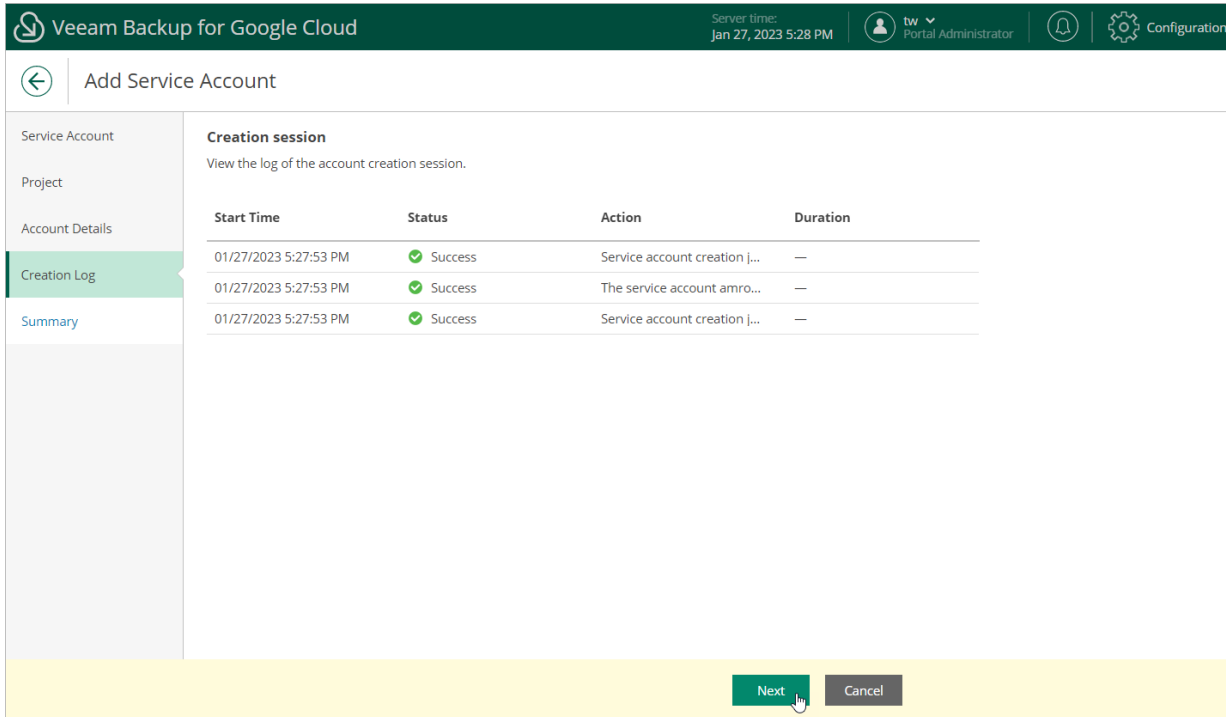
The account under which you run the script must have the permissions described in [Google Cloud documentation](#).

The screenshot shows the 'Add Service Account' wizard in Veeam Backup for Google Cloud. The interface is divided into a top header, a left sidebar, and a main content area. The top header displays the Veeam logo, the product name, the server time (Jan 27, 2023 5:11 PM), the user (tw Portal Administrator), and a Configuration icon. The left sidebar contains navigation links: Service Account, Project, Account Details (highlighted), Creation Log, and Summary. The main content area is titled 'Add Service Account' and contains the 'Enter account details' section. This section includes three input fields: 'Account ID' with the value 'amrozsa', 'Email address' with the value 'amrozsa@backup-2.iam.gserviceaccount.com', and 'Account description' with the value 'new service account'. Below the description field is a blue information icon and the text 'The description will be visible in the Veeam Backup UI only.' A green 'Download Script' button is located below the description field. At the bottom of the wizard, there is a yellow bar containing three buttons: 'Previous', 'Next', and 'Cancel'.

Step 5. Track Account Creation Progress

[This step applies only if you have selected the **Create new account** option at the **Service Account** step of the wizard]

Veeam Backup for Google Cloud will display the results of every step performed while creating the service account. At the **Creation Log** step of the wizard, wait for the creation process to complete and click **Next**.



The screenshot shows the 'Add Service Account' wizard in Veeam Backup for Google Cloud. The 'Creation Log' step is active, displaying a table of account creation sessions. The table has four columns: Start Time, Status, Action, and Duration. Three rows are visible, all showing a 'Success' status at 01/27/2023 5:27:53 PM. The 'Action' column contains truncated text: 'Service account creation j...', 'The service account amro...', and 'Service account creation j...'. The 'Duration' column shows dashes. At the bottom right, there are 'Next' and 'Cancel' buttons.

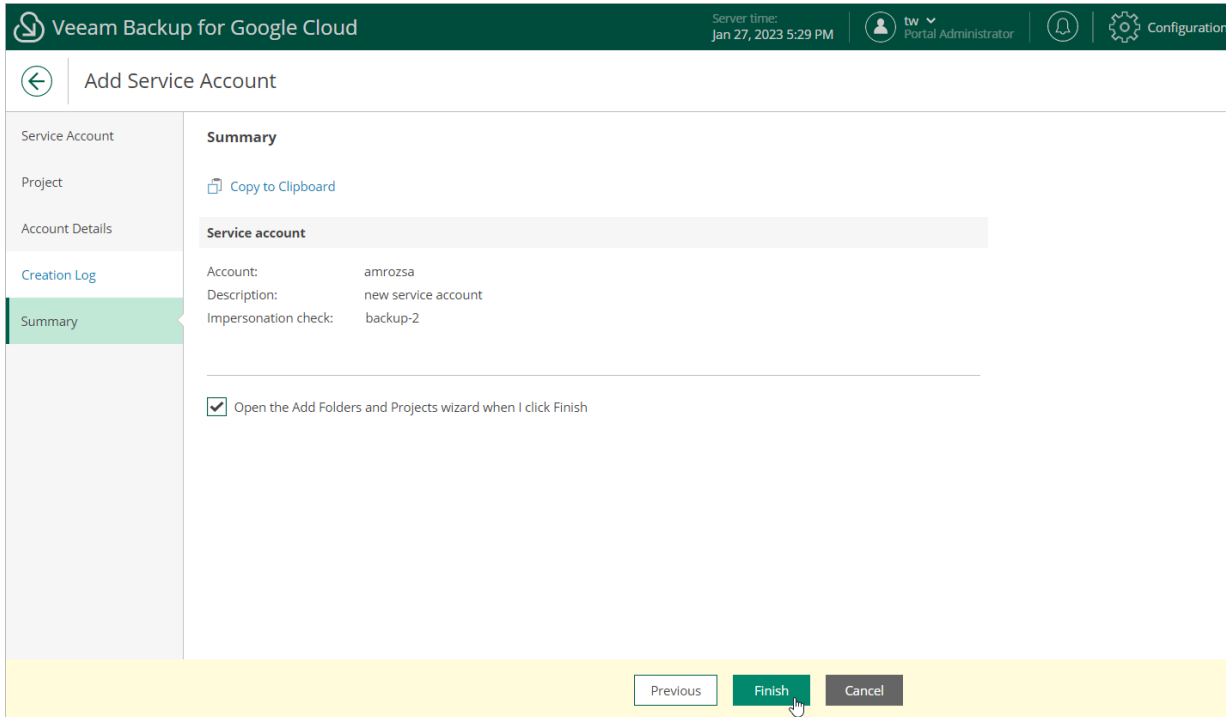
Start Time	Status	Action	Duration
01/27/2023 5:27:53 PM	Success	Service account creation j...	—
01/27/2023 5:27:53 PM	Success	The service account amro...	—
01/27/2023 5:27:53 PM	Success	Service account creation j...	—

Step 6. Finish Working with Wizard

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

TIP

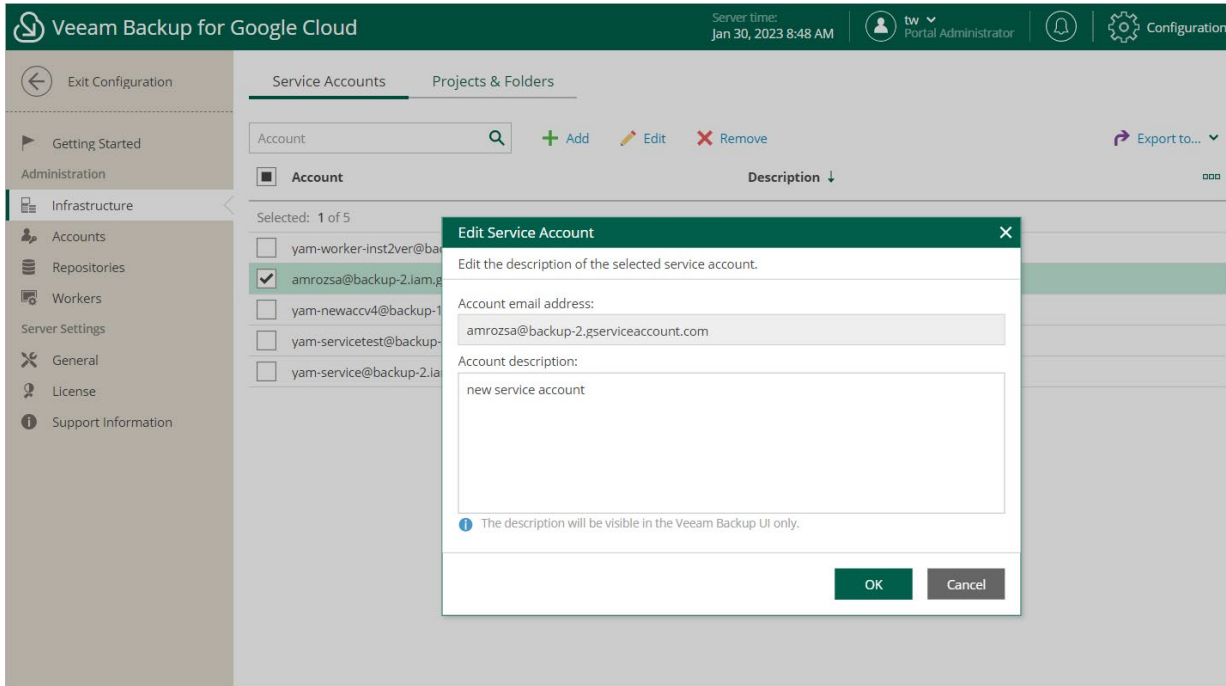
If you want to associate the newly added service account with a project or folder, select the **Open the Add Projects and Folders wizard when I click Finish** check box.



Editing Service Accounts

For each service account, you can only edit the description provided while adding the account:

1. Switch to the **Configuration** page.
2. Navigate to **Infrastructure > Service Accounts**.
3. Select the service account and click **Edit**.
4. In the **Edit Service Account** window, modify the description of the account and click **OK**.



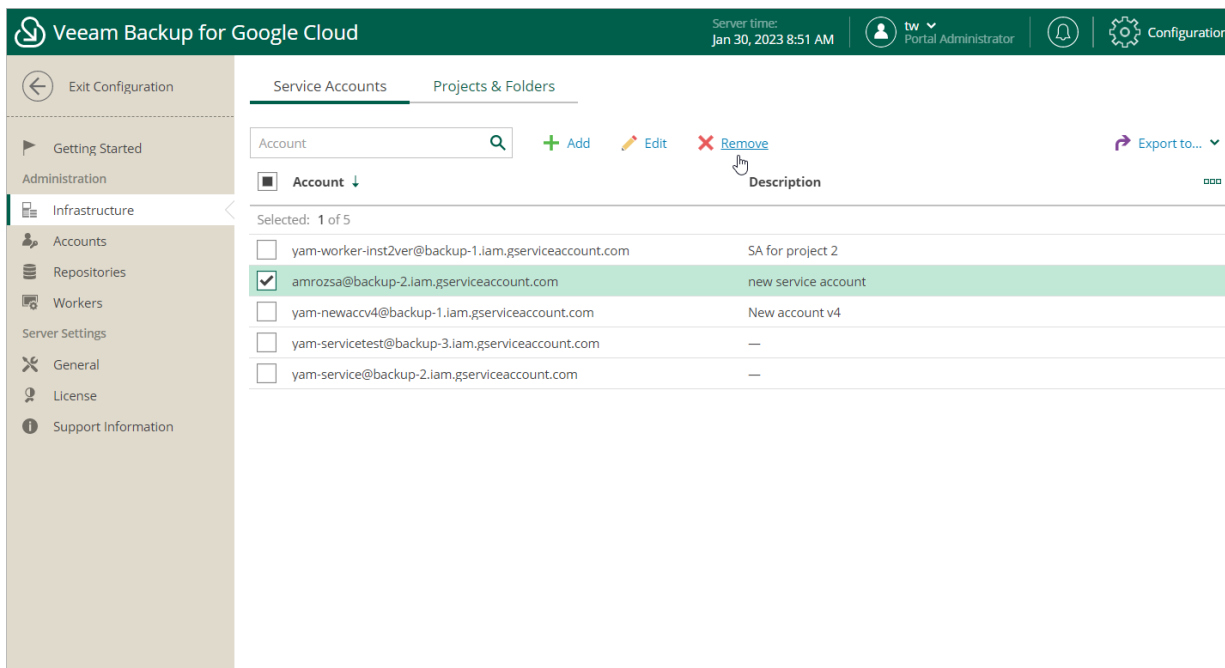
Removing Service Accounts

Veeam Backup for Google Cloud allows you to permanently remove a service account from the configuration database if you no longer need it:

1. Switch to the **Configuration** page.
2. Navigate to **Infrastructure > Service Accounts**.
3. Select the account and click **Remove**.

NOTE

You cannot remove a service account that is associated with any project or folder. [Remove all the related projects and folders](#) – and then try removing the account again.



Managing Projects and Folders

To be able to perform data protection and disaster recovery tasks for Google Cloud resources, you must first add to Veeam Backup for Google Cloud a project or folder that manages these resources, choose a service account that will be used to access this project or folder, and then specify a set of permissions that will be granted to the service account to perform the necessary operations.

TIP

When adding projects and folders, you can grant either a wide scope of permissions to one service account (to perform operations in different projects and folders) or granular scopes of permissions to different service accounts (to perform specific operations in one project or folder).

Adding Projects and Folders

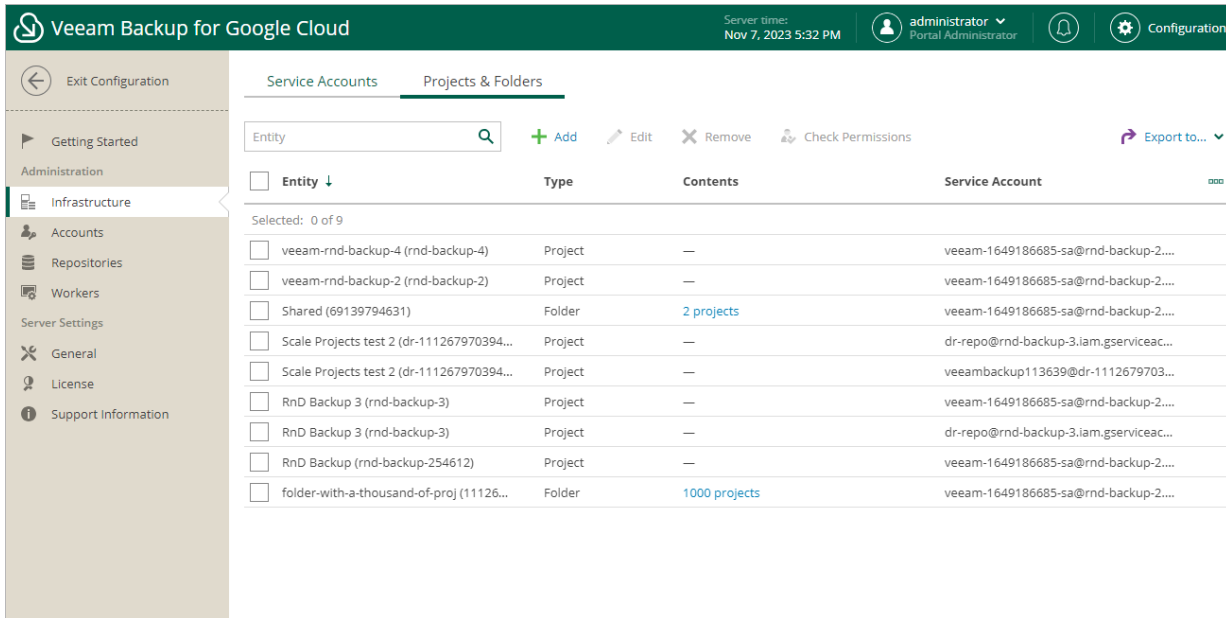
To add a new project or folder, do the following:

1. [Launch the Add Projects and Folders wizard.](#)
2. [Specify a service account to access the project or folder.](#)
3. [Define operations to perform in the project or folder.](#)
4. [Select the project or folder.](#)
5. [Check the required permissions.](#)
6. [Finish working with the wizard.](#)

Step 1. Launch Add Projects and Folders Wizard

To launch the **Add Projects and Folders** wizard, do the following:

1. Switch to the **Configuration** page.
2. Navigate to **Infrastructure > Projects & Folders**.
3. Click **Add**.



Step 2. Specify Service Account

At the **Service Account** step of the wizard, specify a service account that Veeam Backup for Google Cloud will use to access the project or folder, and choose whether you want to define operations that Veeam Backup for Google Cloud will be able to perform with resources managed by the project or folder.

NOTE

If you choose not to define the available operations, the specified service account will be assigned a wide scope of permissions required to perform all data protection operations in the selected project or folder.

To specify a service account, click the link in the **Service account** field. For a service account to be displayed in the list of available accounts, it must be added to Veeam Backup for Google Cloud as described in section [Adding Service Accounts](#). If you have not added the necessary service account to Veeam Backup for Google Cloud beforehand, you can do it without closing the **Add Projects and Folders** wizard. To add an account, click **Add** and complete the **Add Service Account** wizard.

The screenshot shows the Veeam Backup for Google Cloud interface. The main window is titled "Add Projects and Folders" and is in the "Specify service account" step. A modal window titled "Choose service account" is open, displaying a list of service accounts. The first account is highlighted.

Service Account	Description
dr-repo@rnd-backup-3.iam.gserviceaccount.com	—
dr-to-delete@rnd-backup-2.iam.gserviceaccount.com	—
veeam-1649186685-sa@rnd-backup-2.iam.gserviceac...	—
veeambackup113639@dr-111267970394-2.iam.gservi...	vbr account

Step 3. Define Operations

[This step applies only if you have selected the **Specify granular roles** check box at the **Service Account** step of the wizard]

At the **Roles** step of the wizard, define operations that Veeam Backup for Google Cloud will be able to perform for the resources managed by the project or folder: choose whether Veeam Backup for Google Cloud will be able to protect VM, Cloud SQL and Cloud Spanner instances that belong to this project or folder using cloud-native snapshots and image-level backups, to deploy backup repositories and workers in the project or folder, and to restore VM and Cloud SQL and Cloud Spanner instances to this project or folder from the created backups and snapshots.

In the **Veeam management roles** section, choose a type of the account role:

- **Repository access role** — permissions of this account role will be used to create new repositories in target Google Cloud buckets and further to access the repositories during data protection and disaster recovery operations. If you create an account role of this type, you will be able to select it [when configuring repository settings](#).
- **Worker deployment role** — permissions of this account role will be used to launch worker instances in the worker project. If you create a role of this type, you will be able to select it [when adding worker configurations](#).
- **File-level recovery to original location** — permissions of this account role will be used to launch worker instances during file-level restore operations. If you create a role of this type, you will be able to select it when performing file-level restore.

In the **Workload permissions** section, choose workloads that will be protected using permissions of the account role, and operations that will be performed with these workloads:

- If you select the **Backup** and **Snapshot** operations, you will be able to specify the service account when performing [VM backup](#), [SQL backup](#) and [Spanner backup](#).
- If you select the **Restore** operation, you will be able to specify the service account when performing [entire VM instance restore](#), [disk-level restore](#), [entire SQL instance restore](#), [SQL database restore](#), [entire Spanner instance restore](#) and [Spanner database restore](#).
- If you select the **File-level recovery to original location** operation, you will be able to specify the service account when performing [file-level recovery to the original location](#).

IMPORTANT

Keep in mind that the specified options apply only to the role selection for restore operations — they do not grant any permissions (unless you have selected the **Create new account** option at [step 2](#) of the **Adding Service Account** wizard). That is why it is recommended that you check whether the added service account has all the permissions required to perform operations with the selected workloads.

← Add Projects and Folders

Service Account

Roles

Project or Folder

Permissions

Summary

Specify granular roles and permissions

For the specific permissions required for data protection and disaster recovery tasks in Google Cloud, see [the User Guide](#).

Veeam management roles

- Repository access role
- ▼ Worker deployment role
 - File-level recovery to original location

Workload permissions

- ▼ VM instances
 - Snapshot
 - Backup
 - Restore
 - File-level recovery to original location
- ▼ Cloud SQL instances
 - Snapshot
 - Backup
 - Restore
- ▼ Cloud Spanner instances
 - Snapshot
 - Backup
 - Restore

Previous

Next

Cancel

Step 4. Specify Project or Folder

At the **Project or Folder** step of the wizard, specify the ID of a project or folder that manages the resources that you want to protect. If you choose a folder, Veeam Backup for Google Cloud will be able to access all resources in all projects that belong to this folder.

You can find the project and folder IDs on the **Dashboard** page in the Google Cloud console. For more information, see [Google Cloud documentation](#).

TIP

To help you choose a folder, Veeam Backup for Google Cloud provides information on the Google Cloud resource hierarchy in your organization. However, this option is available for authorized users only. To authorize in Google Cloud and to display the hierarchy, do the following:

1. Click **Sign in with Google**.

For Veeam Backup for Google Cloud to be able to authorize in Google Cloud, the OAuth consent screen must be configured as described in section [Registering Applications](#).

2. Specify credentials of a Google account with the Organization Viewer and Folder Viewer roles assigned.

Note that Veeam Backup for Google Cloud does not store in the configuration database the provided Google account credentials and access tokens received during authorization.

3. Click **Browse**.

The screenshot shows the 'Add Projects and Folders' step in the Veeam Backup for Google Cloud wizard. The interface includes a top navigation bar with the Veeam logo, the product name, server time (Nov 7, 2023 5:39 PM), and user information (administrator, Portal Administrator). A left sidebar contains navigation options: Service Account, Roles, Project or Folder (highlighted), Permissions, and Summary. The main content area is titled 'Specify target entity' and features a 'Sign in with Google' button. Below this, there are two radio button options: 'Folder name or ID' (with an 'Enter...' text box and a 'Browse...' button) and 'Project ID' (with a text box containing 'rnd-backup-2'). A note states: 'All subfolders and projects contained in the folder will be added automatically.' At the bottom of the wizard, there are three buttons: 'Previous', 'Next' (highlighted), and 'Cancel'.

Step 5. Check Required Permissions

At the **Permissions** step of the wizard, Veeam Backup for Google Cloud will verify whether the specified service account has all the permissions required to perform the defined operations in the selected project or folder. For more information on the required permissions, see [Service Account Permissions](#).

TIP

The service account specified at [step 2](#) of the wizard should not necessarily belong to the project or folder selected at [step 4](#). You can specify a service account created in another project or folder – and grant it the permissions required to access the entity.

To see the list of missing permissions that must be granted to the service account in order to perform an operation, click the link in the **Details** column. You can grant the missing permissions to the service account [using the Google Cloud console](#) or instruct Veeam Backup for Google Cloud to do it:

- To grant the missing permissions manually, click **Download Script**. Veeam Backup for Google Cloud will generate a gcloud script that you can run in the Google Cloud console to assign all the necessary permissions to the service account.

The account under which you run the script must have the permissions required both to get and set project IAM policies and to create custom IAM roles (for example, it can have the *iam.securityAdmin* and *iam.roleAdmin* roles assigned). To learn what permissions and roles are required to create custom roles in IAM, see [Google Cloud documentation](#).

- To let Veeam Backup for Google Cloud grant the missing permissions automatically, click **Grant** and then click **Sign in with Google** in the **Grant permissions** window. You will be redirected to the OAuth consent screen authorization page. Sign in using credentials of a Google account that will be used to grant the permissions.

The account under which you sign in to Google Cloud must have the permissions required both to get and set project IAM policies and to create custom IAM roles (for example, it can have the *iam.securityAdmin* and *iam.roleAdmin* roles assigned). To learn what permissions and roles are required to create service account, see [Google Cloud documentation](#).

NOTE

For Veeam Backup for Google Cloud to be able to authorize in Google Cloud, the OAuth consent screen must be configured as described in section [Registering Applications](#). Note that Veeam Backup for Google Cloud does not store in the configuration database the provided Google account credentials and access tokens received during authorization.

To make sure that the missing permissions have been successfully granted, click **Recheck**.

The screenshot shows the 'Check permissions' dialog in the Veeam Backup for Google Cloud interface. The dialog title is 'Add Projects and Folders'. The main content area is titled 'Check permissions' and contains the instruction: 'Verify whether all the required permissions are granted.' Below this instruction are three buttons: 'Recheck' (highlighted with a green checkmark icon), 'Download Script' (with a green download icon), and 'Grant' (with a blue plus icon). A table lists the following checks and their results:

Check	Result	Details
Appliance	Passed	All the required permissions are gra...
VM Backup	Passed	All the required permissions are gra...
VM Snapshot	Passed	All the required permissions are gra...
VM Restore	Passed	All the required permissions are gra...
Repository	Passed	All the required permissions are gra...
Worker	Passed	All the required permissions are gra...
Cloud SQL Snapshot	Passed	All the required permissions are gra...
Cloud SQL Backup	Passed	All the required permissions are gra...
Cloud SQL Restore	Passed	All the required permissions are gra...
Cloud SQL Staging Server	Passed	All the required permissions are gra...
VM File-Level Recovery to Original L...	Passed	All the required permissions are gra...
Cloud Spanner Snapshot	Passed	All the required permissions are gra...

At the bottom of the dialog are three buttons: 'Previous', 'Next' (highlighted in green), and 'Cancel'.

Step 6. Finish Working with Wizard

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

The screenshot shows the 'Add Projects and Folders' wizard in the Veeam Backup for Google Cloud interface. The 'Summary' step is selected in the left-hand navigation menu. The main content area displays 'Review configured settings' with a 'Copy to Clipboard' button. The settings are organized into several sections:

- General**
 - Project: veeam-rnd-backup-2 (rnd-backup-2)
 - Service account: veeam-1649186685-sa@rnd-backup-2.iam.gserviceaccount.com
- Veeam management roles**
 - Repository: Enabled
 - Worker: Enabled
- Workload permissions**
- VM instances**
 - Snapshot: Enabled
 - Backup: Enabled
 - Restore: Enabled
 - Restore (File-level recovery): Enabled
- Cloud SQL instances**
 - Snapshot: Enabled
 - Backup: Enabled
 - Restore: Enabled
- Cloud Spanner instances**
 - Snapshot: Enabled
 - Backup: Enabled
 - Restore: Enabled
- Validation:**
 - Permission checks: ✔ Passed

At the bottom of the wizard, there are three buttons: 'Previous' (disabled), 'Finish' (active), and 'Cancel' (disabled).

Editing Projects and Folders

For each project or folder, you can modify settings configured while adding the entity:

1. Switch to the **Configuration** page.
2. Navigate to **Infrastructure > Projects & Folders**.
3. Select the project or folder and click **Edit**.
4. Complete the **Edit Projects and Folders** wizard:
 - a. To modify the list of operations that Veeam Backup for Google Cloud can perform for the project or folder, follow the instructions provided in section [Adding Projects and Folders](#) (step 3).
 - b. To check and assign the required permissions to the selected service account, follow the instructions provided in section [Adding Projects and Folders](#) (step 5).

NOTE

The service account that is used to access the project to which the backup appliance belongs (that is, the project specified during the product installation) can only be changed in the Google Cloud console, as described in [Google Cloud documentation](#).

- c. At the **Summary** step of the wizard, review configuration information and click **Finish** to confirm the changes.

The screenshot shows the 'Edit Projects and Folders' wizard in the 'Summary' step. The interface includes a top navigation bar with the Veeam logo, 'Veeam Backup for Google Cloud', server time (Nov 18, 2023 6:11 PM), user profile (administrator), and a 'Configuration' button. The main content area is divided into a left sidebar with navigation options (Service Account, Roles, Permissions, Summary) and a main panel titled 'Review configured settings'. The main panel contains a 'Copy to Clipboard' button and several sections of configuration details: 'General' (Project: Scale Projects test 2, Service account: veeambackup113639@dr-111267970394-2.iam.gserviceaccount.com), 'Veeam management roles' (Repository: Enabled, Worker: Disabled), 'Workload permissions' (VM instances, Cloud SQL instances, Cloud Spanner instances), and 'Validation' (Permission checks: Passed). At the bottom, there are three buttons: 'Previous', 'Finish', and 'Cancel'.

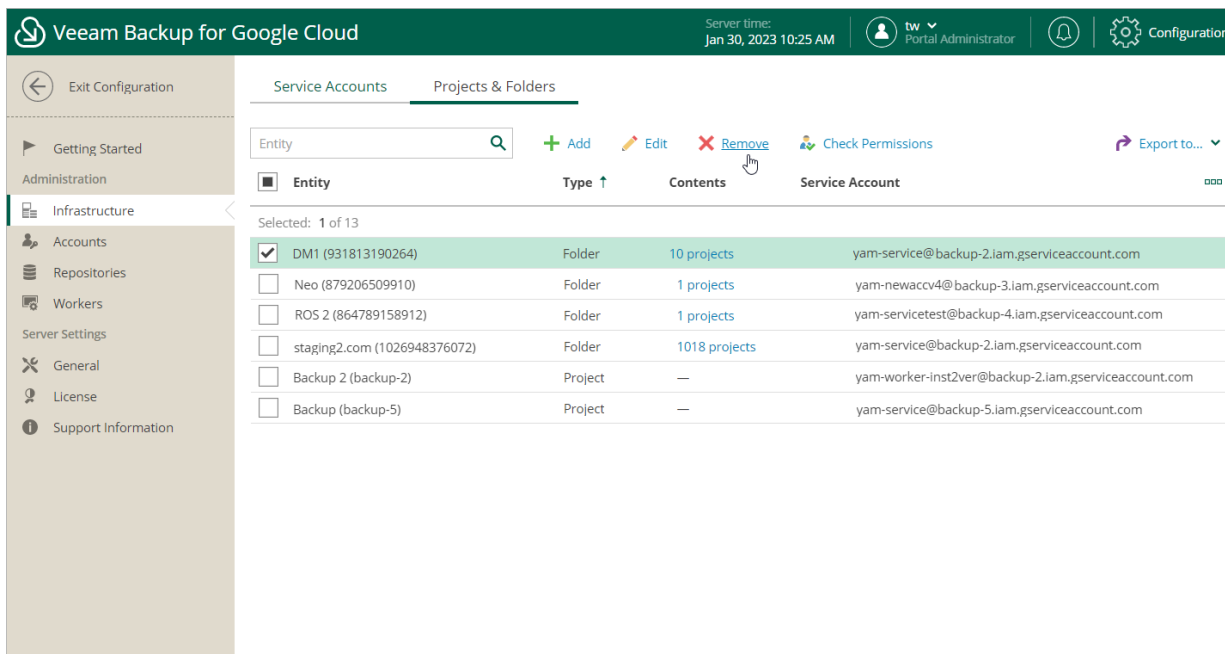
Removing Projects and Folders

Veeam Backup for Google Cloud allows you to permanently remove a project or folder from the configuration database if you no longer need it:

1. Switch to the **Configuration** page.
2. Navigate to **Infrastructure > Projects & Folders**.
3. Select the project or folder and click **Remove**.

NOTE

You cannot remove a project or folder that is used by any backup policy, backup repository or worker configuration. [Disable and remove all the related policies](#), [remove all the related repositories](#), [remove all the related worker configurations](#) — and then try removing the project again.



The screenshot displays the Veeam Backup for Google Cloud configuration interface. The top navigation bar includes the Veeam logo, server time (Jan 30, 2023 10:25 AM), and user information (tw Portal Administrator). The left sidebar shows the navigation menu with 'Infrastructure' selected. The main content area is titled 'Projects & Folders' and contains a table of entities. The table has columns for 'Entity', 'Type', 'Contents', and 'Service Account'. The first row, 'DM1 (931813190264)', is selected and highlighted in green. Above the table, there are buttons for '+ Add', 'Edit', 'Remove', and 'Check Permissions', along with an 'Export to...' dropdown. The table data is as follows:

Entity	Type	Contents	Service Account
<input checked="" type="checkbox"/> DM1 (931813190264)	Folder	10 projects	yam-service@backup-2.iam.gserviceaccount.com
<input type="checkbox"/> Neo (879206509910)	Folder	1 projects	yam-newaccv4@backup-3.iam.gserviceaccount.com
<input type="checkbox"/> ROS 2 (864789158912)	Folder	1 projects	yam-servicetest@backup-4.iam.gserviceaccount.com
<input type="checkbox"/> staging2.com (1026948376072)	Folder	1018 projects	yam-service@backup-2.iam.gserviceaccount.com
<input type="checkbox"/> Backup 2 (backup-2)	Project	—	yam-worker-inst2ver@backup-2.iam.gserviceaccount.com
<input type="checkbox"/> Backup (backup-5)	Project	—	yam-service@backup-5.iam.gserviceaccount.com

Managing User Accounts

Veeam Backup for Google Cloud 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 the Veeam Backup for Google Cloud UI.

There are 3 roles that you can assign to users working with Veeam Backup for Google Cloud:

- **Portal Administrator** – can perform all configuration actions, can manage user roles, and can also act as a Portal Operator and Restore Operator.
- **Portal Operator** – can create, edit and start 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.

The following table describes the functionality available to users with different roles in the Veeam Backup for Google Cloud 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	Execute
	File-level recovery	Full	Full	Execute
	Remove	Full	Full	N/A
Session Logs	Session logs	Full	Full	Read
	Stop session execution	Full	Full	N/A
Configuration				
Infrastructure	Service accounts, projects and folders	Full	N/A	N/A
Accounts	Portal users and SMTP accounts	Full	N/A	N/A

Tab	Functionality	Portal Administrator	Portal Operator	Restore Operator
Repositories	Backup repositories	Full	N/A	N/A
Workers	Worker instances	Full	N/A	N/A
General	General settings	Full	N/A	N/A
License	Licensing	Full	N/A	N/A
Support Information	Updates and logs	Full	N/A	N/A

Adding User Accounts

To add a new user account, do the following:

1. [Launch the Add Account wizard.](#)
2. [Specify an account name and description.](#)
3. [Specify a password.](#)
4. [Finish working with the wizard.](#)

Step 1. Launch Add Account Wizard

To launch the **Add Account** wizard, do the following:

1. Switch to the **Configuration** page.
2. Navigate to **Accounts > Portal Users**.
3. Click **Add**.

The screenshot shows the Veeam Backup for Google Cloud configuration interface. The top navigation bar includes the Veeam logo, the product name, the server time (Nov 17, 2023 6:54 PM), the user profile (administrator, Portal Administrator), and the Configuration page indicator. The left sidebar contains a navigation menu with categories like Administration, Infrastructure, Accounts, Repositories, Workers, and Server Settings. The main content area is titled 'Portal Users' and contains a table of existing portal users. The table has columns for Account, Role, Description, and MFA Enabled. The 'Add' button is visible at the top of the table.

Account	Role	Description	MFA Enabled
<input type="checkbox"/> wendy_may	Restore Operator	Dasha	No
<input type="checkbox"/> vitality	Portal Administrator		No
<input type="checkbox"/> vbr	Portal Administrator	integration vbr	No
<input type="checkbox"/> vasilisa	Portal Administrator		No
<input type="checkbox"/> sup	Portal Administrator		No
<input type="checkbox"/> sara_baker	Portal Administrator		No
<input type="checkbox"/> sam	Restore Operator	Sam Nicholls	No
<input type="checkbox"/> niels	Portal Administrator		No
<input type="checkbox"/> nastya	Portal Administrator		No
<input type="checkbox"/> lena	Portal Administrator		No
<input type="checkbox"/> dry	Portal Administrator		No
<input type="checkbox"/> aws	Portal Administrator		No
<input type="checkbox"/> alina	Portal Administrator		No
<input type="checkbox"/> alec	Portal Administrator	Alec The King	No

Step 2. Specify Account Name and Description

At the **Account Info** step of the wizard, use the **Name** and **Description** fields to enter a name for the new user account and to provide a description for future reference.

The maximum length of the account name is 32 characters. The following characters are supported: lowercase Latin letters, numeric characters, underscores and dashes. The following characters are not supported: \ / " ' [] : | < > + = ; , ? * @ & \$.

The screenshot shows the 'Add Account' wizard in Veeam Backup for Google Cloud. The interface includes a top navigation bar with the Veeam logo, server time (Nov 17, 2023 6:55 PM), and user information (administrator, Portal Administrator). The main content area is titled 'Add Account' and contains a sidebar with 'Account Info', 'General Settings', and 'Summary'. The 'Account Info' section is active, showing the 'Specify account name and description' step. Below this, there are two input fields: 'Name' with the value 'john_smith' and 'Description' with the value 'john_smith@veeam.com'. At the bottom of the form, there are 'Next' and 'Cancel' buttons.

Step 3. Specify Password

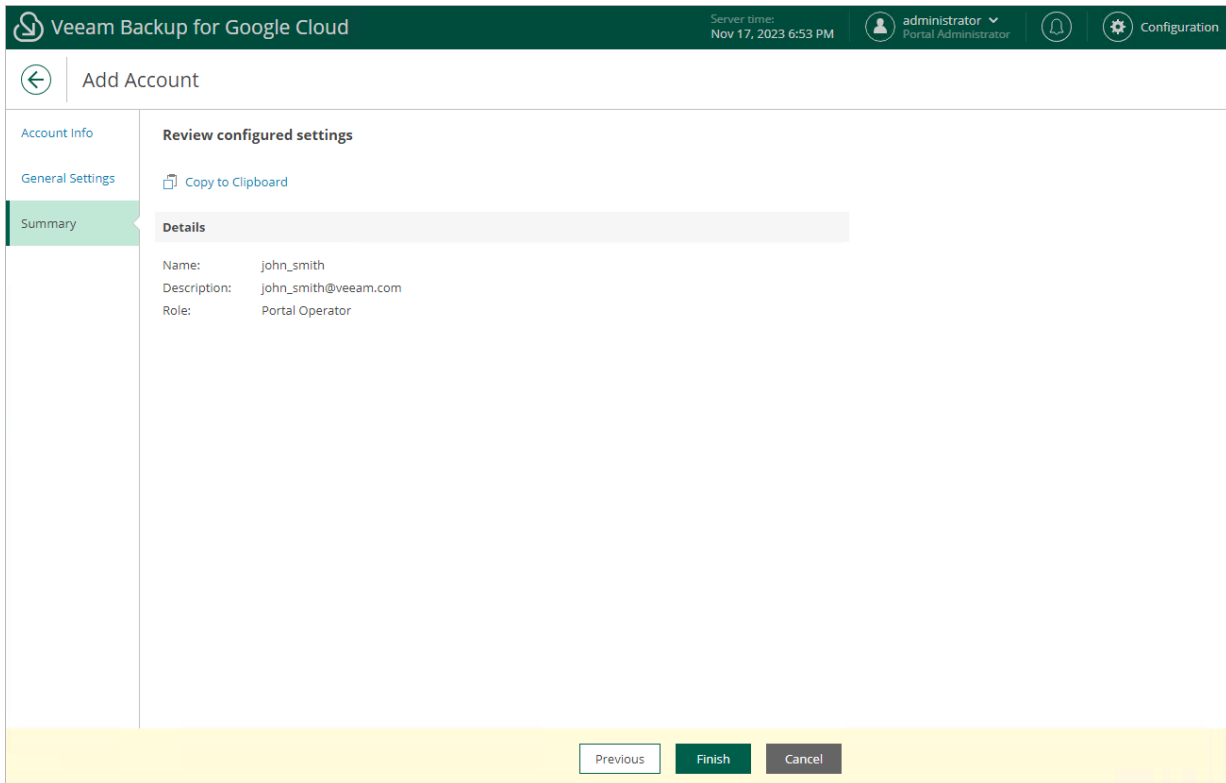
At the **General Settings** step of the wizard, choose a role for the user account and specify a password that the user will use to access Veeam Backup for Google Cloud.

The screenshot shows the 'Add Account' wizard in the Veeam Backup for Google Cloud interface. The top navigation bar includes the Veeam logo, the product name, the server time (Nov 17, 2023 6:55 PM), the user role (administrator / Portal Administrator), and a Configuration icon. The main content area is titled 'Add Account' and features a left-hand navigation pane with 'Account Info', 'General Settings' (selected), and 'Summary'. The 'Specify account settings' section is active, showing a 'Role' dropdown menu set to 'Portal Operator'. Below this is the 'Password' section, which includes two input fields: 'Password:' and 'Repeat password:', both masked with dots and having visibility toggles. A blue information icon and a text box provide password requirements: 'The password must be at least 8 characters long. It must contain at least 1 numeric character (0-9), 1 uppercase letter (A-Z) and 1 lowercase letter (a-z). Monotonic sequences (such as 1234) are not allowed.'

Previous Next Cancel

Step 4. Finish Working with Wizard

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



Editing User Accounts

For each user account, you can modify settings configured while adding the account:

1. Switch to the **Configuration** page.
2. Navigate to **Accounts > Portal Users**.
3. Select the account and click **Edit**.
4. Complete the **Edit Account** wizard:
 - a. To specify a new name and description for the account, follow the instructions provided in section [Adding User Accounts](#) (step 2).
 - b. To choose a new role for the account, follow the instructions provided in section [Adding User Accounts](#) (step 3).
 - c. At the **Summary** step of the wizard, review summary information and click **Finish** to confirm the changes.

The screenshot shows the 'Edit Account' wizard in the Veeam Backup for Google Cloud interface. The top navigation bar includes the Veeam logo, 'Veeam Backup for Google Cloud', server time 'Nov 17, 2023 6:56 PM', and user information 'administrator Portal Administrator'. The main content area is titled 'Edit Account' and features a sidebar with 'Account Info', 'General Settings', and 'Summary' (the active step). The 'Summary' step displays 'Review configured settings' with a 'Copy to Clipboard' button and a 'Details' section containing the following information:

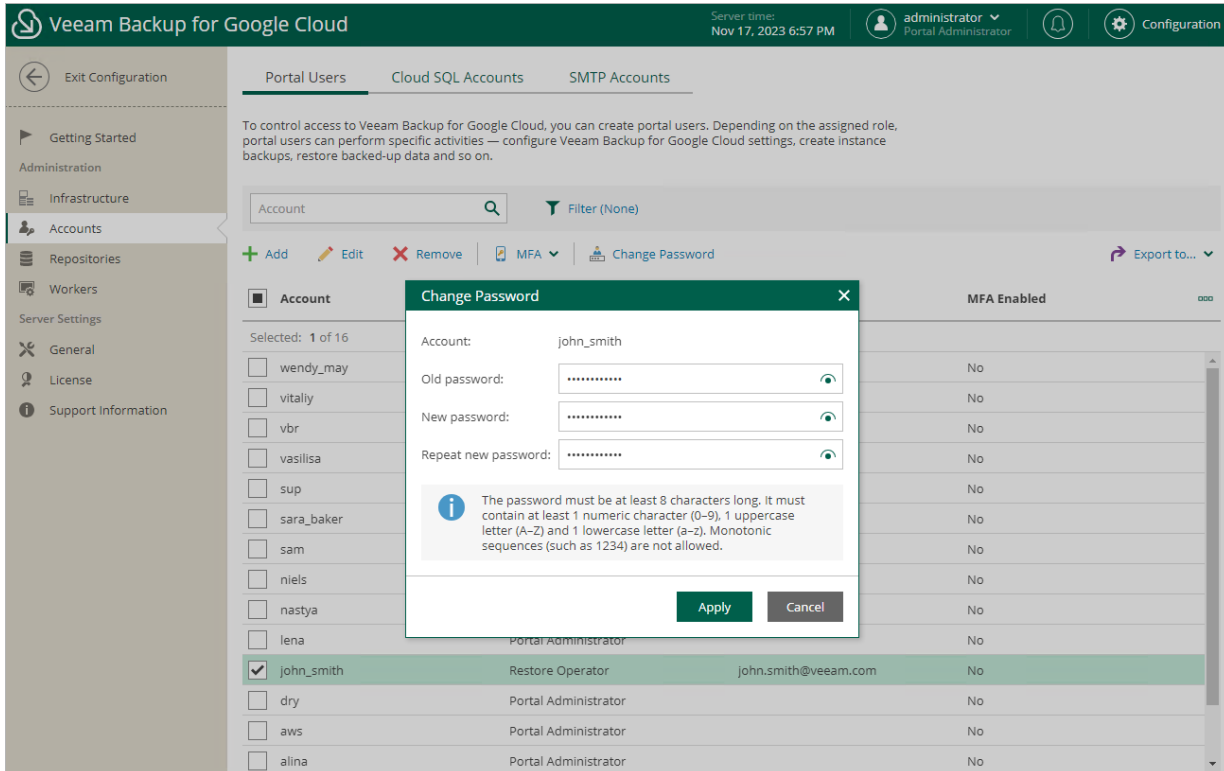
Name:	john_smith
Description:	john.smith@veeam.com
Role:	Restore Operator

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

Changing User Passwords

For each user account, you can change the password specified while adding the account:

1. Switch to the **Configuration** page.
2. Navigate to **Accounts > Portal Users**.
3. Select the 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 **Apply**.



Enabling Multi-Factor Authentication

Multi-factor authentication (MFA) in Veeam Backup for Google Cloud is based on the Time-based One-Time Password (TOTP) method that requires the user to verify their identity by providing a temporary six-digit code generated by an authentication application running on a trusted device.

To enable MFA for a user account, do the following:

1. Switch to the **Configuration** page.
2. Navigate to **Accounts > Portal Users**.
3. Select the account and click **MFA > Enable**.
4. Follow the instructions provided in the **Enabling MFA** window:
 - a. Install a supported authentication application on a trusted device. To view the list of authentication applications supported by Veeam Backup for Google Cloud, click **See the full list of compatible apps**.

NOTE

Only Google Authenticator is fully supported by Veeam Backup for Google Cloud.

- b. 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**.

The screenshot displays the Veeam Backup for Google Cloud configuration interface. The top navigation bar includes the Veeam logo, server name 'Veeam Backup for Google Cloud', server time 'Nov 17, 2023 6:58 PM', and user information 'administrator Portal Administrator'. The main content area is divided into three tabs: 'Portal Users', 'Cloud SQL Accounts', and 'SMTP Accounts'. The 'Portal Users' tab is active, showing a list of users with columns for 'Account' and 'Role'. The user 'john_smith' is selected, and the 'MFA' dropdown menu is open, showing 'Enable'. An 'Enabling MFA' dialog box is overlaid on the interface, containing the following steps:

1. Install a compatible app on your mobile device or computer. See the full list of compatible apps.
2. Use the app to scan the QR code below. A QR code is displayed in the center of the dialog.
- Alternatively, enter the following secret key manually: INROVINFKWWWYWSV3XJZI6WYMFJ
3. Enter the one-time password that currently appears in the app.

The dialog box has 'Apply' and 'Cancel' buttons at the bottom right. The background user list includes the following entries:

Account	Role
<input type="checkbox"/> wendy_may	Restore Operator
<input type="checkbox"/> vitality	Portal Administrator
<input type="checkbox"/> vbr	Portal Administrator
<input type="checkbox"/> vasilisa	Portal Administrator
<input type="checkbox"/> sup	Portal Administrator
<input type="checkbox"/> sara_baker	Portal Administrator
<input type="checkbox"/> sam	Restore Operator
<input type="checkbox"/> niels	Portal Administrator
<input type="checkbox"/> nastya	Portal Administrator
<input type="checkbox"/> lena	Portal Administrator
<input checked="" type="checkbox"/> john_smith	Restore Operator
<input type="checkbox"/> dry	Portal Administrator
<input type="checkbox"/> aws	Portal Administrator
<input type="checkbox"/> allina	Portal Administrator

Managing Cloud SQL Accounts

To allow Veeam Backup for Google Cloud to authenticate against Cloud SQL instances protected by backup policies, you must specify credentials that will be used to access the instances.

Out of the box, Veeam Backup for Google Cloud comes with the default IAM account. Credentials of this account allow Veeam Backup for Google Cloud to automatically detect unique email addresses associated with service accounts that are used to access Cloud SQL instances added to backup policies. However, you can create additional Cloud SQL accounts to granularly define credentials that will be used to access specific Cloud SQL instances.

IMPORTANT

To be able to use the default IAM credentials, you must configure Cloud SQL IAM database authentication for Cloud SQL instances in the Google Cloud console in advance, as described in [Google Cloud documentation](#). Note that Cloud IAM database authentication method is supported for MySQL instances only.

Adding Cloud SQL Accounts

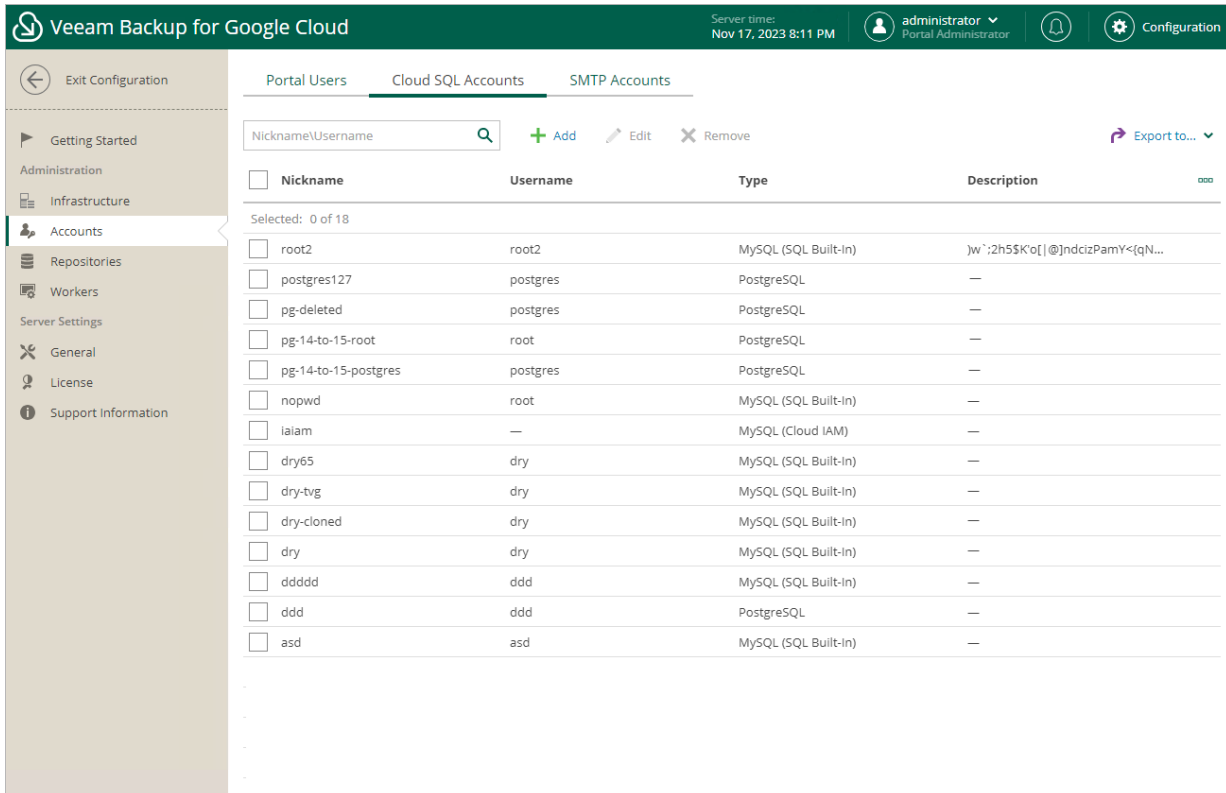
To add a new Cloud SQL account, do the following:

1. [Launch the Add Account wizard.](#)
2. [Specify an account name and description.](#)
3. [Specify general settings.](#)
4. [Finish working with the wizard.](#)

Step 1. Launch Add Account Wizard

To launch the **Add Account** wizard, do the following:

1. Switch to the **Configuration** page.
2. Navigate to **Accounts > Cloud SQL Accounts**.
3. Click **Add**.



Step 2. Specify Account Name and Description

At the **Account Info** step of the wizard, use the **Nickname** and **Description** fields to enter a name for the new Cloud SQL account and to provide a description for future reference.

The maximum length of the account nickname is 32 characters. The following characters are supported: lowercase Latin letters, numeric characters, underscores and dashes. The following characters are not supported: \ / " ' [] : | < > + = ; , ? * @ & \$.

The screenshot shows the 'Add Account' wizard in the Veeam Backup for Google Cloud interface. The top navigation bar includes the Veeam logo, the product name 'Veeam Backup for Google Cloud', the server time 'Nov 17, 2023 8:18 PM', and the user 'administrator Portal Administrator'. The main content area is titled 'Add Account' and has a left sidebar with 'Account Info' selected. The current step is 'Specify account nickname and description', with instructions to 'Enter a nickname and description for the Cloud SQL account.' There are two input fields: 'Nickname' containing 'root_postgres_account' and 'Description' containing 'postgresql account'. At the bottom, there are 'Next' and 'Cancel' buttons.

Step 3. Specify General Settings

At the **General Settings** step of the wizard, choose whether you plan to use this account in PostgreSQL or MySQL backup policies, and specify credentials that the account will use to access instances protected by these policies.

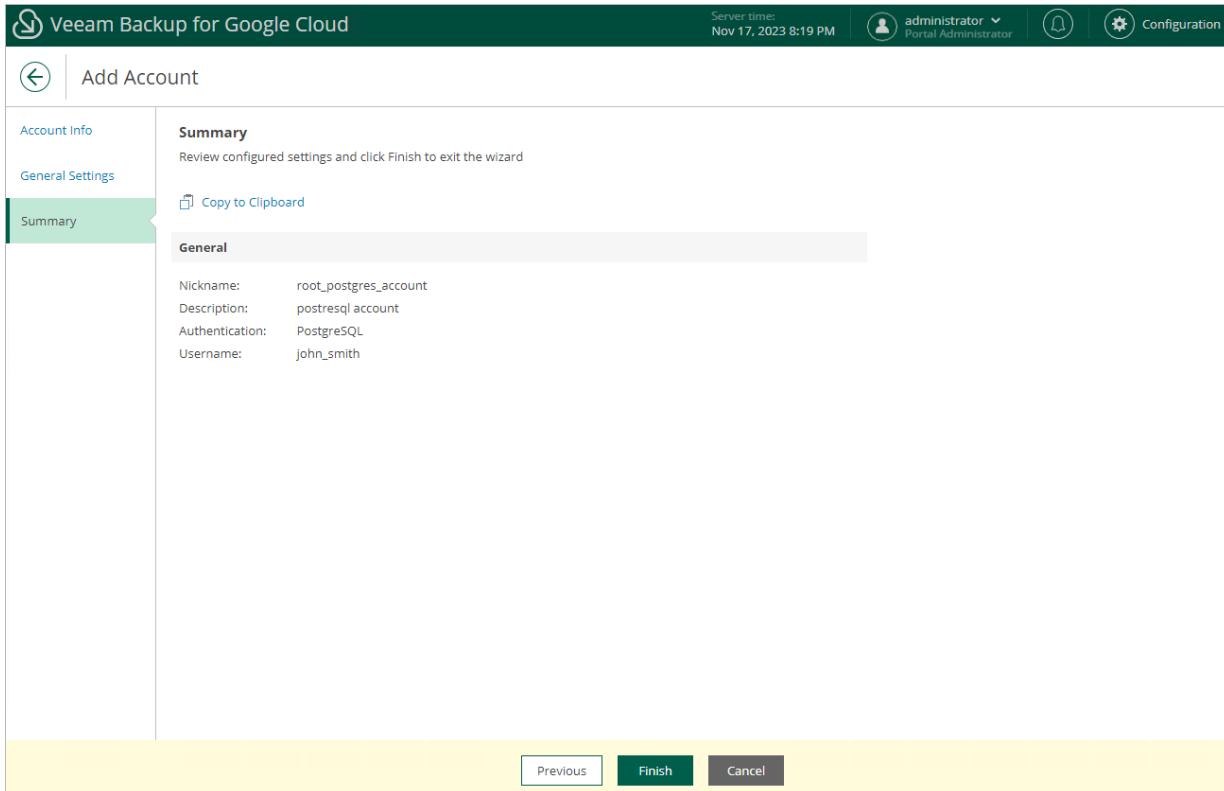
The screenshot shows the 'Add Account' wizard in the Veeam Backup for Google Cloud interface. The top navigation bar includes the Veeam logo, the product name 'Veeam Backup for Google Cloud', the server time 'Nov 17, 2023 8:18 PM', and the user 'administrator Portal Administrator'. The main content area is titled 'Add Account' and features a sidebar with three tabs: 'Account Info', 'General Settings' (which is active), and 'Summary'. The 'Specify account settings' section is active, with the instruction 'Enter credentials for the Cloud SQL account.' Below this, there are three input fields: 'Authentication' is a dropdown menu set to 'PostgreSQL'; 'Username' is a text input field containing 'john_smith'; and 'Password' is a masked text input field with a toggle icon to show or hide the password. At the bottom of the form, there are three buttons: 'Previous', 'Next' (highlighted in green), and 'Cancel'.

Step 4. Finish Working with Wizard

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

TIP

After you add the Cloud SQL account, you will be able to specify this account while creating backup policies to allow Veeam Backup for Google Cloud to access source Cloud SQL instances. For more information, see [Performing SQL Backup](#).



Editing Cloud SQL Accounts

For each Cloud SQL account, you can modify settings configured while adding the account:

1. Switch to the **Configuration** page.
2. Navigate to **Accounts > Cloud SQL Accounts**.
3. Select the account and click **Edit**.
4. Complete the **Edit Account** wizard:
 - a. To specify a new description and nickname for the account, follow the instructions provided in section [Adding Cloud SQL Accounts](#) (step 2).
 - b. To modify the credentials that are used to access Cloud SQL instances added to backup policies, follow the instructions provided in section [Adding Cloud SQL Accounts](#) (step 3).
 - c. At the **Summary** step of the wizard, review summary information and click **Finish** to confirm the changes.

The screenshot shows the 'Edit Account' wizard for a Cloud SQL account named 'root_postgres_account'. The interface is in the 'Summary' step, where the user is prompted to review the configured settings and click 'Finish' to exit the wizard. The 'General' section displays the following account details:

Nickname:	root_postgres_account
Description:	postgres account
Authentication:	PostgreSQL
Username:	wendy_may

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

Removing Cloud SQL Accounts

Veeam Backup for Google Cloud allows you to permanently remove a Cloud SQL account from the configuration database if you no longer need it:

1. Switch to the **Configuration** page.
2. Navigate to **Accounts > Cloud SQL Accounts**.
3. Select the account and click **Remove**.

NOTES

- You cannot remove the default *IAM Credentials* account.
- You cannot remove a Cloud SQL account that is associated with any backup policy. Delete all of the affected policies or [edit their settings](#) — and then try removing the account again.

The screenshot shows the Veeam Backup for Google Cloud configuration interface. The top navigation bar includes the Veeam logo, the text 'Veeam Backup for Google Cloud', the server time 'Nov 17, 2023 8:19 PM', and the user 'administrator Portal Administrator'. The left sidebar contains navigation options: Exit Configuration, Getting Started, Administration (Infrastructure), Accounts (selected), Repositories, Workers, Server Settings (General, License, Support Information). The main content area is titled 'Cloud SQL Accounts' and features a search bar, '+ Add', 'Edit', 'Remove', and 'Export to...' buttons. A table lists 19 accounts with columns for Nickname, Username, Type, and Description. The 'root_postgres_account' is selected, indicated by a green row and a checked checkbox.

<input type="checkbox"/>	Nickname	Username	Type	Description	...
<input type="checkbox"/>	root2	root2	MySQL (SQL Built-in));2h5\$K0[@JndcizPamY<{qN...	
<input checked="" type="checkbox"/>	root_postgres_account	john_smith	PostgreSQL	postgres account	
<input type="checkbox"/>	postgres127	postgres	PostgreSQL	—	
<input type="checkbox"/>	pg-deleted	postgres	PostgreSQL	—	
<input type="checkbox"/>	pg-14-to-15-root	root	PostgreSQL	—	
<input type="checkbox"/>	pg-14-to-15-postgres	postgres	PostgreSQL	—	
<input type="checkbox"/>	nopwd	root	MySQL (SQL Built-in)	—	
<input type="checkbox"/>	ialiam	—	MySQL (Cloud IAM)	—	
<input type="checkbox"/>	dry65	dry	MySQL (SQL Built-in)	—	
<input type="checkbox"/>	dry-tvg	dry	MySQL (SQL Built-in)	—	
<input type="checkbox"/>	dry-cloned	dry	MySQL (SQL Built-in)	—	
<input type="checkbox"/>	dry	dry	MySQL (SQL Built-in)	—	
<input type="checkbox"/>	dddd	ddd	MySQL (SQL Built-in)	—	
<input type="checkbox"/>	ddd	ddd	PostgreSQL	—	
<input type="checkbox"/>	asd	asd	MySQL (SQL Built-in)	—	

Managing Worker Instances

To perform most data protection and disaster recovery operations (such as creating image-level backups in backup repositories and restoring backed-up data), Veeam Backup for Google Cloud uses worker instances.

Each worker instance is deployed in a specific Google Cloud region for the duration of the backup or restore process. For more information on regions in which Veeam Backup for Google Cloud deploys worker instances, see [Architecture Overview](#).

Managing Worker Configurations

A configuration is a group of network settings that Veeam Backup for Google Cloud uses to deploy worker instances in a specific Google Cloud region to perform data protection and disaster recovery operations. Veeam Backup for Google Cloud deploys one worker instance per each VM, Cloud SQL or Cloud Spanner instance added to a backup policy or restore task.

By default, Veeam Backup for Google Cloud deploys worker instances with the same network configurations as those specified for the processed instances. However, to optimize infrastructure costs and to ensure better performance of backup and restore processes, you can add worker configurations to specify network settings for each region in which worker instances will be deployed.

NOTE

You can tell worker instances from other VM instances running in your environment by their names – the names of all worker instances deployed by Veeam Backup for Google Cloud will contain the word *worker*, a GUID and the name of the processed resource, and will be assigned the label *veeamvbid*.

Specifying Project for Worker Instances

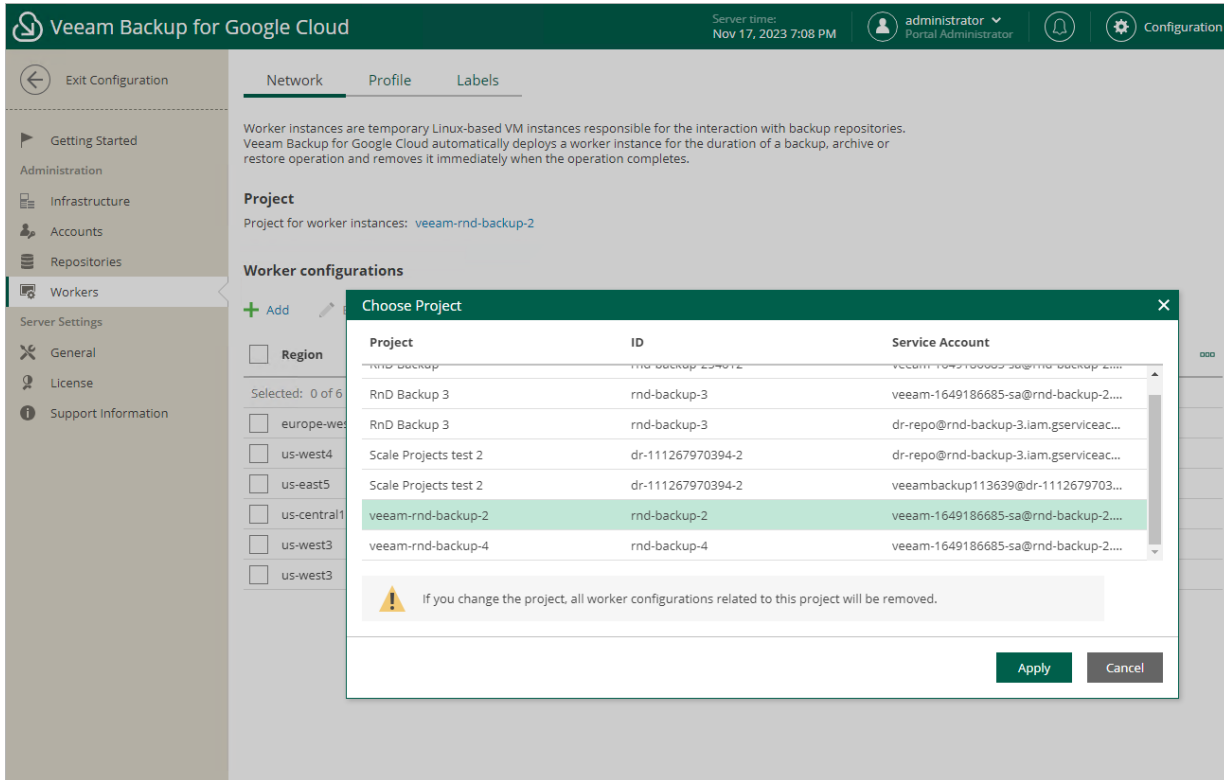
To specify a project in which worker instances will be created, do the following:

1. Switch to the **Configuration** page.
2. Navigate to **Workers > Network**.
3. Click the link in the **Project** section.
4. In the **Choose Project** window, select the project associated with a service account whose permissions will be used to deploy worker instances. Then, click **Apply**.

Note that Veeam Backup for Google Cloud does not automatically check whether the service account has all the permissions required to deploy worker instances. That is why you must select the project carefully.

IMPORTANT

It is recommended that you do not use a production project for worker instances. Production projects are not suitable for worker instances, since they could use too many network and storage resources when added to workloads in large environments, and thus could trigger the [Google Cloud quota limits](#).



Adding Worker Configurations

To add a new worker configuration, do the following:

1. [Launch the Add Worker Configuration wizard.](#)
2. [Specify general settings for the worker configuration.](#)
3. [Specify network settings for the worker configuration.](#)
4. [Check the required prerequisites.](#)
5. [Finish working with the wizard.](#)

Step 1. Launch Add Worker Configuration Wizard


To launch the Add Worker Configuration wizard, click Add in the Worker configurations section.

The screenshot shows the Veeam Backup for Google Cloud configuration interface. The top navigation bar includes the Veeam logo, the product name, the server time (Nov 17, 2023 7:08 PM), and the user role (administrator Portal Administrator). The left sidebar contains navigation options: Exit Configuration, Getting Started, Administration (Infrastructure, Accounts, Repositories), Workers (selected), and Server Settings (General, License, Support Information). The main content area is titled 'Worker configurations' and includes a description: 'Worker instances are temporary Linux-based VM instances responsible for the interaction with backup repositories. Veeam Backup for Google Cloud automatically deploys a worker instance for the duration of a backup, archive or restore operation and removes it immediately when the operation completes.' Below this, the 'Project' is identified as 'veeam-rnd-backup-2'. The 'Worker configurations' section has '+ Add', 'Edit', and 'Remove' actions. A table lists configurations with columns for Region, Availability Zone, Virtual Private Cloud, Subnet, and Firewall Rule. The table shows 6 rows of configurations, all with checkboxes in the left margin.

<input type="checkbox"/>	Region	Availability Zone	Virtual Private Cloud	Subnet	Firewall Rule	...
<input type="checkbox"/>	europa-west3	—	yam-worker-prj2	yam-sub-west3	Worker network con...	
<input type="checkbox"/>	us-west4	—	tvg-net	tvg-net	https	
<input type="checkbox"/>	us-east5	—	tvg-net	tvg-net	tvg-vb-v3-conf-b-vm...	
<input type="checkbox"/>	us-central1	us-central1-a	rnd-shared	rnd-usc1	dr-rule	
<input type="checkbox"/>	us-west3	us-west3-a	rnd-shared	rnd-eu-west3	dr-rule	
<input type="checkbox"/>	us-west3	us-west3-b	rnd-shared	rnd-eu-west3	dr-rule	

Step 2. Specify General Settings

At the **Region** step of the wizard, select a region where new worker instances will operate and an availability zone for which you want to configure network settings.



Step 3. Specify Network Settings

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

1. Select a VPC network and a subnet to which you want to connect worker instances created based on the new worker configuration.

For a VPC network and a subnet to be displayed in the lists of available networks, they must be created in the Google Cloud console for the region specified at [step 2](#) of the wizard, as described in [Google Cloud documentation](#).

IMPORTANT

- A route whose destination IP address range is 0 . 0 . 0 . 0 / 0 and whose next hop is the default internet gateway must exist for the selected VPC network. To learn how to add and remove routes for a network, see [Google Cloud documentation](#).
- The selected subnet must have Private Google Access enabled. To learn how to enable Private Google Access for a subnet, see [Google Cloud documentation](#).
- If you plan to back up Cloud SQL instances, you must configure network access between the subnets of the worker instances and the subnets of the processed Cloud SQL instances. Alternatively, you can configure the worker instances to allow public IP access as described in section [Configuring Deployment Mode](#).
- If you plan to back up Cloud SQL instances using a [staging server](#), the selected VPC network must have private services access configured. To learn how to configure private services access for a VPC network, see [Google Cloud documentation](#).
- If you want to connect worker instances created based on the worker configuration to a Shared VPC network, the [service account used to launch worker instances](#) must have the permissions described in [Worker Permissions](#).

2. Select a firewall rule that will be used to access worker instances deployed based on the configuration during file-level recovery operations.

For a firewall rule to be displayed in the list of available rules, it must be created in the Google Cloud console as described in [Google Cloud documentation](#).

IMPORTANT

- The selected firewall rule must allow direct network traffic to Google Cloud resources. Proxy redirect and setting a proxy in the Veeam Backup for Google Cloud configuration are not supported.
- If you plan to [perform file-level recovery](#), the selected firewall rule must allow HTTPS traffic to all VM instances on the specified VPC network. To learn how to create firewall rules that allow HTTPS connections, see [Google Cloud documentation](#).

Add Worker Configuration

Specify network settings
Configure network settings for worker instances that will run in

Region

Network

Verification

Summary

VPC: tvg-net

Subnet: tvg-net

Firewall rule: tvg-vb-v5-1677592743-https

Choose firewall rule

Rule 🔍 [Rescan](#)

Rule	ID
rule-5900	1755578712986310651
ssh	167163596410727104
tvg-big-vb-1699892463-https	6520257071449895452
tvg-vb-regress-v4-1674739330-https	356051613623776997
tvg-vb-regress-v4-1674739330-public-api	5824380315488993382
tvg-vb-regress-v4-1674739330-ssh	7862108436471726182
tvg-vb-regress-v5-1700128386-https	5971972455658969195
tvg-vb-temp-for-del-1697802732-https	7911887158940522270
tvg-vb-v3-conf-b-vm-1698929207-https	2755411193587705006
tvg-vb-v3-conf-b-vm-1698929207-public-api	4380214733461922990
tvg-vb-v3-conf-b-vm-1698929207-ssh	4256442116817691822
tvg-vb-v5-1677592743-https	2410165835274630725
tvg-vb-v5-1677592743-public-api	8834787085833713221
tvg-vb-v5-1677592743-ssh	4995604373575022149

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

Step 4. Check Required Prerequisites

At the **Verification** step of the wizard, Veeam Backup for Google Cloud will verify whether all the necessary prerequisites required to deploy worker instances based on the new worker configuration are met. For more information on the prerequisites, see [Specifying Network Settings](#).

The screenshot shows the 'Add Worker Configuration' wizard in the Veeam Backup for Google Cloud interface. The 'Verification' step is active, displaying a list of checks and their results. The interface includes a top navigation bar with the Veeam logo, server time (Nov 17, 2023 7:10 PM), and user information (administrator, Portal Administrator). A left sidebar contains navigation tabs for Region, Network, Verification, and Summary. The main content area shows the 'Run verification checks' section with a 'Recheck' button and a list of four items: Google Private Access is enabled (OK), Port 443 for HTTPS traffic is open (Warning: Firewall rule does not allow the HTTPS traffic), System-generated default route is enabled (OK), and Private services access is enabled (OK). At the bottom, there are 'Previous', 'Next', and 'Cancel' buttons.

Veeam Backup for Google Cloud

Server time: Nov 17, 2023 7:10 PM

administrator Portal Administrator

Configuration

← Add Worker Configuration

Region

Network

Verification

Summary

Run verification checks

Verify that permissions and configuration are correct.

🔄 Recheck

Google Private Access is enabled: ✓ OK

Port 443 for HTTPS traffic is open: ⚠ Warning: Firewall rule does not allow the HTTPS traffic.

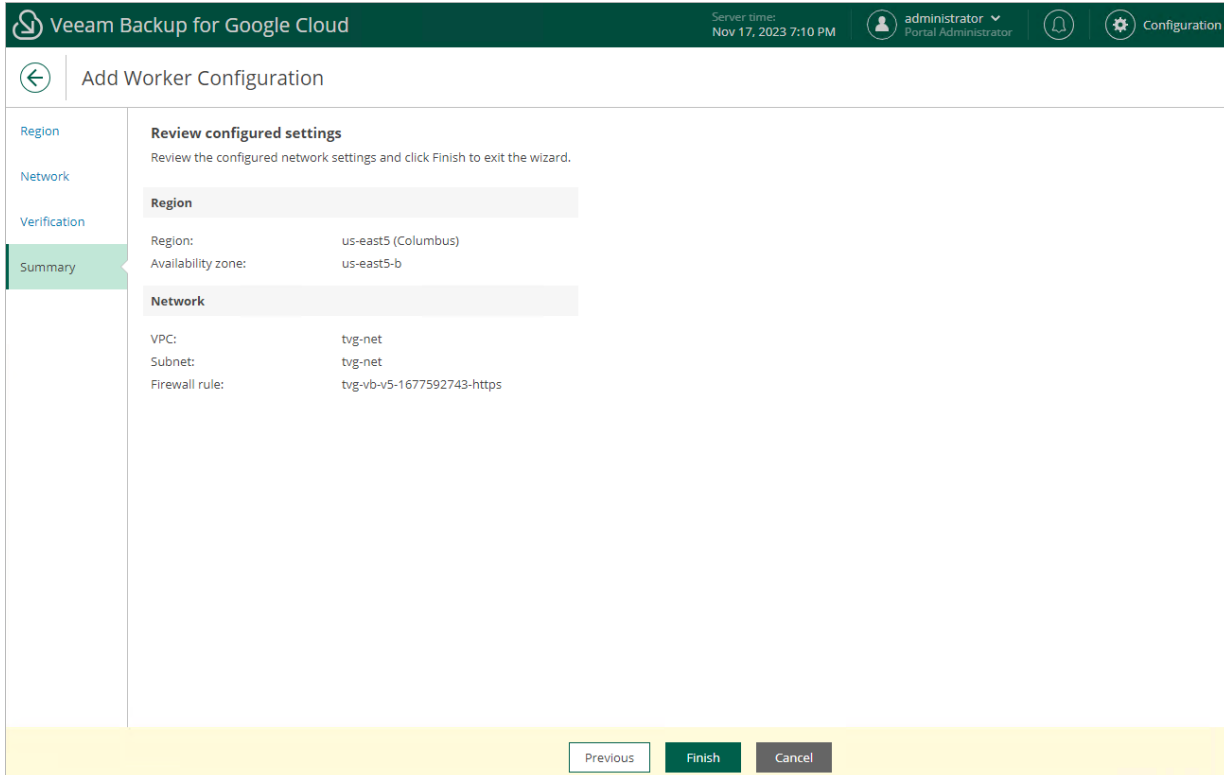
System-generated default route is enabled: ✓ OK

Private services access is enabled: ✓ OK

Previous Next Cancel

Step 5. Finish Working with Wizard

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



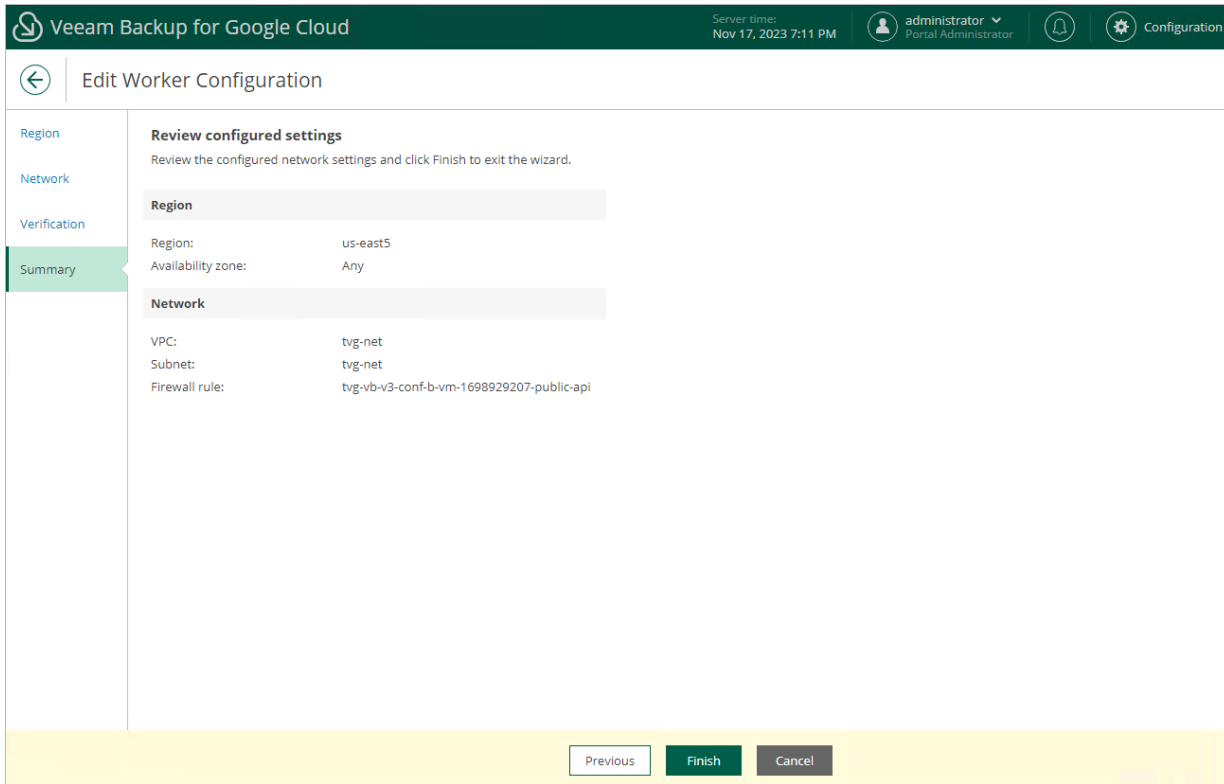
Editing Worker Configurations

For each worker configuration, you can modify settings specified while adding the worker configuration to Veeam Backup for Google Cloud:

1. Switch to the **Configuration** page.
2. Navigate to **Workers > Network**.
3. Select the worker configuration and click **Edit**.
4. Complete the **Edit Worker Configuration** wizard:
 - a. To modify the VPC network and subnet to which the related worker instances are connected, and to change the firewall rule associated with the specified network, follow the instructions provided in section [Adding Worker Configurations](#) (step 3).
 - 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 created based on the selected configuration that are currently involved in a backup or restore process, the changes will be applied only when the process completes.



Removing Worker Configurations

Veeam Backup for Google Cloud allows you to permanently remove worker configurations if you no longer need them. When you remove a worker configuration, Veeam Backup for Google Cloud 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 Google Cloud, do the following:

1. Switch to the **Configuration** page.
2. Navigate to **Workers > Network**.
3. 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 Google Cloud configuration interface. The top navigation bar includes the Veeam logo, the product name, the server time (Nov 17, 2023 7:11 PM), the user (administrator), and a Configuration button. The left sidebar contains navigation options: Exit Configuration, Getting Started, Administration (Infrastructure, Accounts, Repositories), Workers, and Server Settings (General, License, Support Information). The main content area is titled 'Network' and contains a description of worker instances, a 'Project' section (veeam-rnd-backup-2), and a 'Worker configurations' table. The table has columns for Region, Availability Zone, Virtual Private Cloud, Subnet, and Firewall Rule. One configuration for 'us-east5' is selected.

Worker instances are temporary Linux-based VM instances responsible for the interaction with backup repositories. Veeam Backup for Google Cloud automatically deploys a worker instance for the duration of a backup, archive or restore operation and removes it immediately when the operation completes.

Project
Project for worker instances: [veeam-rnd-backup-2](#)

Worker configurations
+ Add Edit Remove

Region	Availability Zone	Virtual Private Cloud	Subnet	Firewall Rule	
Selected: 1 of 6					
<input type="checkbox"/>	europe-west3	—	yam-worker-prj2	yam-sub-west3	Worker network con...
<input type="checkbox"/>	us-west4	—	tvg-net	tvg-net	https
<input checked="" type="checkbox"/>	us-east5	—	tvg-net	tvg-net	tvg-vb-v3-conf-b-vm-...
<input type="checkbox"/>	us-central1	us-central1-a	rnd-shared	rnd-usc1	dr-rule
<input type="checkbox"/>	us-west3	us-west3-a	rnd-shared	rnd-eu-west3	dr-rule
<input type="checkbox"/>	us-west3	us-west3-b	rnd-shared	rnd-eu-west3	dr-rule

Managing Worker Profiles

A profile is the machine type of a worker instance that Veeam Backup for Google Cloud deploys in a specific Google Cloud region to perform a backup or archive operation. Veeam Backup for Google Cloud deploys one worker instance per each Google Cloud resource (whether it is a VM instance, a Cloud SQL instance or a Cloud Spanner instance) added to a backup policy. The profile of each deployed worker instance is selected based on the regional quota.

There are 3 types of worker profiles in Veeam Backup for Google Cloud:

- **Primary** – a profile that Veeam Backup for Google Cloud uses for creating image-level backups if the regional disk quota has not been reached yet.
- **Secondary** – a profile that Veeam Backup for Google Cloud uses for creating image-level backups if you have run or about to run out of the regional disk quota.
- **Archiving** – a profile that Veeam Backup for Google Cloud uses for creating archived backups.

Out of the box, Veeam Backup for Google Cloud comes with the default set of worker profiles where the primary profile is *e2-highcpu-8*, the secondary profile is *e2-highcpu-2*, and the archiving profile is *e2-standard-4*. However, to boost operational performance and to guarantee that you do not breach Google Cloud quota limits, you can add custom sets of worker profiles to specify machine types of VM instances that will operate as worker instances in different regions.

IMPORTANT

Veeam Backup for Google Cloud does not allow you to change the default worker profiles that are used to launch worker instances performing restore, file-level recovery, health check and retention operations – the default machine types of these instances are listed in section [Architecture Overview](#). To customize the default worker profiles, open a [support case](#).

In This Section

- [Adding Worker Profiles](#)
- [Editing Worker Profiles](#)
- [Removing Worker Profiles](#)

Adding Worker Profiles

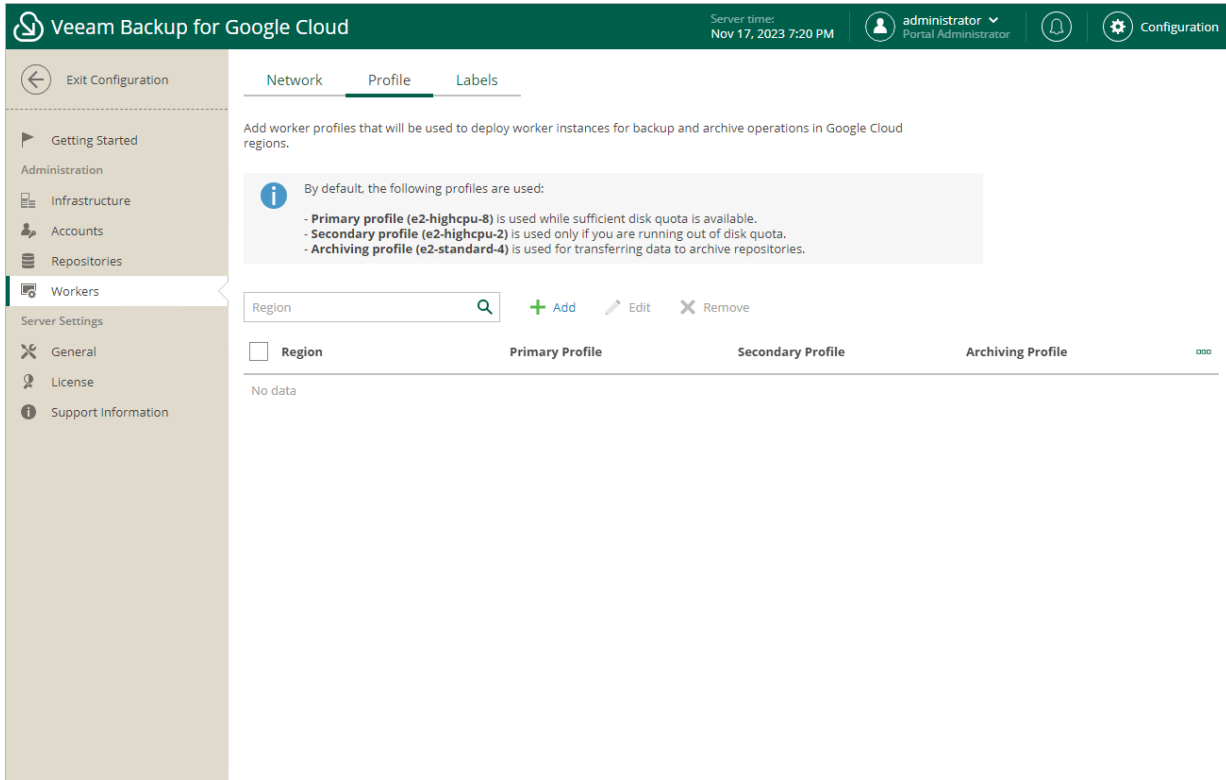
To add a new custom set of worker profiles for one or more regions, do the following:

1. [Launch the Add Worker Profiles wizard](#).
2. [Choose the necessary regions](#).
3. [Choose the default, secondary and archiving profiles for worker instances in these regions](#).
4. [Finish working with the wizard](#).

Step 1. Launch Add Worker Profiles Wizard

To launch the **Add Worker Profiles** wizard, do the following:

1. Switch to the **Configuration** page.
2. Navigate to **Workers > Profile**.
3. Click **Add**.



Step 2. Choose Regions

At the **Regions** step of the wizard, select regions for which you want to specify worker profiles.

The screenshot shows the 'Add Worker Profiles' wizard in the Veeam Backup for Google Cloud interface. The top navigation bar includes the Veeam logo, the product name, the server time (Nov 17, 2023 7:20 PM), the user role (administrator / Portal Administrator), and a Configuration icon. The main content area is titled 'Add Worker Profiles' and features a left-hand navigation menu with 'Regions' selected. The main panel is titled 'Choose one or more regions' and is divided into two sections: 'Available Regions (38)' and 'Selected Regions (1)'. The 'Available Regions' list includes various global locations such as europe-southwest1 (Madrid), europe-west1 (Belgium), europe-west2 (London), europe-west3 (Frankfurt), europe-west4 (Netherlands), europe-west6 (Zürich), europe-west8 (Milan), europe-west9 (Paris), europe-west12 (Turin), me-central1 (Doha), me-central2 (Dammam), me-west1 (Tel Aviv), northamerica-northeast1 (Montréal), northamerica-northeast2 (Toronto), southamerica-east1 (São Paulo), southamerica-west1 (Santiago), and us-central1 (Iowa). The 'Selected Regions' list currently contains 'europe-west10 (Berlin)'. Between the two lists are 'Add' and 'Remove' buttons. At the bottom of the wizard, there are 'Next' and 'Cancel' buttons.

Step 3. Choose Worker Profiles

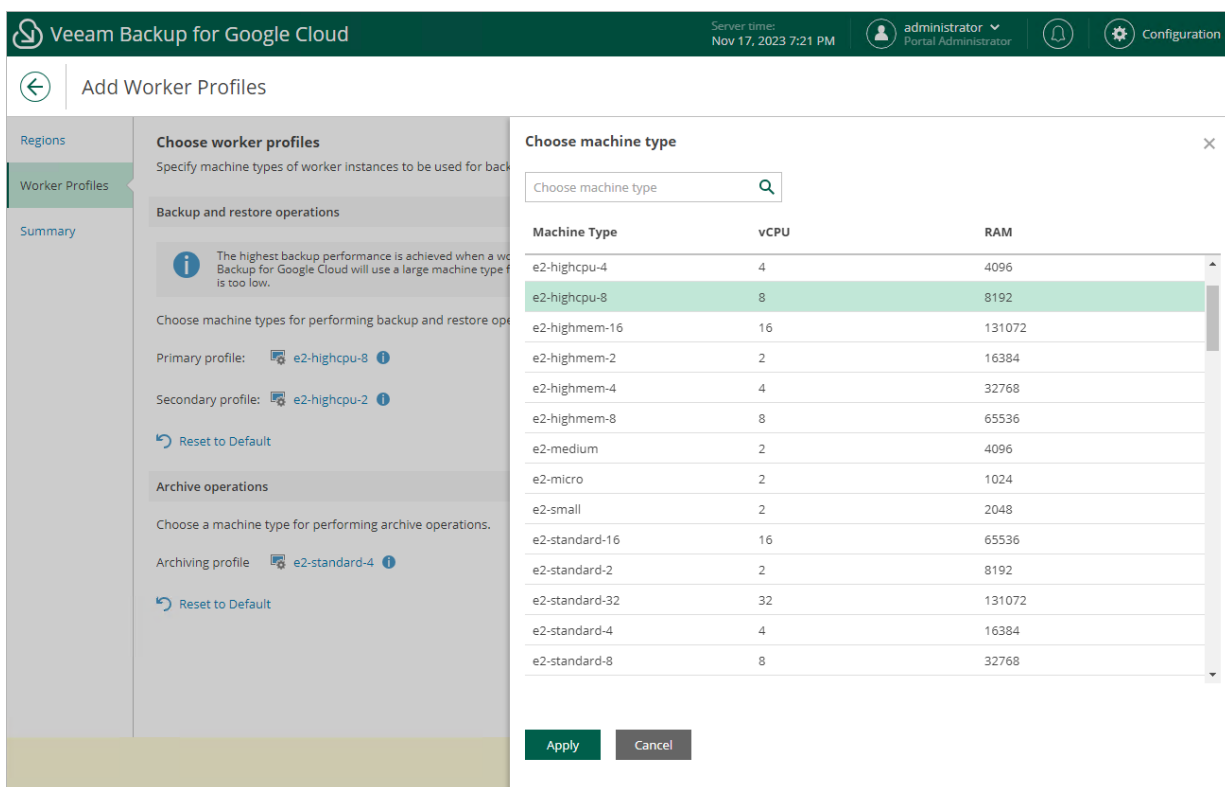
At the **Worker Profiles** step of the wizard, choose profiles that will be used to deploy workers in the selected regions. To help you choose, tables in the **Choose machine type** sections will provide information on the number of vCPU cores and the amount of system RAM for each available machine type.

IMPORTANT

Due to technical limitations, the list of available machine types is automatically filtered to show:

- For the primary profile, only those machine types that allow mounting persistent disks with at least 4 TB of total disk space attached.
- For the archiving profile, only those machine types that come with at least 8 GB RAM.

For the full description of machine types that can be used to deploy VM instances in Google Cloud, see [Google Cloud documentation](#).



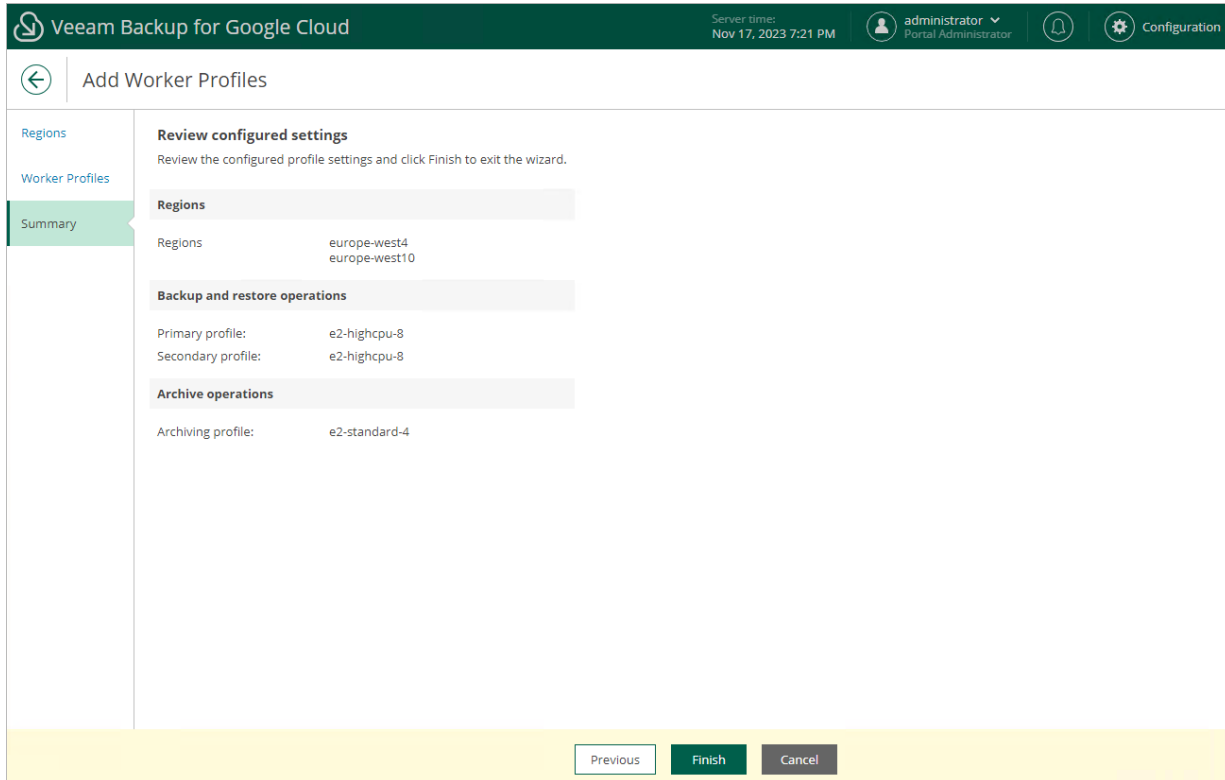
The screenshot shows the Veeam Backup for Google Cloud interface. The main window is titled "Add Worker Profiles" and has a sidebar with "Regions", "Worker Profiles", and "Summary". The "Worker Profiles" section is active, showing "Choose worker profiles" and "Choose machine type" options. The "Choose machine type" dialog is open, displaying a table of machine types. The table has columns for "Machine Type", "vCPU", and "RAM". The "e2-highcpu-8" machine type is selected, highlighted in green. Below the table are "Apply" and "Cancel" buttons.

Machine Type	vCPU	RAM
e2-highcpu-4	4	4096
e2-highcpu-8	8	8192
e2-highmem-16	16	131072
e2-highmem-2	2	16384
e2-highmem-4	4	32768
e2-highmem-8	8	65536
e2-medium	2	4096
e2-micro	2	1024
e2-small	2	2048
e2-standard-16	16	65536
e2-standard-2	2	8192
e2-standard-32	32	131072
e2-standard-4	4	16384
e2-standard-8	8	32768

Step 4. Finish Working with Wizard

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

As soon as you click **Finish**, Veeam Backup for Google Cloud will create a separate set of worker profiles for each of the selected regions.



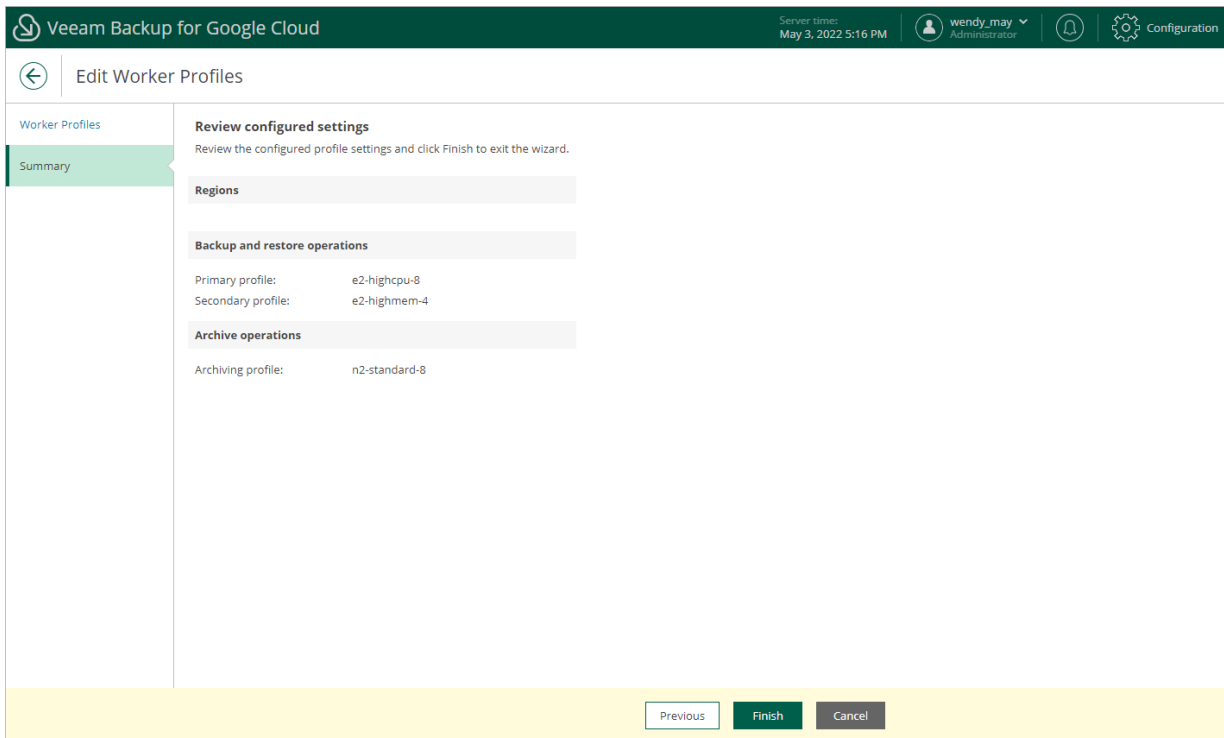
Editing Worker Profiles

For each set of worker profiles created for a Goggle Cloud region, you can modify settings specified while creating the profile set:

1. Switch to the **Configuration** page.
2. Navigate to **Workers > Profile**.
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 worker instances in the selected region, follow the instructions provided in section [Adding Worker Profiles](#) (step 3).
 - 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 process in the selected region, the changes will be applied only when the process completes.



Removing Worker Profiles

Veeam Backup for Google Cloud allows you to permanently remove sets of worker profiles if you no longer need them. When you remove a profile set, Veeam Backup for Google Cloud 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.

NOTE

After you remove a profile set, all worker instances that Veeam Backup for Google Cloud will further use to perform backup and archive operations in the region specified in the set will be deployed with the [default profiles](#).

To remove a profile set from Veeam Backup for Google Cloud, do the following:

1. Switch to the **Configuration** page.
2. Navigate to **Workers > Profile**.

3. Select the profile set and click **Remove**.

The screenshot shows the Veeam Backup for Google Cloud configuration interface. The top navigation bar includes the Veeam logo, the product name, the server time (Nov 17, 2023 7:23 PM), the user (administrator), and the Configuration icon. The left sidebar contains navigation options: Exit Configuration, Getting Started, Administration (Infrastructure, Accounts, Repositories), Workers (selected), and Server Settings (General, License, Support Information). The main content area is titled 'Profile' and contains an information box with default profile settings. Below this is a search bar for 'Region' and buttons for '+ Add', 'Edit', and 'Remove'. A table lists worker profiles for two regions: europe-west10 and europe-west4. The 'europe-west10' row is selected.

Server time: Nov 17, 2023 7:23 PM | administrator | Portal Administrator | Configuration

Exit Configuration

Getting Started

Administration

- Infrastructure
- Accounts
- Repositories

Workers

Server Settings

- General
- License
- Support Information

Network Profile Labels

Add worker profiles that will be used to deploy worker instances for backup and archive operations in Google Cloud regions.

By default, the following profiles are used:

- Primary profile (e2-highcpu-8) is used while sufficient disk quota is available.
- Secondary profile (e2-highcpu-2) is used only if you are running out of disk quota.
- Archiving profile (e2-standard-4) is used for transferring data to archive repositories.

Region [Search] + Add Edit Remove

Region	Primary Profile	Secondary Profile	Archiving Profile
<input checked="" type="checkbox"/> europe-west10	e2-highcpu-8	e2-highmem-8	e2-standard-4
<input type="checkbox"/> europe-west4	e2-highcpu-8	e2-highcpu-8	e2-standard-4

Selected: 1 of 2

Assigning Worker Instance Labels

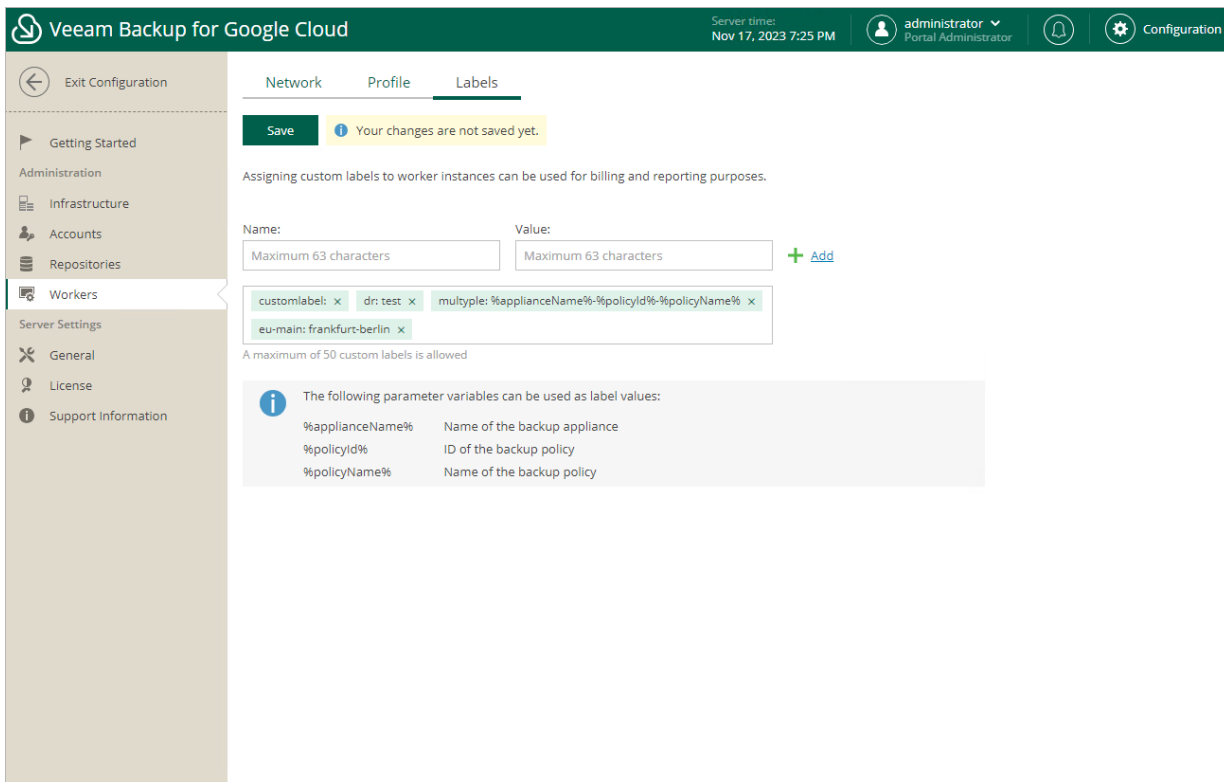
Veeam Backup for Google Cloud allows you to assign labels to worker instances deployed during backup and restore operations. You can then use these labels to track worker instances in Google Cloud for billing and reporting purposes.

To add a new label, do the following:

1. Switch to the **Configuration** page.
2. Navigate to **Workers > Labels**.
3. Use the **Name** and **Value** fields to specify a name and a value for the label, and then click **Add**. Note that you cannot add more than 50 labels.
3. Click **Save**.

NOTE

The `veeamvid` label is assigned to all newly deployed worker instances automatically and is reserved by Veeam Backup for Google Cloud for internal purposes.



The screenshot displays the Veeam Backup for Google Cloud configuration interface. The top navigation bar shows the server time as Nov 17, 2023 7:25 PM and the user as administrator (Portal Administrator). The left sidebar contains navigation options: Exit Configuration, Getting Started, Administration (Infrastructure, Accounts, Repositories), Workers, Server Settings (General, License, Support Information), and Configuration. The main content area is titled 'Labels' and includes a 'Save' button with a warning: 'Your changes are not saved yet.' Below this, a message states: 'Assigning custom labels to worker instances can be used for billing and reporting purposes.' There are two input fields for 'Name' and 'Value', both with a 'Maximum 63 characters' limit and an '+ Add' button. The 'Name' field contains 'customlabel: x', 'dr: test x', and 'multiple: %applianceName%-%policyId%-%policyName% x'. The 'Value' field contains 'eu-main: frankfurt-berlin x'. A note below the fields states: 'A maximum of 50 custom labels is allowed.' An information icon (i) is followed by the text: 'The following parameter variables can be used as label values:'

Parameter Variable	Description
%applianceName%	Name of the backup appliance
%policyId%	ID of the backup policy
%policyName%	Name of the backup policy

Configuring General Settings

Veeam Backup for Google Cloud allows you to configure general settings that are applied to all performed operations and deployed architecture components:

- [Define for how long obsolete snapshots and session records will be retained.](#)
- [Provide certificates to secure connections between Veeam Backup for Google Cloud architecture components.](#)
- [Configure notification settings for automated delivery of reports.](#)
- [Change the time zone set on the backup appliance.](#)
- [Register applications to be able to grant permissions to service accounts in the Google Cloud console automatically.](#)

Configuring Global Retention Settings

You can configure global retention settings to specify for how long the following data will be retained in the configuration database:

- [Obsolete snapshots](#)
- [Session records](#)

Configuring Retention Settings for Obsolete Snapshots

If an instance (whether it is a VM instance, a Cloud SQL instance or a Cloud Spanner 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 become obsolete. These snapshots are removed from the configuration database according to their own retention settings.

NOTE

Global retention settings apply to all cloud-native snapshots created by the Veeam backup service. If an instance is still processed by a backup policy, but some of its cloud-native snapshots are older than the number of days (or months) specified in the global retention settings, these cloud-native snapshots will be removed from the configuration database.

To configure retention settings for obsolete snapshots, do the following:

1. Switch to the **Configuration** page.
2. Navigate to **General > Retention**.
3. In the **Obsolete snapshots retention** section, select either of the following options:
 - Select the **Never** option if you do not want Veeam Backup for Google Cloud to remove obsolete snapshots.
 - Select the **After** option if you want to specify the number of days (or months) during which Veeam Backup for Google Cloud will keep obsolete snapshots in the configuration database. The number must be between 15 and 36135 for days, and between 1 and 1188 for months.

If you select this option, Veeam Backup for Google Cloud will remove obsolete instance snapshots from the configuration database as soon as the specified period of time is over – even if the instances are still processed by backup policies.
4. Click **Save**.

NOTE

When Veeam Backup for Google Cloud removes an obsolete snapshot from the configuration database, it also removes the snapshot from Google Cloud Storage.

Configuring Retention Settings for Session Records

Veeam Backup for Google Cloud 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:

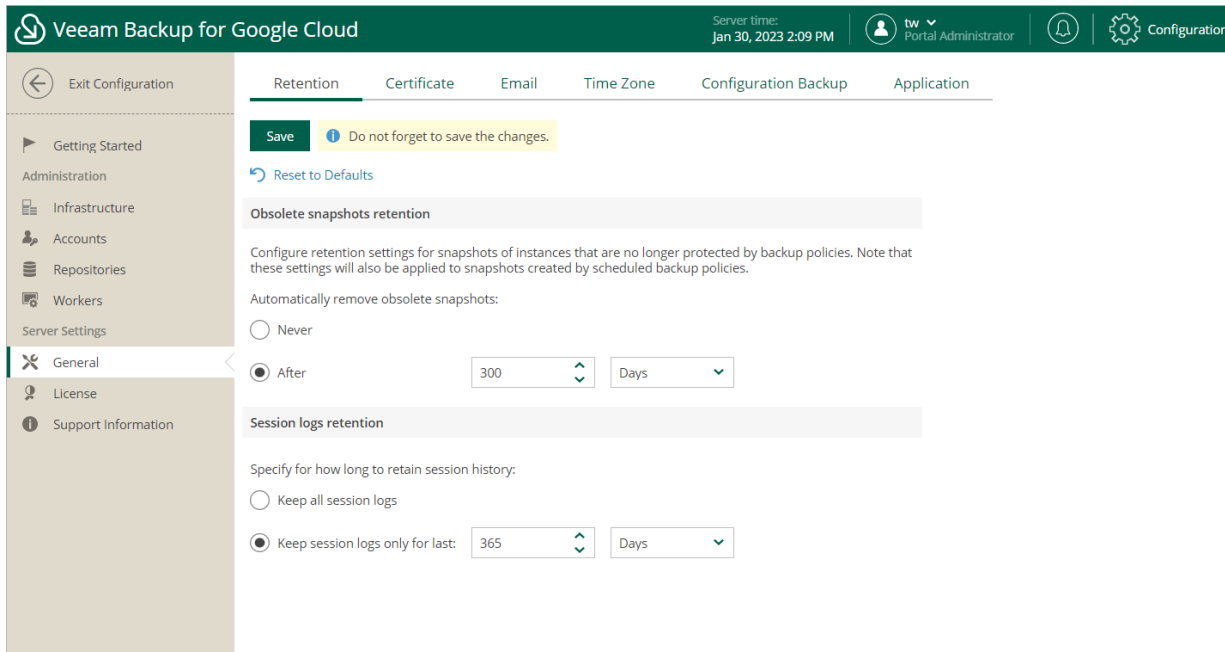
1. In the **Session logs retention** section, select either of the following options:
 - o Select the **Keep all session logs** option if you do not want Veeam Backup for Google Cloud to remove session records.
 - o Select the **Keep session logs only for last** option if you want to specify the number of days (or months) during which Veeam Backup for Google Cloud will keep session records in the configuration database.

If you select this option, Veeam Backup for Google Cloud 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 disk. By default, the disk comes with 20 GB of storage capacity. If you choose not to remove sessions records at all, consider increasing the disk space 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 for all backup policy and snapshot retention sessions run within the past 24-hour period.

To connect an email server that will be used for sending email notifications:

1. Switch to the **Configuration** page.
2. Navigate to **General > Email**.
3. Select the **Enable email notifications** check box.
4. Click the link in the **Email server** field and configure [email server settings](#).
5. In the **From** field, enter an email address of the notification sender. This email address will be displayed in the **From** field of notifications.
6. In the **To** field, enter an email address of a recipient. Use a semicolon to separate multiple recipient addresses.

For each particular policy, you can specify additional recipients. For more information, see [Creating Backup Policies](#).

NOTE

If you specify the same email recipient in both backup policy notification and global notification settings, Veeam Backup for Google Cloud will send each notification to this recipient twice.

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.
8. In the **Notify 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 Google Cloud allows you to send a test message to check whether you have configured the settings correctly. To do that, click **Send Test Email**. A test message will be sent to the specified email address.

Configuring Email Server Settings

To configure email server settings, choose whether you want to employ [Basic \(SMTP\)](#) or [Modern \(OAuth 2.0\)](#) authentication for your email server.

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 **Email 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 587.
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 Google Cloud 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 **Save**.

Using Modern Authentication


To employ the Modern authentication to connect to your email server:

1. In **Email Server Settings** window, copy the URL from the **Redirect URL** field.
2. For Veeam Backup for Google Cloud 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 Google Cloud Web UI.

3. Back to the Veeam Backup for Google Cloud Web UI, do the following in **Email Server Settings** window:
 - a. From the **Authentication** drop-down list, select *Modern*.
 - b. Use the **Email server** drop-down list to choose whether the server that you want to use to send email notifications is a *Google* or *Microsoft* email server.
 - c. 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](#).
 - d. [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.

- e. Click **Authorize**. You will be redirected to the authorization page. Sign in using a Google or Microsoft Azure account to validate the configured settings.



Adding SMTP Accounts

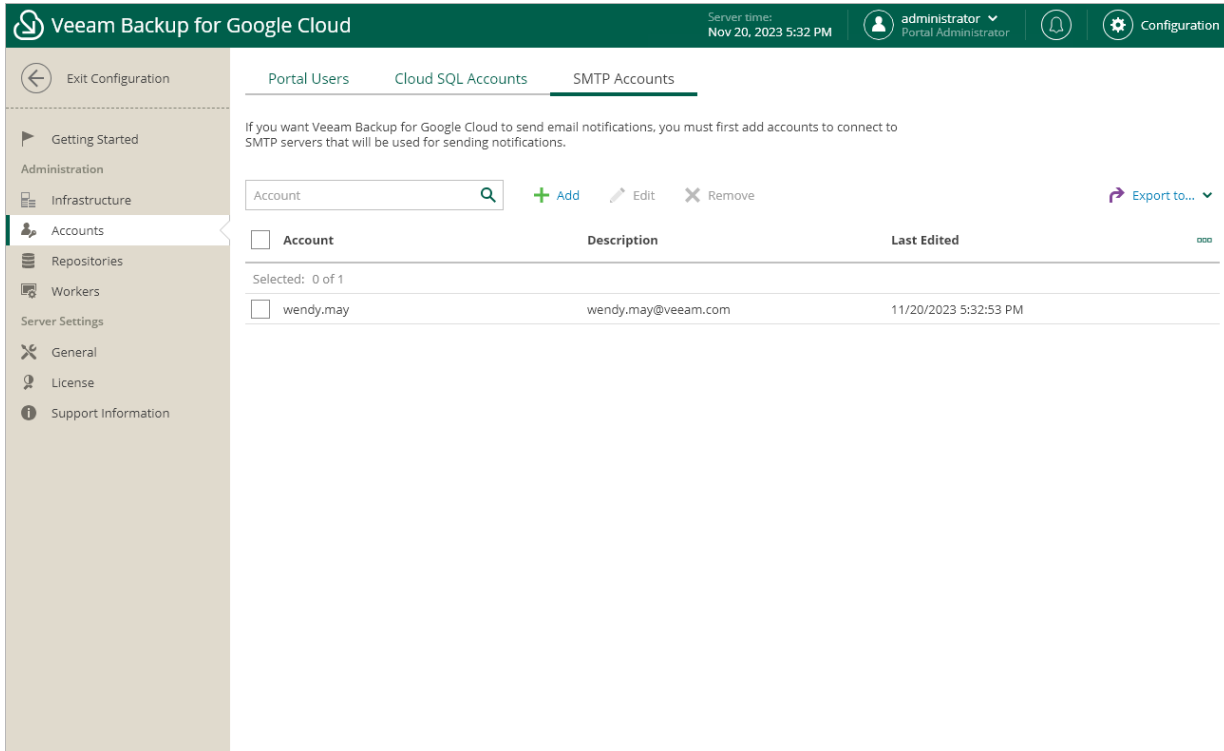
To add an account that will be used to connect to an SMTP server, do the following:

1. [Launch the Add Account wizard.](#)
2. [Specify an account display name and description.](#)
3. [Provide credentials.](#)
4. [Finish working with the wizard.](#)

Step 1. Launch Add Account Wizard

To launch the **Add Account** wizard, do the following:

1. Switch to the **Configuration** page.
2. Navigate to **Accounts > SMTP Accounts**.
3. Click **Add**.

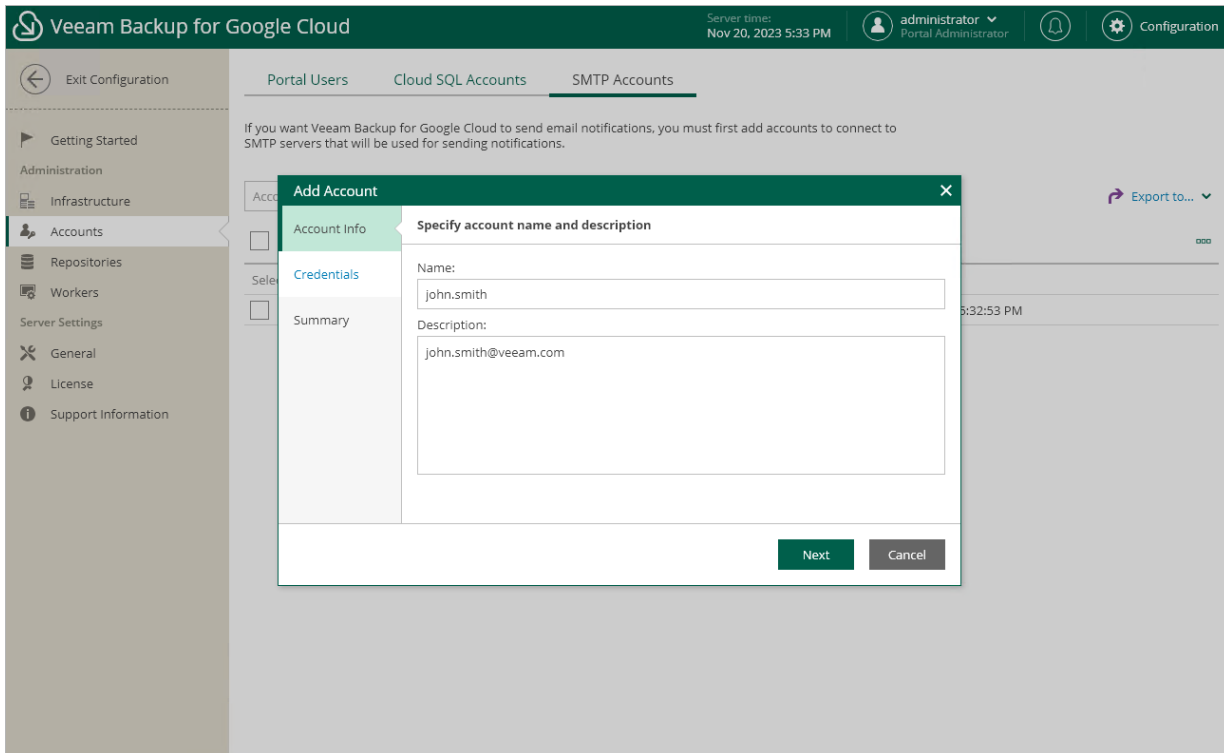


The screenshot displays the Veeam Backup for Google Cloud configuration interface. The top navigation bar includes the Veeam logo, the product name, the server time (Nov 20, 2023 5:32 PM), the user role (administrator, Portal Administrator), and the Configuration page indicator. The left sidebar shows the navigation menu with 'Accounts' selected. The main content area is titled 'SMTP Accounts' and contains a search bar, an '+ Add' button, and an 'Export to...' dropdown. Below this is a table with columns for 'Account', 'Description', and 'Last Edited'. One account is listed: 'wendy.may' with description 'wendy.may@veeam.com' and last edited '11/20/2023 5:32:53 PM'. The 'Selected: 0 of 1' indicator is visible above the table.

Account	Description	Last Edited	
<input type="checkbox"/>	wendy.may	wendy.may@veeam.com	11/20/2023 5:32:53 PM


Step 2. Specify Account Name and Description

At the **Account Info** step of the wizard, use the **Name** and **Description** fields to enter a name for the new SMTP account and to provide a description for future reference. The maximum length of the name is 255 characters.




Step 3. Provide Credentials

At the **Account** step of the wizard, specify credentials of a user account that will be used to authenticate against the SMTP server.



Step 4. Finish Working with Wizard

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

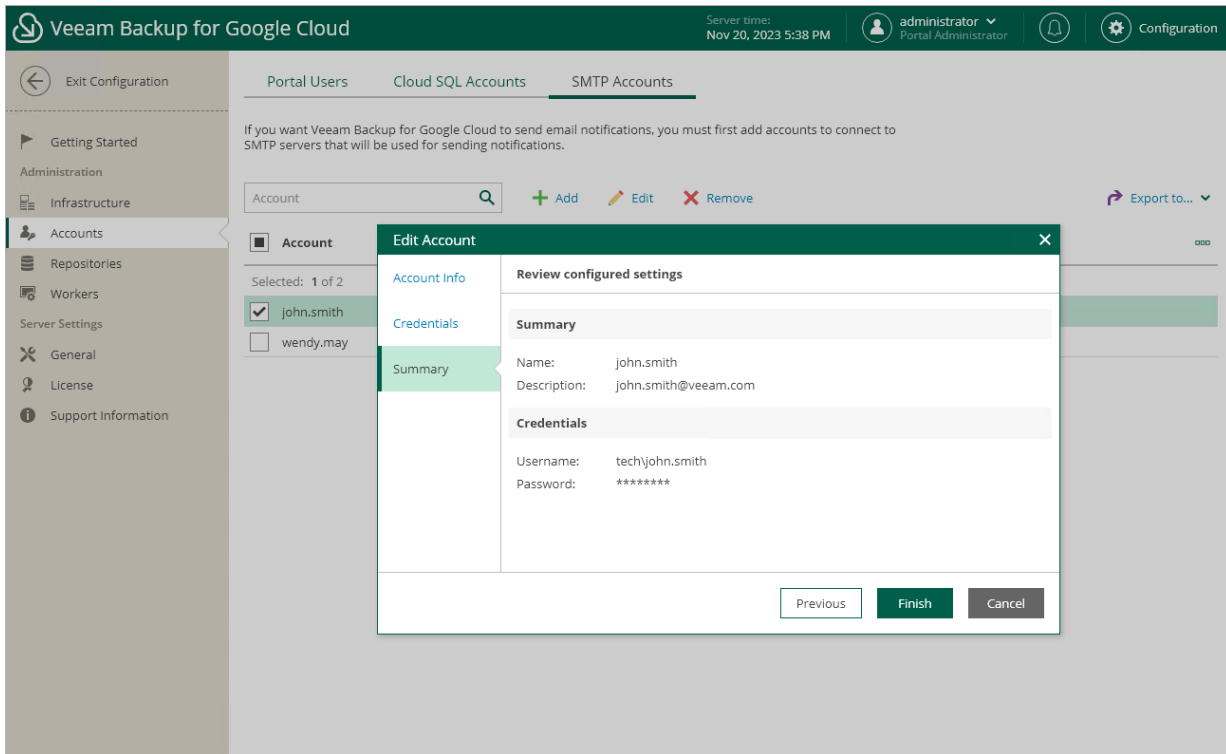


Editing SMTP Accounts

For each SMTP account, you can modify settings configured while creating the account:

1. Switch to the **Configuration** page.
2. Navigate to **Accounts > SMTP Accounts**.
3. Select the account and click **Edit**.
4. 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 2).
 - 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 3).

- c. At the **Summary** step of the wizard, review summary information and click **Finish** to confirm the changes.



Replacing Security Certificates

To establish secure data communications between the backup appliance and web browsers running on user workstations, Veeam Backup for Google Cloud uses Transport Layer Security (TLS) certificates.

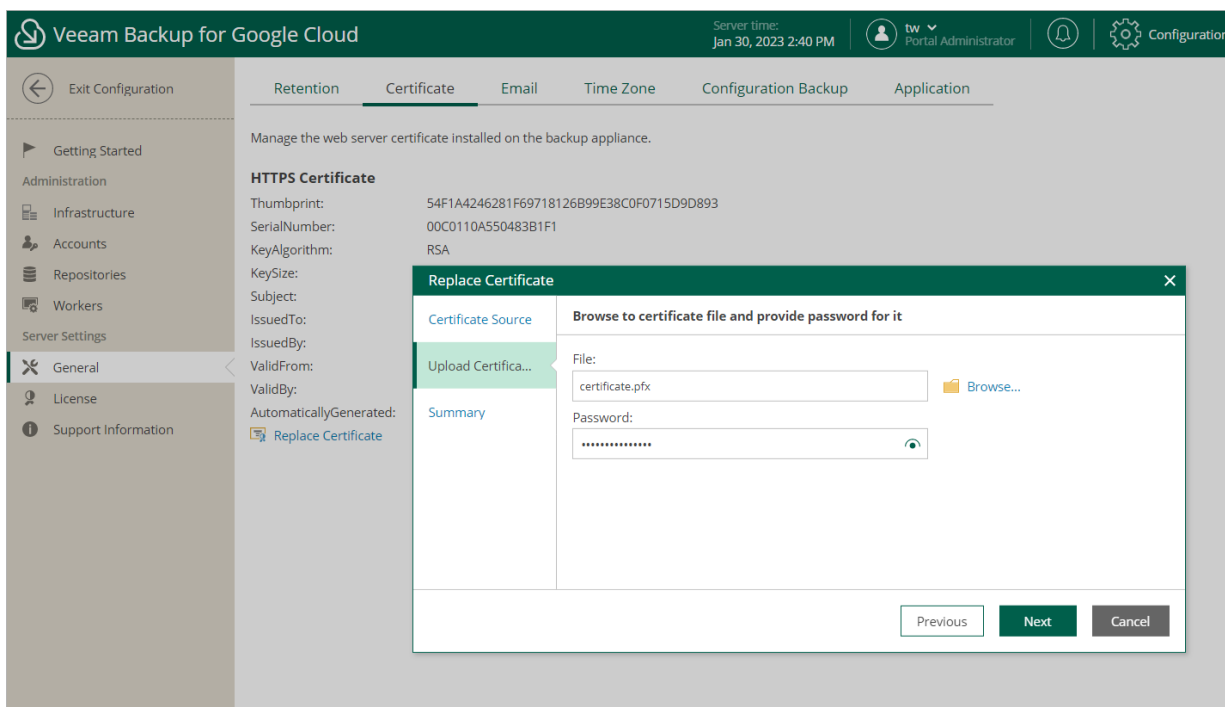
When you install Veeam Backup for Google Cloud, it automatically generates a default self-signed certificate. You can replace this default certificate with your own self-signed certificate or with a certificate obtained from a Certificate Authority (CA). To replace the currently used TLS certificate, do the following:

1. Switch to the **Configuration** page.
2. Navigate to **General > Certificate**.
3. Click **Replace Certificate**.
4. Complete the **Replace Certificate** wizard:
 - a. At the **Certificate Source** step of the wizard, do the following:
 - § Select the **Re-create the self-signed certificate** option if you want to replace the existing certificate with a new self-signed certificate automatically generated by Veeam Backup for Google Cloud.
 - § Select the **Upload a new 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 a new certificate** option] At the **Upload Certificate** step of the wizard, browse to the certificate that you want to install, and provide a password for the certificate file.

NOTE

Only the PFX and P12 certificate formats are supported.

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



Changing Time Zone

Veeam Backup for Google Cloud runs daily reports and performs all data protection and disaster recovery operations according to the time zone set on the backup appliance.

IMPORTANT

If Daylight Saving Time (DST) is used in the time zone set on the backup appliance, consider the following:

- When DST starts (clocks are set one hour forward), all policy sessions scheduled to launch at the skipped hour on this day do not run.
- When DST ends (clocks are set one hour back), all policy sessions scheduled to launch at the duplicated hour on this day run only once.

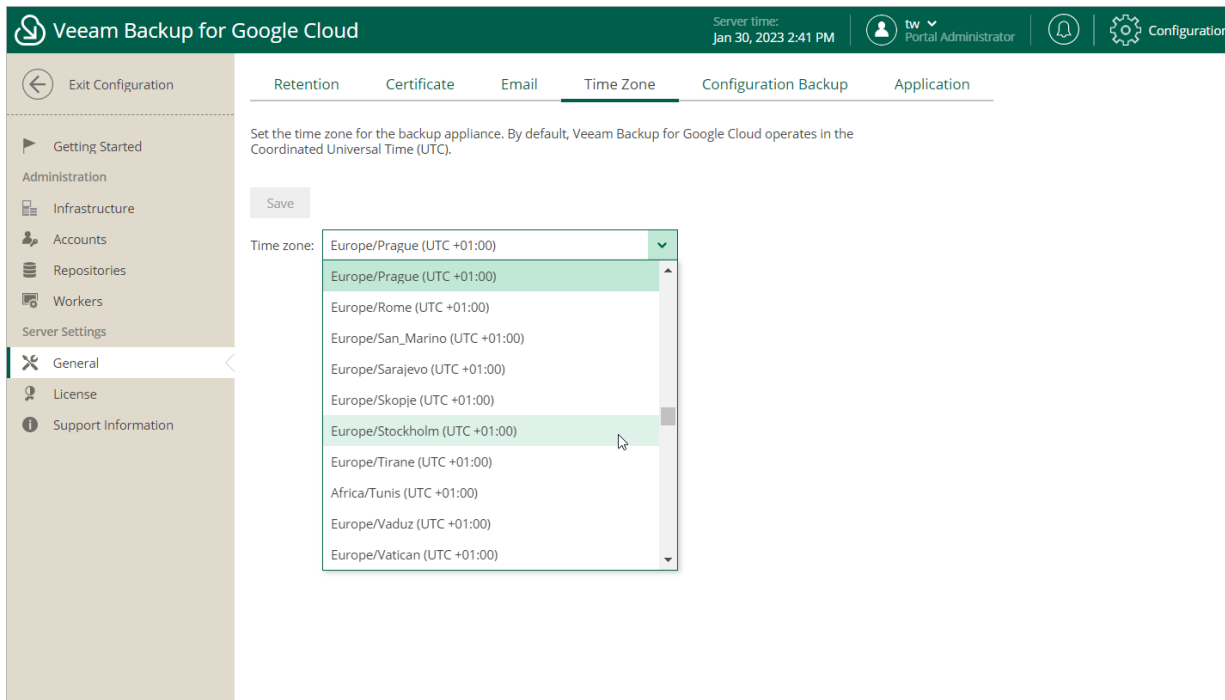
Since the backup appliance is deployed on a VM instance in Google Cloud, the time zone is set to Coordinated Universal Time (UTC) by default. However, you can change the time zone if required. For example, you may want the time on the backup appliance to match the time on the workstation from which you access Veeam Backup for Google Cloud.

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

You cannot 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 try changing the time zone again.



Registering Application

To allow Veeam Backup for Google Cloud to perform data protection and disaster recovery operations for resources in Google Cloud projects and folders, service accounts associated with the projects and folders must have specific permissions required to access these resources. If any of the permissions listed in section [Planning and Preparation](#) are missing for a service account, you can grant them in the Google Cloud console automatically, without leaving the Veeam Backup for Google Cloud UI. However, since this functionality employs the OAuth 2.0 protocol to access Google Cloud APIs, you must do the following:

1. In the Google Cloud console, configure the OAuth consent screen as described in [Google Cloud documentation](#).

Consider that Veeam Backup for Google Cloud requires the `https://www.googleapis.com/auth/cloud-platform` scope to be identified for the application in the OAuth consent screen. For more information on OAuth 2.0 Scopes for Google APIs, see [Google Cloud documentation](#).

2. Set up a DNS hostname for the VM instance running Veeam Backup for Google Cloud (for example, using [Cloud DNS](#)).

Due to Google Cloud limitations, the OAuth consent screen cannot use public IP addresses as redirect URIs for OAuth 2.0 authorization. For more information on redirect URI validation rules, see [Google Cloud documentation](#).

3. Access the Veeam Backup for Google Cloud UI using the DNS hostname of the backup appliance, switch to the **Configuration** page, navigate to **General** > **Application**, set the **Register** toggle to *On*, and copy the address displayed in the **Redirect URL** field.

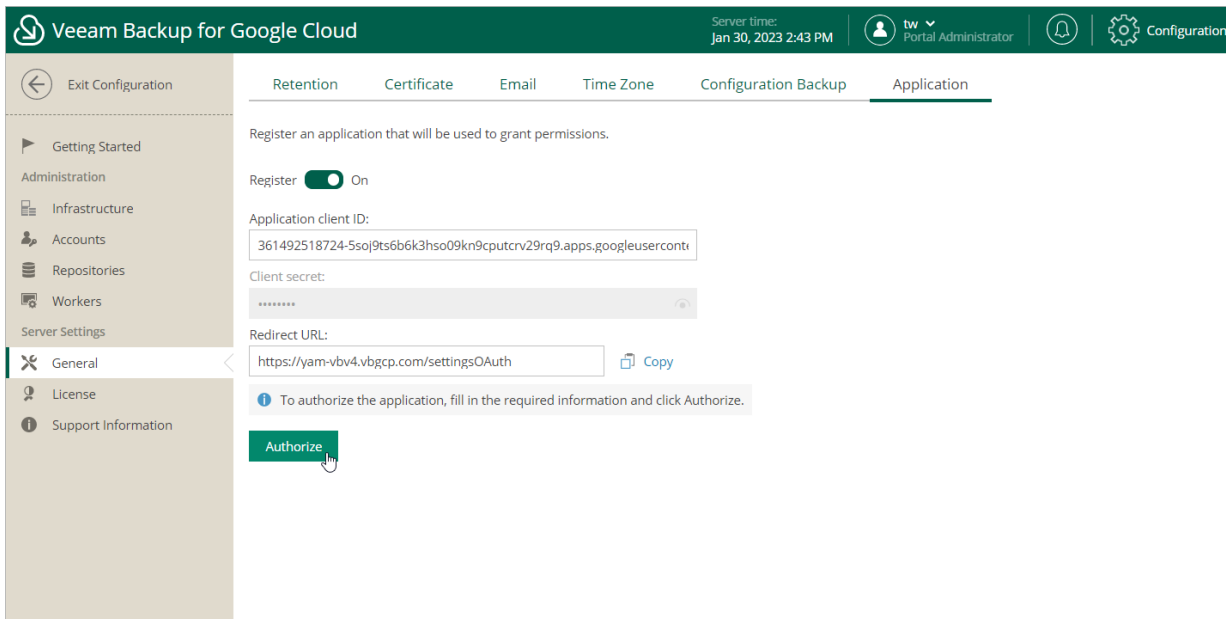
To learn how to access Veeam Backup for Google Cloud UI, see [Accessing Veeam Backup for Google Cloud](#).

4. Back to the Google Cloud console, create OAuth client ID credentials as described in [Google Cloud documentation](#).

In the **Authorized redirect URIs** section of the **Create OAuth client ID** page, add the address copied from the Veeam Backup for Google Cloud UI.

5. Back to the Veeam Backup for Google Cloud UI, on the **Application** tab, provide the Client ID and Client secret used to authorize access to the configured OAuth consent screen, and then click **Authorize**.

You will be redirected to the OAuth consent screen authorization page. Sign in using a Google account to validate the configured settings.



Performing Configuration Backup and Restore

You can back up and restore the configuration database that stores data collected from Veeam Backup for Google Cloud for the existing backup policies, protected VM, Cloud SQL and Cloud Spanner instances, connected Google Cloud projects, 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 configuration backup. You can also use a configuration backup to migrate the configuration of one backup appliance to another backup appliance in Google Cloud.

It is recommended that you regularly perform configuration backup for every backup appliance present in Google Cloud. 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 Google Cloud to do it automatically on a regular basis.

Performing Configuration Backup

During the configuration backup, Veeam Backup & Replication exports data from the configuration database of an appliance and saves it to a backup file in a repository. The configuration database contains the following information: the existing backup policies, protected VM, Cloud SQL and Cloud Spanner instances, connected Google Cloud projects, logged session records and so on.

Performing Configuration Backup Using Console

When Veeam Backup & Replication performs configuration backup, it 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 backups, 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 Google Cloud 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 can process 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 the configuration backup settings. Backups are saved to the \\VeeamConfigBackup\GCP folder.

NOTES

- It is not recommended to store configuration backups in any folder on the backup server. Otherwise, you will not be able to restore the configurations 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 Google Cloud 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 Google Cloud 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 **General > 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 Google Cloud as described in section [Adding Backup Repositories](#). The **Repository** list shows only backup repositories of the *Standard* and *Nearline* storage classes that have encryption enabled.

As soon as you click **Create**, Veeam Backup for Google Cloud 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 Google Cloud creates a successful configuration backup, you can click **Export Last Backup** to download the backup file and then use it to [restore configuration data](#).


Performing Configuration Backup Automatically

To instruct Veeam Backup for Google Cloud to back up the configuration database automatically by schedule, do the following:

1. Switch to the **Configuration** page.
2. Navigate to **General > Configuration Backup**.
3. In the **Backup schedule** section, set the **Enable scheduling** toggle to *On*.
4. Click **Choose** in the **Repository** field, and use the list of available repositories in the **Choose Repository** window to select a repository where configuration backups will be stored.

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

5. In the **Keep restore points for** field, specify the number of days for which you want to keep restore points in the selected backup repository.
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 Google Cloud 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, do the following:

1. Switch to the **Configuration** page.
2. Navigate to **General > Configuration Backup**.

3. Use either 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 Google Cloud will save the exported backup file to the default download directory on the local machine.



Performing Configuration Restore

Veeam Backup for Google Cloud 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.

Restoring Configuration Data Using Console

To restore the configuration database of a backup appliance using the Veeam Backup & Replication console, do the following:

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:

- Configuration restore of backup appliances that run Veeam Backup for Google Cloud version 2.0 is not supported.
- Make sure that there are no sessions currently running on the backup appliance. Also, make sure that 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](#).
- 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 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.

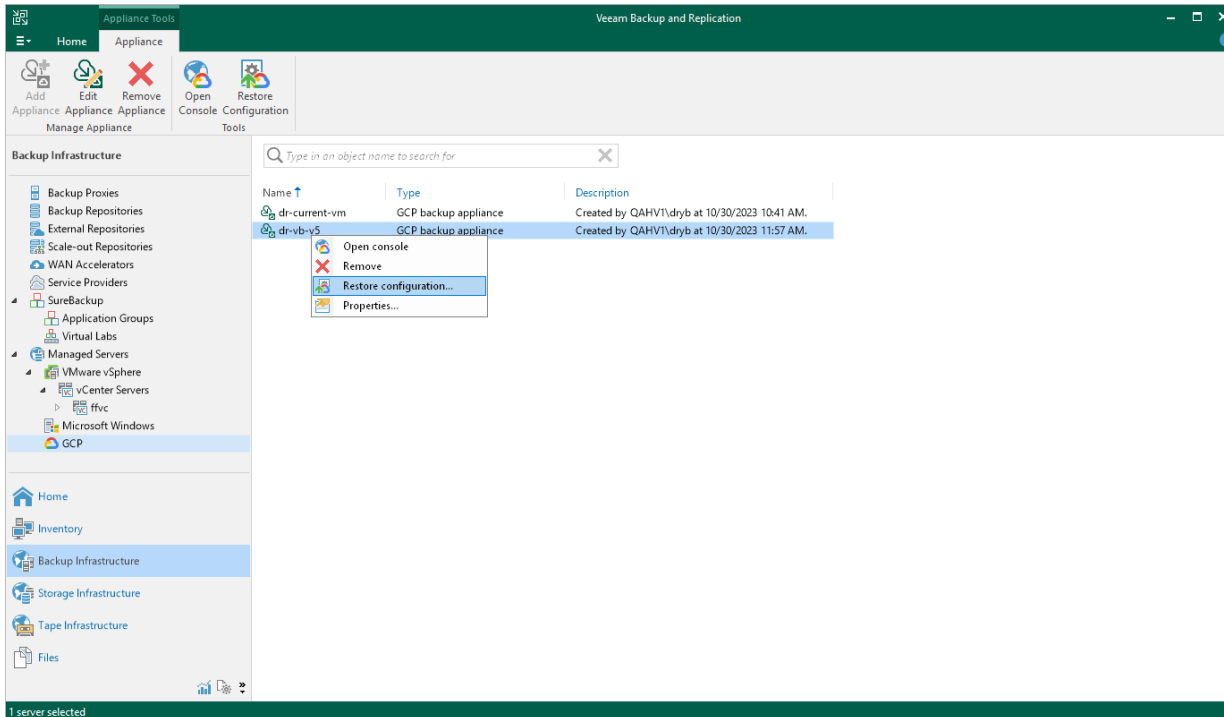
- Make sure that repositories added to the backup appliance are not managed by any other appliances. Otherwise, retention sessions running on different appliances may corrupt backup files stored in the repositories, which may result in unpredictable data loss.
- The backup appliance to which you restore the configuration preserves its TLS certificate.
- During configuration restore, Veeam Backup & Replication will overwrite custom settings of the Linux configuration file on the backup appliance with the settings saved in the configuration backup file.

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 > GCP**.
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**.



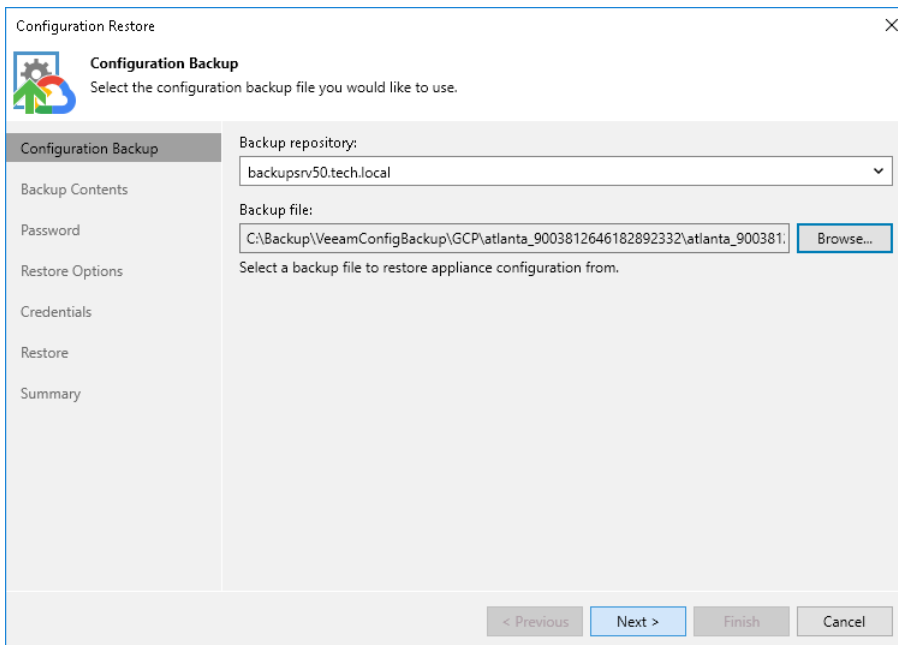
Step 2. Choose Backup File

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 in the 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.




Step 3. Review Backup File Info

Veeam Backup & Replication will analyze the content of the selected backup file 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

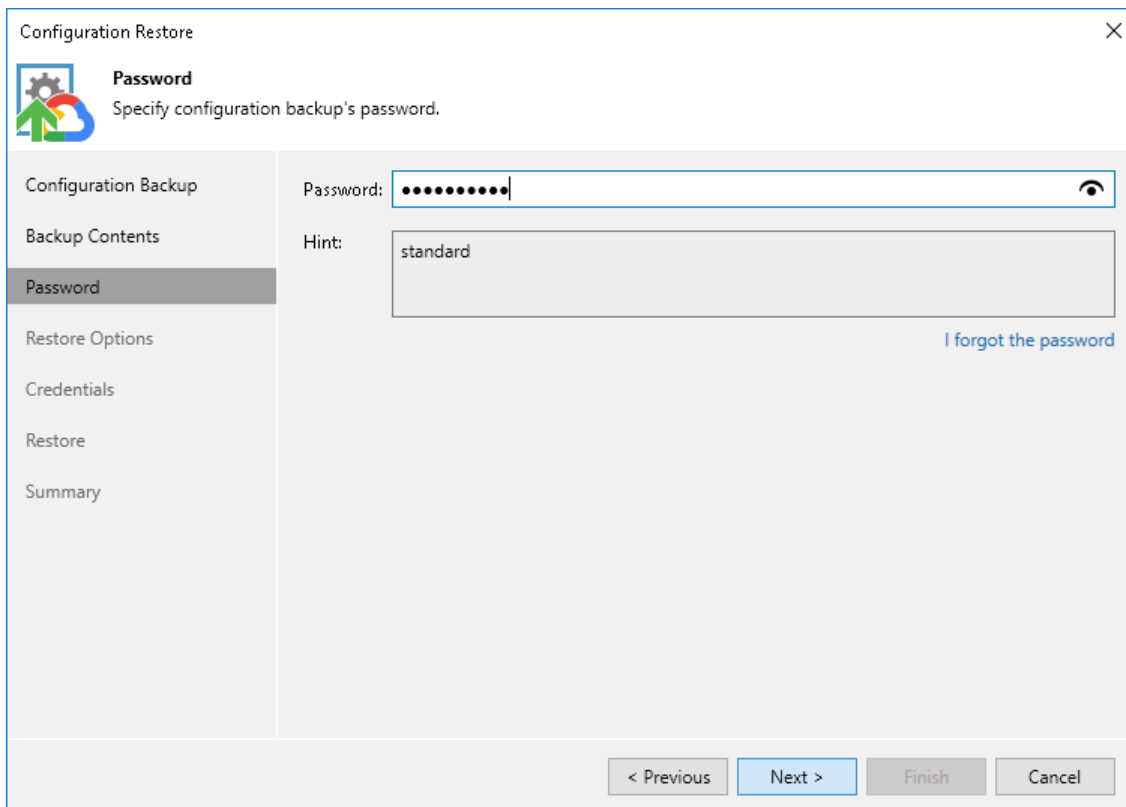
At the **Password** step of the wizard, specify the password used to encrypt the configuration backup file.

If you do not remember the password, you can restore configuration backup data without providing it. To do that, click the **I forgot the password** link and follow the instructions provided in the Veeam Backup & Replication User Guide, section [Decrypting Data Without Password](#).

NOTE

To restore configuration data without a password, the following requirements must be met:

- You must have either the Veeam Universal License or a legacy socket-based license (Enterprise edition or higher) installed on the backup server.
- The backup server must be connected to Veeam Backup Enterprise Manager, and password loss protection must be enabled on the Veeam Backup Enterprise Manager side for the duration of both the backup and restore operations. For more information, see the [Veeam Backup Enterprise Manager Guide](#).



The screenshot shows the 'Configuration Restore' wizard window. The title bar reads 'Configuration Restore' with a close button (X) on the right. Below the title bar is a navigation pane on the left with the following items: 'Configuration Backup', 'Backup Contents', 'Password' (highlighted), 'Restore Options', 'Credentials', 'Restore', and 'Summary'. The main area of the window is titled 'Password' and contains the instruction 'Specify configuration backup's password.' Below this instruction, there is a 'Password:' field with a masked password of ten dots and a visibility toggle icon (an eye). Underneath the password field is a 'Hint:' field containing the text 'standard'. To the right of the hint field is a blue link that says 'I forgot the password'. At the bottom of the window, there are four buttons: '< Previous' (disabled), 'Next >' (active), 'Finish' (disabled), and 'Cancel' (disabled).

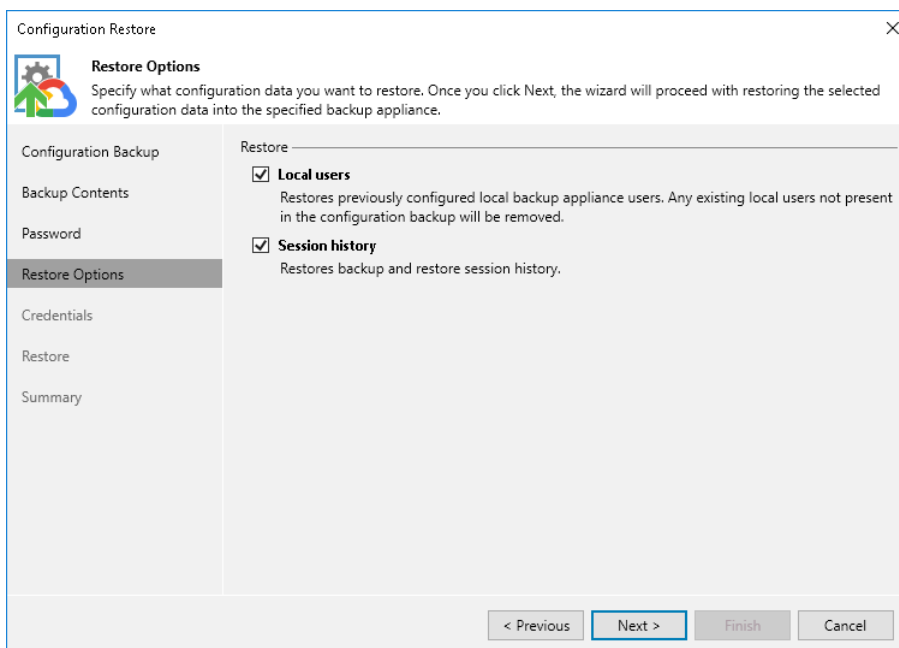
Step 5. Choose Restore Options

By default, Veeam Backup & Replication 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 and portal users of the initial backup appliance as well.

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]

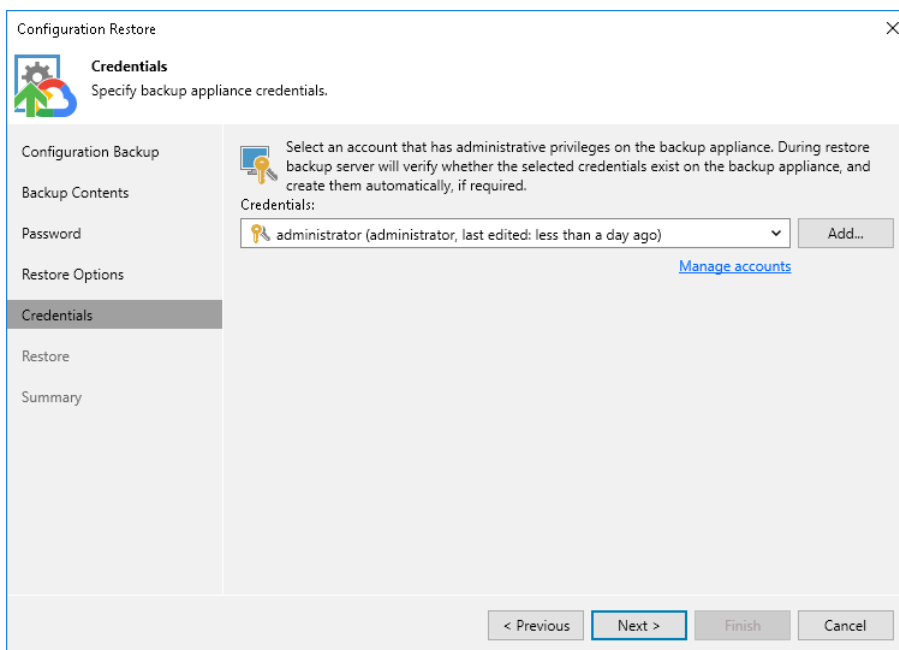
After the configuration restore process completes, Veeam Backup & Replication will try to connect to the backup appliance using credentials of the user specified [when adding the appliance](#) to the backup infrastructure. However, since you have chosen to restore all users saved to the configuration 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. Below the title bar, there is a Veeam logo and the text 'Credentials' and 'Specify backup appliance credentials.' The main area is divided into a left sidebar and a main content area. The sidebar contains a list of steps: 'Configuration Backup', 'Backup Contents', 'Password', 'Restore Options', 'Credentials' (which is highlighted), 'Restore', and 'Summary'. The main content area contains the following text: 'Select an account that has administrative privileges on the backup appliance. During restore backup server will verify whether the selected credentials exist on the backup appliance, and create them automatically, if required.' Below this text is a 'Credentials:' label and a dropdown menu showing 'administrator (administrator, last edited: less than a day ago)'. To the right of the dropdown is an 'Add...' button. Below the dropdown is a blue link labeled 'Manage accounts'. At the bottom of the window, there are four buttons: '< Previous', 'Next >' (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**.



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 Google Cloud Web UI where you can view the details of the occurred issue. 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 Google Cloud Web UI, do the following:

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.](#)


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 Google Cloud from the Web UI. In this case, you can perform configuration restore using the Veeam Backup & Replication console as described in section [Performing Configuration Restore Using Console](#).
- Before you start the restore process, stop all backup policies that are currently running.

Step 1. Launch Configuration Restore Wizard

To launch the **Configuration Restore** wizard, do the following:

1. Switch to the **Configuration** page.
2. Navigate to **General > Configuration Backup**.
3. In the **Configuration restore** section, click **Restore**.



Step 2. Choose Backup File

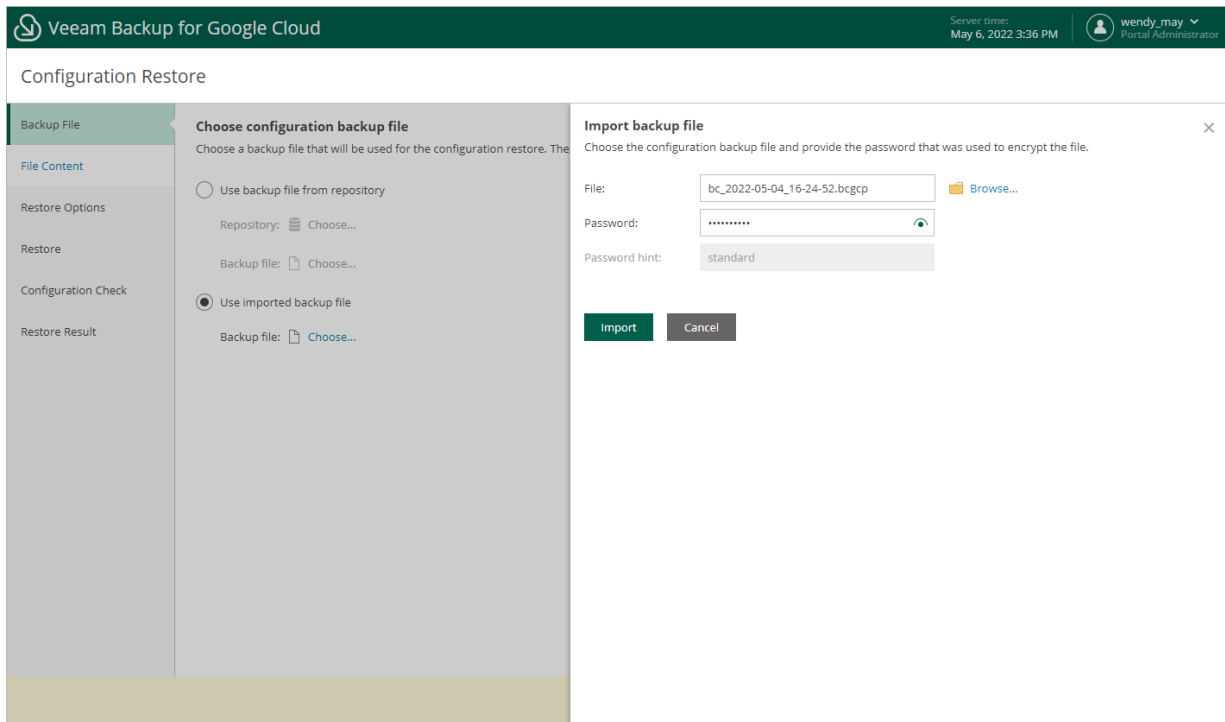
At the **Backup File** step of the wizard, choose whether you want to use an exported backup file or a backup file stored in a backup repository:

- 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 **Choose** in the **Repository** field, and use the list of available repositories in the **Choose repository** window to select the repository where the necessary configuration backup file is stored.

For a backup repository to be displayed in the **Repository** list, it must be added to Veeam Backup for Google Cloud as described in section [Adding Backup Repositories](#). The list shows only backup repositories that store configuration backup files.
 - b. Click **Choose** in the **Backup file** field, and select the necessary file in the **Choose backup file** window.
- 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 **Choose** 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**.

IMPORTANT

The size of an uploaded backup file must not exceed 10 GB. To upload a file of a bigger size, open a [support case](#).



Step 3. Review Backup File Info

Veeam Backup for Google Cloud 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 Google Cloud that was installed on the initial backup appliance and the version of the File-Level Recovery Service that was running on the appliance.
- **Product configuration** — configuration data saved in the file (such as the number of configured backup policies, added user accounts, created backup 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 Google Cloud. 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:	05/04/2022 10:24:53 AM
----------------	------------------------

Product information

Product name:	Veeam Backup for Google Cloud
Product version:	3.0.0.823
File-level recovery service version:	5.0.0.579

Product configuration

Standard repositories:	7
Archive repositories:	1
VM backup policies:	3
Cloud SQL backup policies:	1
Portal users:	12
Sessions:	1183

At the bottom of the wizard, there are three buttons: 'Previous', 'Next', and 'Cancel'. The 'Next' button is highlighted in green.

Step 4. Choose Restore Options

By default, Veeam Backup for Google Cloud restores only configuration data for the existing architecture 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 and user accounts of the initial backup appliance as well.

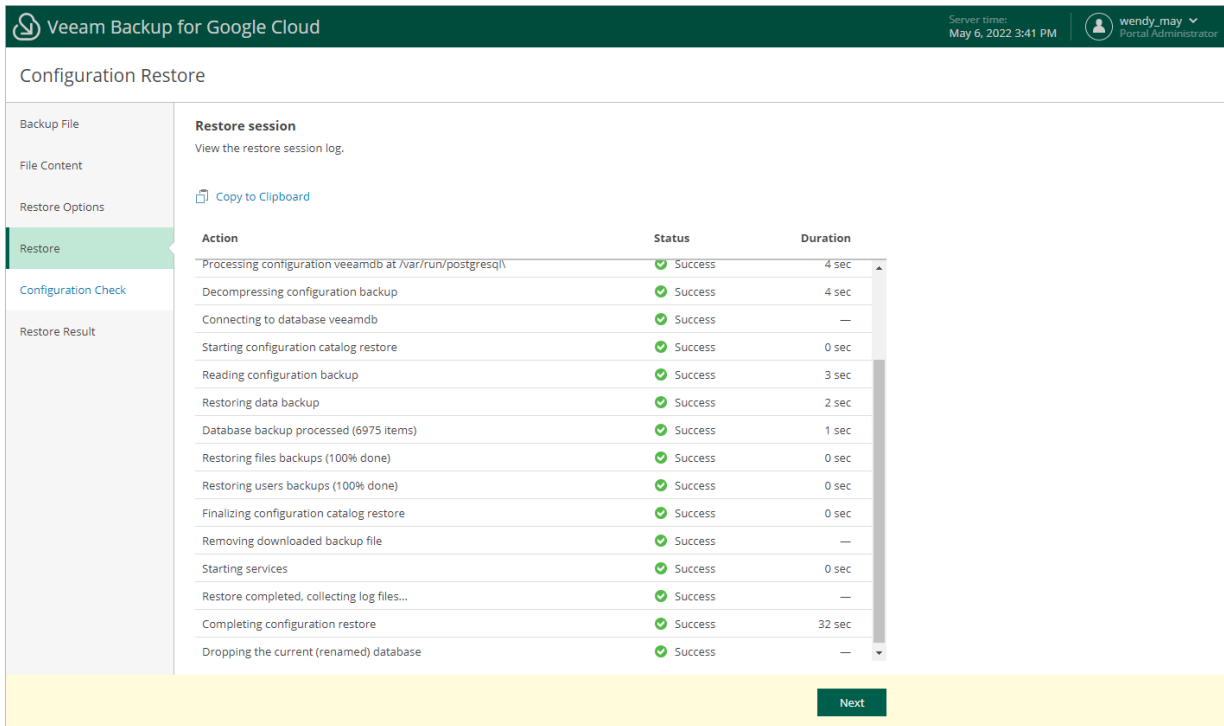
IMPORTANT

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

The screenshot shows the 'Configuration Restore' wizard in Veeam Backup for Google Cloud. The interface is divided into a left sidebar and a main content area. The sidebar contains a vertical list of steps: 'Backup File', 'File Content', 'Restore Options' (highlighted in green), 'Restore', 'Configuration Check', and 'Restore Result'. The main content area is titled 'Specify restore options' and includes the instruction: 'Choose configuration data to restore and click Start Restore to perform the restore operation.' Below this instruction are two checked checkboxes: 'Restore session history' (with the subtext 'Restore all backed-up policy sessions from the configuration backup file.') and 'Restore local users' (with the subtext 'Restore all backed-up local users from the configuration backup file.'). At the bottom of the wizard, there is a yellow bar containing three buttons: 'Previous', 'Restore', and 'Cancel'.

Step 5. Track Restore Progress

Veeam Backup for Google Cloud 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 in Veeam Backup for Google Cloud. The interface includes a top navigation bar with the Veeam logo, the product name, server time (May 6, 2022 3:41 PM), and a user profile (wendy_may, Portal Administrator). The main content area is titled 'Configuration Restore' and features a left-hand navigation pane with options: Backup File, File Content, Restore Options, Restore (selected), Configuration Check, and Restore Result. The 'Restore' section is active, displaying a 'Restore session' log. A 'Copy to Clipboard' button is visible above a table of actions. The table lists various steps, all of which are marked as 'Success' with a green checkmark. The final step is 'Dropping the current (renamed) database'. A 'Next' button is located at the bottom right of the screen.

Action	Status	Duration
Processing configuration veeamdb at /var/run/postgresql/	Success	4 sec
Decompressing configuration backup	Success	4 sec
Connecting to database veeamdb	Success	—
Starting configuration catalog restore	Success	0 sec
Reading configuration backup	Success	3 sec
Restoring data backup	Success	2 sec
Database backup processed (6975 items)	Success	1 sec
Restoring files backups (100% done)	Success	0 sec
Restoring users backups (100% done)	Success	0 sec
Finalizing configuration catalog restore	Success	0 sec
Removing downloaded backup file	Success	—
Starting services	Success	0 sec
Restore completed, collecting log files...	Success	—
Completing configuration restore	Success	32 sec
Dropping the current (renamed) database	Success	—

Step 6. View Configuration Check Results

After the restore process is over, Veeam Backup for Google Cloud 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 click **Next**.

If Veeam Backup for Google Cloud encounters an issue while performing a verification check, the **Result** column will display a description of the issue. Note that some issues are displayed for informational purposes only and do not require any action at this point. If any actions are required, the **Action** column will provide instructions on how to resolve the issue. For example, to resolve the issue with service account permissions, click **View** in the **Project check** permissions field, and then use the **Project checks** window to grant missing permissions to every service account associated with a specific project.

You can grant the missing permissions to the service account [using the Google Cloud console](#) or instruct Veeam Backup for Google Cloud to do it:

- To grant the missing permissions manually, click **Download Script**. Veeam Backup for Google Cloud will generate a gcloud script that you can run in the Google Cloud console to assign all the necessary permissions to the service account.

The account under which you run the script must have the permissions required both to get and set project IAM policies and to create custom IAM roles (for example, it can have the *iam.securityAdmin* and *iam.roleAdmin* roles assigned). To learn what permissions and roles are required to create custom roles in IAM, see [Google Cloud documentation](#).

- To let Veeam Backup for Google Cloud grant the missing permissions automatically, click **Grant** and then click **Sign in with Google** in the **Grant permissions** window. You will be redirected to the OAuth consent screen authorization page. Sign in using credentials of a Google account that will be used to grant the permissions.

The account under which you sign in to Google Cloud must have the permissions required both to get and set project IAM policies and to create custom IAM roles (for example, it can have the *iam.securityAdmin* and *iam.roleAdmin* roles assigned). To learn what permissions and roles are required to create service account, see [Google Cloud documentation](#).

NOTE

For Veeam Backup for Google Cloud to be able to authorize in Google Cloud, the OAuth consent screen must be configured as described in section [Registering Applications](#). Note that Veeam Backup for Google Cloud does not store in the configuration database the provided Google account credentials and access tokens received during authorization.

To make sure that the missing permissions have been successfully granted, click **Recheck**.


The screenshot shows the Veeam Backup for Google Cloud interface during a Configuration Restore process. The top navigation bar includes the Veeam logo, the product name, and the server time: May 6, 2022 3:44 PM. The main content area is titled "Configuration Restore" and features a left-hand navigation menu with options: Backup File, File Content, Restore Options, Restore, Configuration Check (highlighted), and Restore Result. The main panel displays "Verification steps" with a sub-header: "The check will confirm that the configuration has been restored successfully, and the backup appliance is fully functional." Below this, there are two buttons: "Recheck" and "Export". A table lists the verification items, all of which have a "Success" status. The table has four columns: Type, Status, Action, and Result.

Type	Status	Action	Result
Project check permissions	Success	—	—
Project services (APIs)	Success	—	—
HMAC keys status	Success	—	—
Repository settings	Success	—	—
Repository encryption	Success	—	—
Worker configuration	Success	—	—
Portal users	Success	—	—

At the bottom right of the interface, there is a "Next" button.

Step 7. Finish Working with Wizard

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



Viewing Available Resources

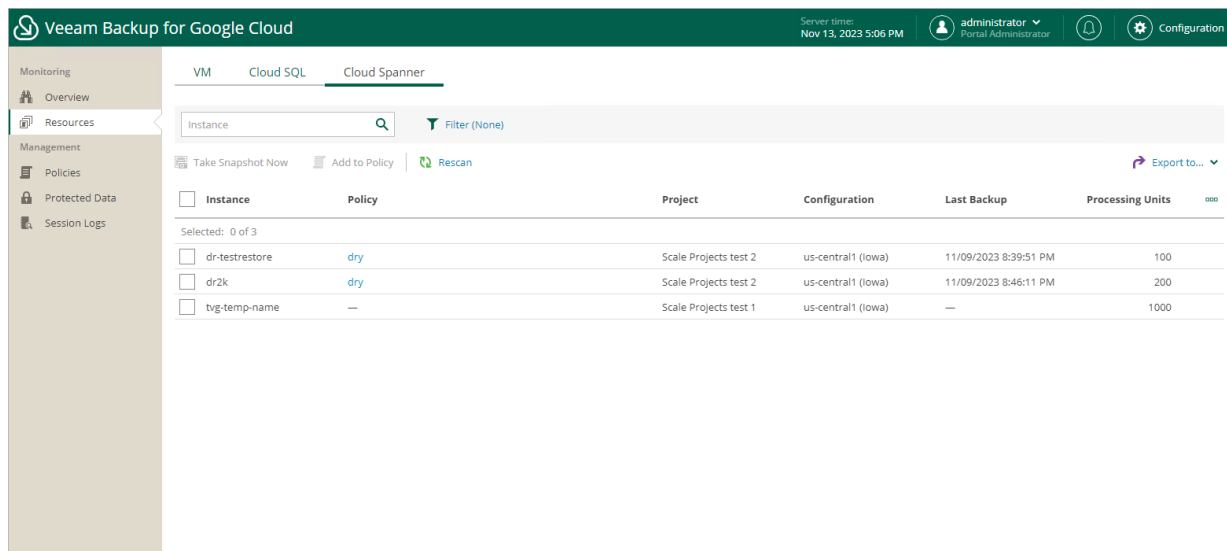
After you create a backup policy to protect a specific type of Google Cloud resources (VM instances, Cloud SQL instances or Cloud Spanner instances), Veeam Backup for Google Cloud rescans Google Cloud regions specified in the policy settings and populates the resource list on the **Resources** tab with all resources of that type residing in these regions. If a region is no longer specified in any backup policy, Veeam Backup for Google Cloud removes all resources residing in the region from the list of available resources.

The **Resources** tab displays Google Cloud resources that can be protected by Veeam Backup for Google Cloud. Each resource is represented with a set of properties, such as:

- **Instance** — the name of the resource.
- **Policy** — the name of the backup policy that protects the resource (if any).
- **Region** — the region in which the resource resides.
- **Project** — the project that manages the resource.
- **Restore Points** — the number of restore points created for the resource (if any).
- **Latest Restore Point** — the date and time of the most recent restore point created for the resource (if any).
- **Destination** — the type of restore points created for the resource (if any).

On the **Resources** tab, you can also perform the following actions:

- Manually create cloud-native snapshots of VM, Cloud SQL and Cloud Spanner instances. For more information, see sections [Performing VM Backup](#), [Performing SQL Backup](#) and [Performing Spanner Backup](#).
- Add VM, Cloud SQL and Cloud Spanner instances to the existing backup policies. For more information, see [Adding Resources to Policies](#).



The screenshot displays the Veeam Backup for Google Cloud interface. The top header shows the server time as Nov 13, 2023 5:06 PM and the user as administrator. The left sidebar contains navigation options: Monitoring (Overview), Resources, and Management (Policies, Protected Data, Session Logs). The main content area is titled 'Cloud Spanner' and features a search bar, a filter dropdown, and buttons for 'Take Snapshot Now', 'Add to Policy', and 'Rescan'. Below these are three resource entries in a table:

Instance	Policy	Project	Configuration	Last Backup	Processing Units
<input type="checkbox"/> dir-testrestore	dry	Scale Projects test 2	us-central1 (Iowa)	11/09/2023 8:39:51 PM	100
<input type="checkbox"/> dir2k	dry	Scale Projects test 2	us-central1 (Iowa)	11/09/2023 8:46:11 PM	200
<input type="checkbox"/> tvg-temp-name	—	Scale Projects test 1	us-central1 (Iowa)	—	1000

Adding Resources to Policies

If you want to include additional resources (VM, Cloud SQL or Cloud Spanner instances) in the existing backup policies, you can either [edit the backup policy settings](#) or quickly add the resources to the policies on the **Resources** tab.

To add a Google Cloud resource to a backup policy, do the following:

1. Navigate to the necessary tab and select the resource.


For a resource to be displayed in the list of available instances, the Google Cloud region in which the resource resides must be specified in any of the configured backup policies, and the service account specified in the backup policy settings must have permissions to access the resource.

2. Click **Add to Policy**.

3. In the **Add to Policy** window, select a backup policy that will protect the resource, and click **Add**.

For a backup policy to be displayed in the list of available policies, the Google Cloud region in which the selected resource resides must be specified in the backup source settings, and the service account used by Veeam Backup for Google Cloud to perform backup must have permissions to access the resource.

4. In the **Results** window, click **OK**.



Performing Backup

With Veeam Backup for Google Cloud, you can protect Google Cloud resources in the following ways:

- **Create cloud-native snapshots of VM instances**

A cloud-native snapshot includes point-in-time snapshots of persistent disks attached to the processed VM instance. Snapshots of persistent disks (also referred to as PD snapshots) are taken using [native Google Cloud capabilities](#). By default, cloud-native snapshots are stored in the multi-regional location closest to the region in which the original instance resides, but the location can be changed in the [backup policy settings](#).

- **Create image-level backups of VM instances**

In addition to cloud-native snapshots, you can protect your VM instances with image-level backups. An image-level backup captures the whole image of the processed VM instance (including OS data, application data and so on) at a specific point in time. The backup is saved as multiple files to a storage bucket in the [native Veeam format](#).

- **Create cloud-native snapshots of Cloud SQL instances**

A cloud-native snapshot is a point-in-time snapshot of the processed Cloud SQL instance. Snapshots of Cloud SQL instances are taken using [native Google Cloud capabilities](#). Cloud-native snapshots are stored in the multi-regional location closest to the region in which the original instance resides.

NOTE

Cloud-native snapshots of Cloud SQL instances are referred to as backups in Google Cloud documentation. However, since all 'backups' of a Cloud SQL instance are automatically deleted after you remove the instance itself, 'backups' of Cloud SQL instances are referred to as snapshots in this guide. In terms of Veeam logic, backups are independent files that are stored in backup repositories and that are not affected by any actions performed with the original instances whatsoever.

- **Create image-level backups of Cloud SQL instances**

In addition to cloud-native snapshots, you can protect your Cloud SQL instances with image-level backups. An image-level backup captures the whole image of the processed Cloud SQL instance (including the instance configuration, databases, triggers, stored procedures and users) at a specific point in time. The backup is saved as multiple files to a storage bucket in the [native Veeam format](#).

NOTE

Veeam Backup for Google Cloud allows you to protect MySQL and PostgreSQL instances. SQL Server instances are not supported. For more information on types of Cloud SQL instances, see [Google Cloud documentation](#).

- **Create cloud-native snapshots of Cloud Spanner instances**

A cloud-native snapshot is a point-in-time snapshot of the processed Cloud Spanner instance. Snapshots of Cloud Spanner instances are taken using [native Google Cloud capabilities](#). Cloud-native snapshots are stored in the location that depends on the [regional configuration](#) of the processed instance.

- **Create image-level backups of Cloud Spanner instances**

In addition to cloud-native snapshots, you can protect your Cloud Spanner instances with image-level backups. An image-level backup captures the whole image of the processed Cloud Spanner instance (including databases schema, data, views, foreign keys) at a specific point in time. The backup is saved as multiple files to a storage bucket in the [native Veeam format](#).

To schedule data protection tasks to run automatically, create backup policies. For VM, Cloud SQL and Cloud Spanner instances residing in any of the regions added to the backup policies, you can also take cloud-native snapshots manually when needed – for more information, see [Creating VM Snapshots Manually](#), [Creating SQL Snapshots Manually](#) and [Creating Spanner Snapshots Manually](#).

TIP

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

Performing Backup Using Console

To produce cloud-native snapshots and image-level backups of VM, Cloud SQL and Cloud Spanner instances, Veeam Backup for Google Cloud 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 backups must be stored, when the backup process must start, and so on.

One backup policy can be used to process multiple VM, Cloud SQL or Cloud Spanner instances within different regions, but you can back up each VM, Cloud SQL or Cloud Spanner instance with one backup policy at a time. 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 Backup Policy Priority](#).


After you install Google Cloud 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 Google Cloud 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 **GCP > GCP VM, GCP SQL, GCP PostgreSQL** or **GCP Spanner**, and select the backup appliance on which you want to create the backup policy.
- Open the **Home** view, right-click **Jobs**, navigate to **Backup > GCP > GCP VM, GCP SQL, GCP PostgreSQL** or **GCP Spanner**, and select the backup appliance on which you want to create the backup policy.

Veeam Backup & Replication will open the **Add VM Policy, Add Cloud SQL Policy** or **Add Cloud Spanner Policy** wizard in a web browser. Complete the wizard as described in section [Creating VM Backup Policies](#), [Creating Cloud SQL Backup Policies](#) or [Creating Cloud Spanner Backup Policies](#).



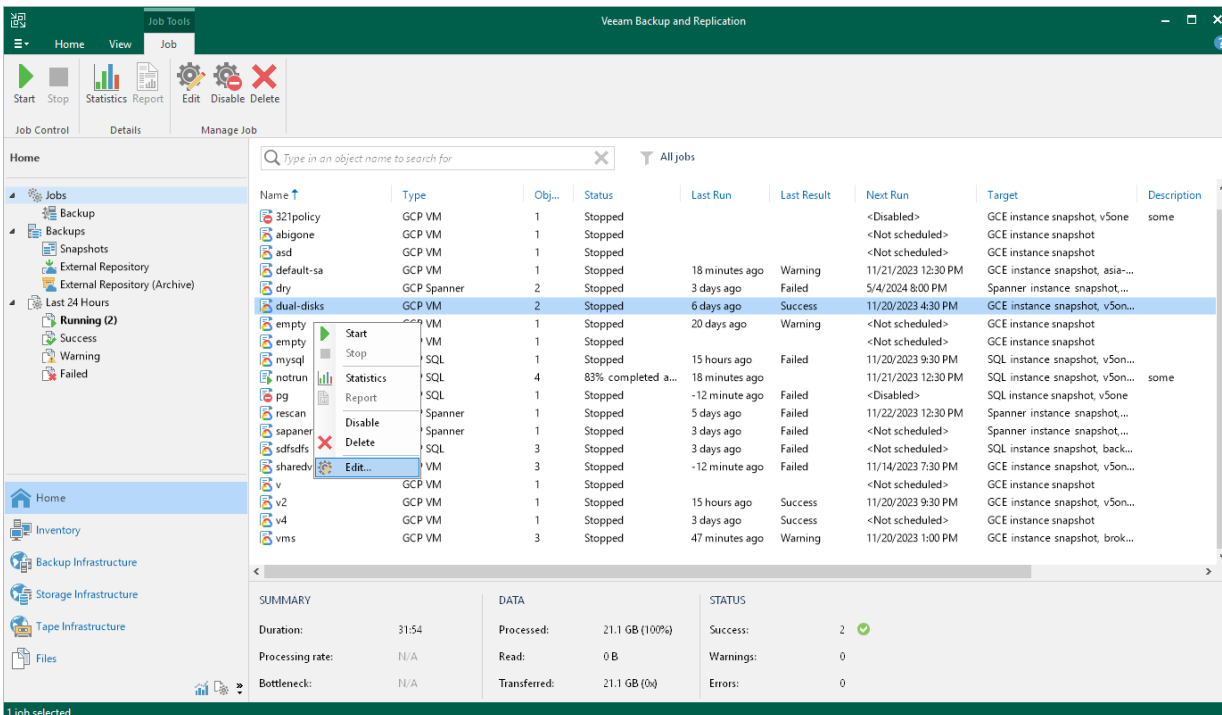
Editing Backup Policy Settings

You can edit backup policies in the Veeam Backup for Google Cloud Web UI only. However, you can launch the edit policy wizard directly from the Veeam Backup & Replication console:

1. In the Veeam Backup & Replication console, open the **Home** view.
2. Navigate to **Jobs**.
3. Select the necessary 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 [Editing Backup Policy Settings](#).



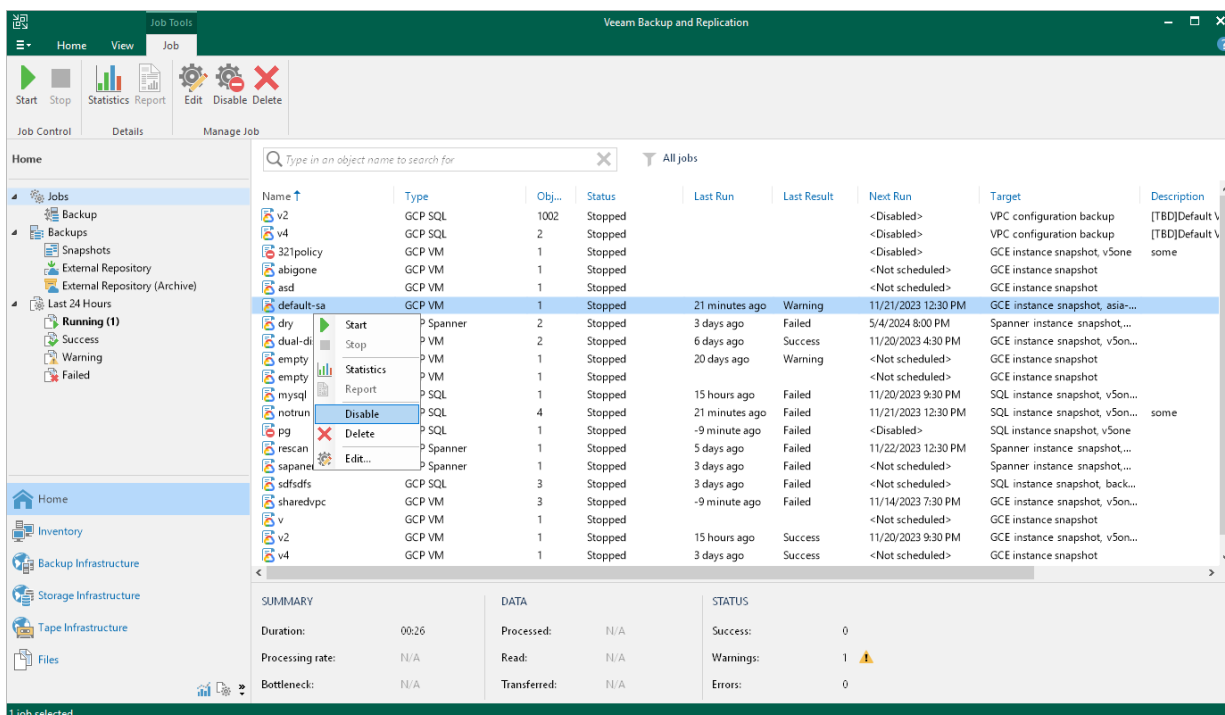
Enabling and Disabling Backup Policies

By default, Veeam Backup for Google Cloud runs all created backup policies according to the specified schedules. However, you can temporarily disable a backup policy so that Veeam Backup for Google Cloud 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**.



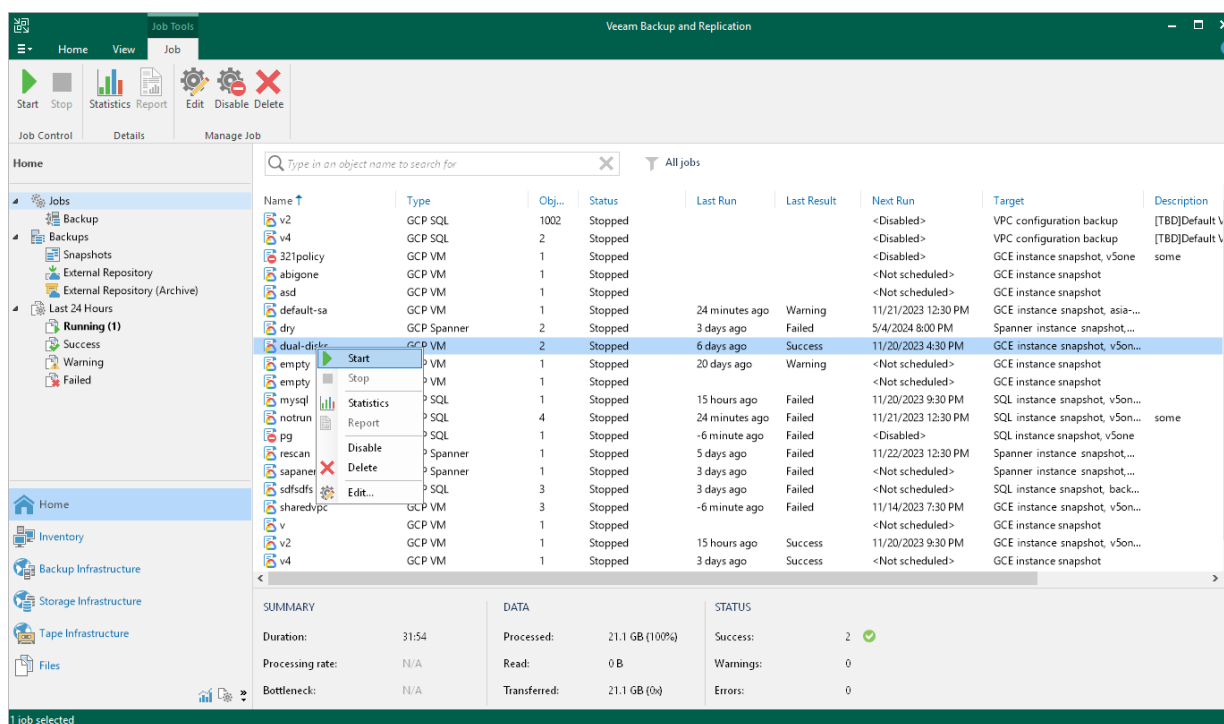
Starting and Stopping Backup Policies

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

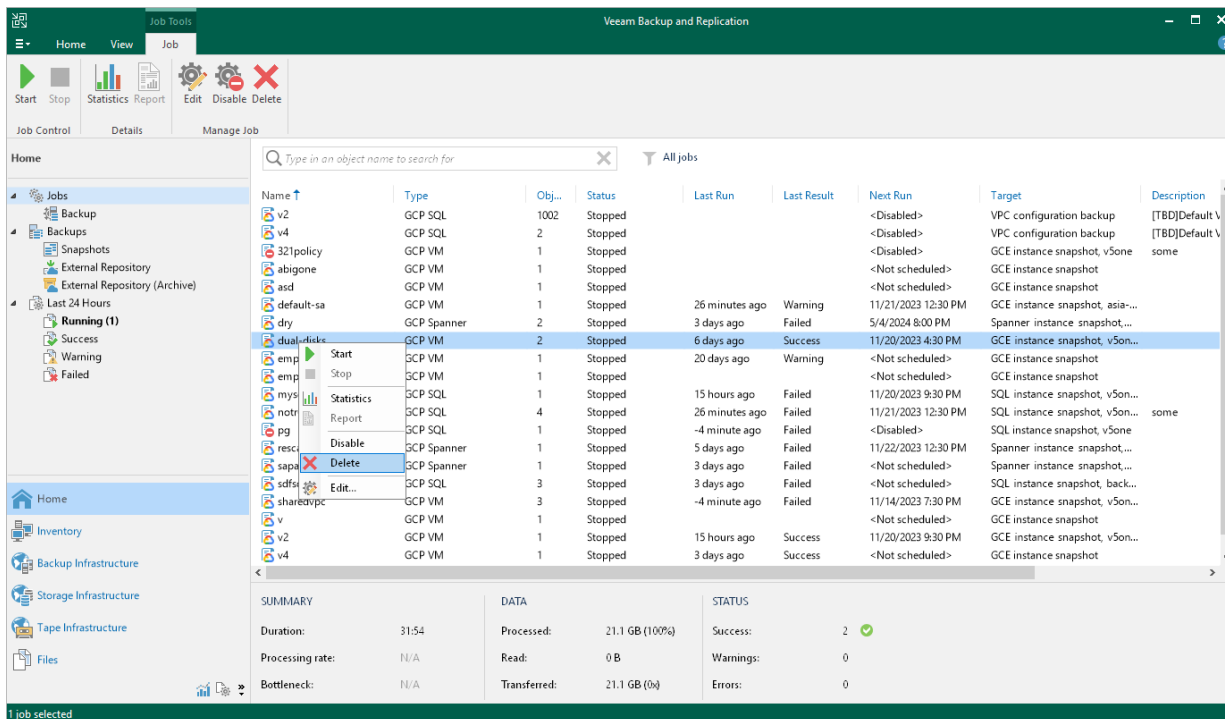
Veeam Backup & Replication allows you to permanently delete backup policies created by Veeam Backup for Google Cloud:

1. In the Veeam Backup & Replication console, open the **Home** view.
2. Navigate to **Jobs**.
3. Select the necessary backup policy and click **Delete** on the ribbon.

Alternatively, you can right-click the necessary backup policy and select **Delete**.

IMPORTANT

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



Creating Backup Copy Jobs

Backup copy is a technology that helps you copy and store backed-up data of VM instances in different locations. Storing data in different locations increases its availability and ensures that data can be recovered in case a disaster strikes.

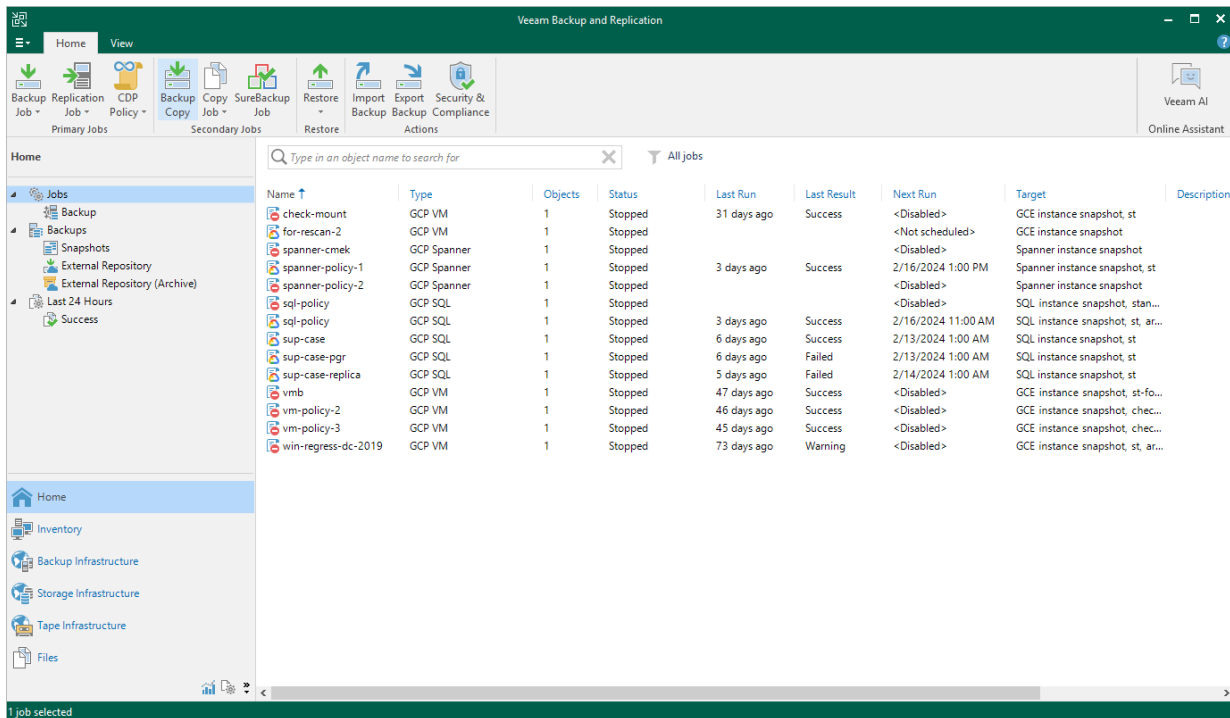
Backup copy is a job-driven process. Veeam Backup & Replication fully automates the backup copy process and lets you specify retention settings to maintain the desired number of restore points, as well as full backups for archival purposes. For more information on the backup copy functionality, see the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

IMPORTANT

Backup copy can be performed only using backup files stored in backup repositories for which you have specified HMAC keys associated with the service accounts that are used to access the repositories. To learn how to specify credentials for repositories, see sections [Creating New Repositories](#) and [Connecting to Existing Repositories](#).

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](#).



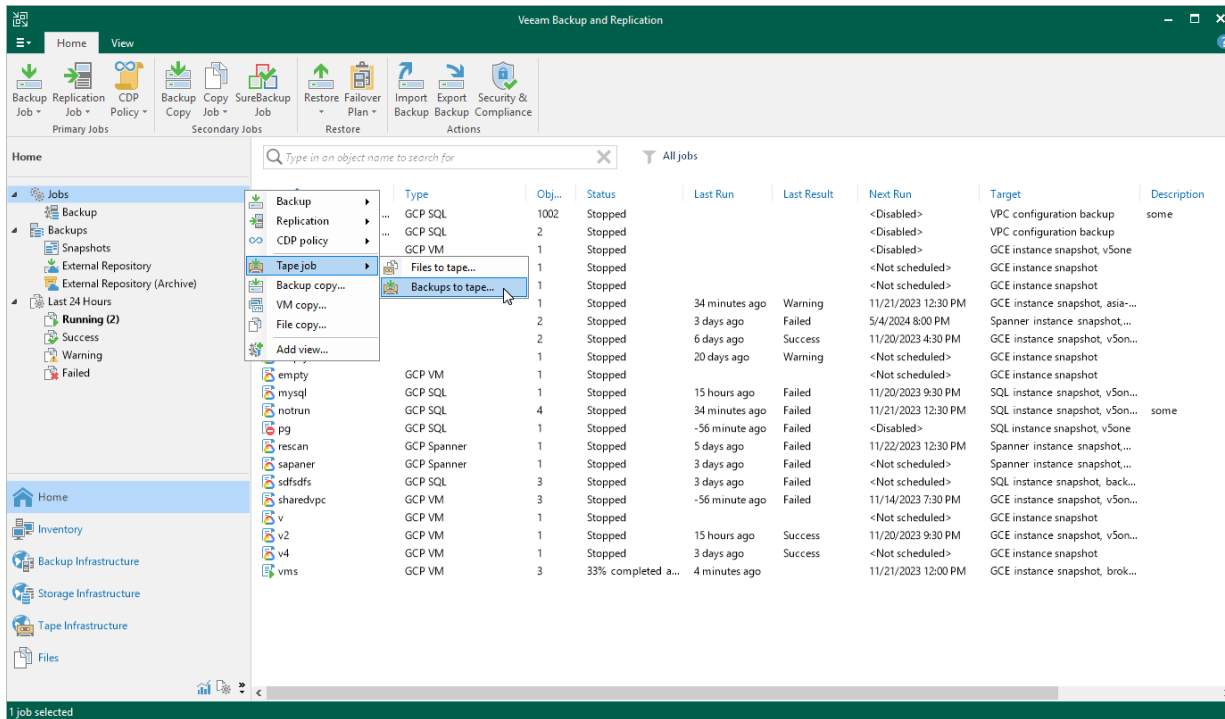
Copying Backups to Tapes

Veeam Backup & Replication allows you to automate copying of image-level backups of VM instances to tape devices and lets you specify scheduling, archiving and media automation options. For more information on the supported tape libraries, see the Veeam Backup & Replication User Guide, section [Tape Devices Support](#).

Before you start copying backup to tapes:

- Copy VM instance backups to on-premises backup repositories as described in section [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 VM 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

To produce cloud-native snapshots and image-level backups of VM, Cloud SQL and Cloud Spanner instances, Veeam Backup for Google Cloud 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, and so on.

One backup policy can be used to process multiple instances within different regions, but you can back up each instance with one backup policy at a time. 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 Backup Policy Priority](#).

Performing VM Backup

One backup policy can be used to process one or more VM instances within one Google Cloud project or folder. The scope of data that you can protect in a project or folder is limited by permissions of a service account that is specified in the backup policy settings.

Before you create a VM backup policy, check the following prerequisites:

- If you plan to create image-level backups of VM 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 the backup policy results, configure SMTP server settings first. For more information, see [Configuring Global Notification Settings](#).

To schedule data protection tasks to run automatically, [create backup policies](#). For each protected VM instance, you can also [take a cloud-native snapshot manually](#) when needed.

Creating Backup Policies

To create a backup policy, do the following:

1. [Launch the Add VM Policy wizard](#).
2. [Specify a backup policy name and description](#).
3. [Choose a project to which VM instances that you plan to back up belong](#).
4. [Configure backup source settings](#).
5. [Configure backup target settings](#).
6. [Create a schedule for the backup policy](#).
7. [Enable label assignment](#).
8. [Specify automatic retry, health check and notification settings for the backup policy](#).
9. [Review the estimated cost of protecting the selected VM instances](#).
10. [Check the required permissions](#).
11. [Finish working with the wizard](#).

Step 1. Launch Add VM Policy Wizard

To launch the **Add VM Policy** wizard, do the following:

1. Navigate to **Policies > VM**.
2. Click **Add**.

The screenshot shows the Veeam Backup for Google Cloud interface. The top navigation bar includes the Veeam logo, the text 'Veeam Backup for Google Cloud', the server time 'Nov 13, 2023 3:55 PM', and the user 'administrator Portal Administrator'. The left sidebar shows a navigation menu with 'Policies' selected. The main content area is titled 'VM' and contains a table of policies. The 'Add' button is highlighted with a mouse cursor. Below the table are sections for 'Instances' and 'Sessions', both showing 'No data'.

Priority	Policy	Snapshots	Backups	Archives	Last Run	Last Duration	Next Run	State	
<input type="checkbox"/>	1	000000	Success	Success	Success	11/13/2023 2:00:30 AM	19 min 59 sec	11/14/2023 12:00:00 AM	Enabled
<input type="checkbox"/>	2	vms	Success	Success	Success	11/09/2023 8:01:58 PM	22 min 58 sec	n/a	Disabled
<input type="checkbox"/>	3	empty	Warning	Not Configured	Not Configured	11/06/2023 11:53:16 PM	2 sec	n/a	Enabled
<input type="checkbox"/>	4	abigone	Failed	Not Configured	Not Configured	11/09/2023 9:02:49 PM	4 min 23 sec	n/a	Enabled

Step 2. Specify Backup Policy Name and Description

At the **Policy Info** step of the wizard, use the **Name** and **Description** fields to enter a name for the new backup policy and to provide a description for future reference. The policy name can contain only uppercase Latin letters, lowercase Latin letters, numeric characters and hyphens; the maximum length of the name is 127 characters.

NOTE

You can tell snapshots created by Veeam Backup for Google Cloud from other snapshots in your infrastructure by their names – the name of every snapshot created by a backup policy will contain the word *veeam*, a GUID and the ordinary number of the processed persistent disk: `veeam- {GUID} - {disk number}`.

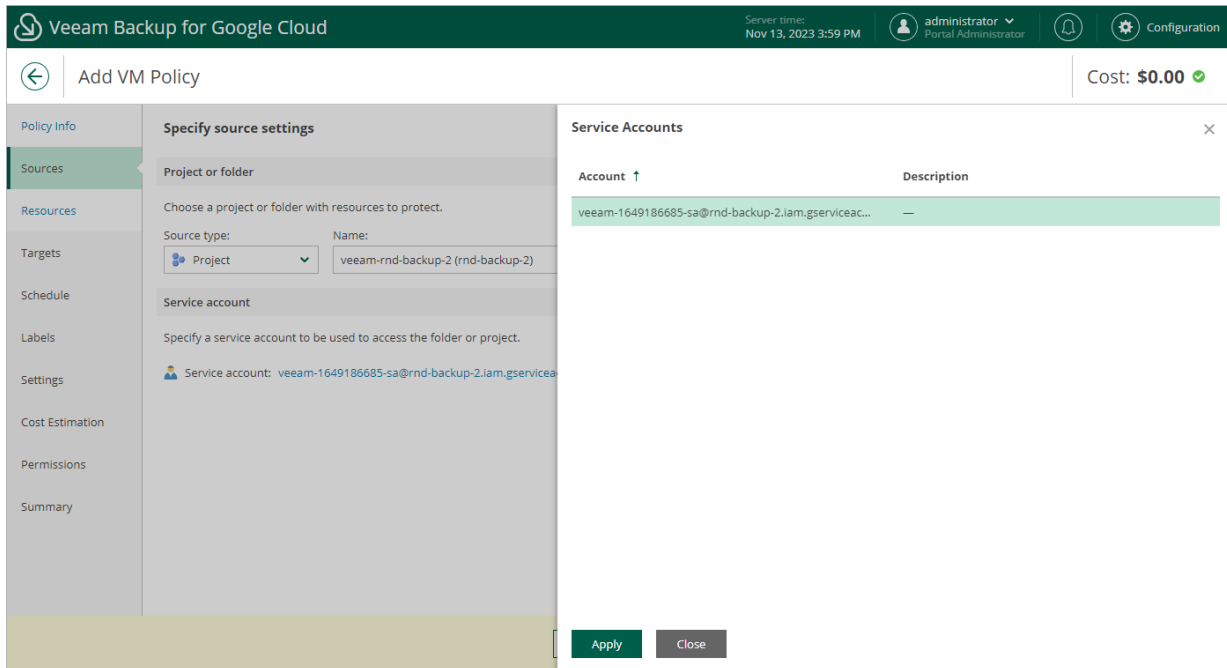
The screenshot shows the 'Add VM Policy' wizard in the Veeam Backup for Google Cloud interface. The top navigation bar includes the Veeam logo, the product name, the server time (Nov 13, 2023 3:58 PM), the user role (administrator Portal Administrator), and a Configuration icon. The main header shows a back arrow, the title 'Add VM Policy', and a cost indicator 'Cost: \$0.00' with a green checkmark. A left sidebar lists the wizard steps: Policy Info (selected), Sources, Resources, Targets, Schedule, Labels, Settings, Cost Estimation, Permissions, 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 'us-west-policy' and 'Description:' with the value 'protecting instances in us-west regions'. At the bottom of the wizard, there are 'Next' and 'Cancel' buttons.

Step 3. Specify Project

At the **Sources** step of the wizard, choose a project or a folder with a project that manages resources that you want to protect, and specify a service account that will be used to access the project or folder.

For a project or folder to be displayed in the list of available entities, it must be added to Veeam Backup for Google Cloud as described in section [Adding Projects and Folders](#). If you have not added the necessary entity to Veeam Backup for Google Cloud beforehand, you can do it without closing the **Add VM Policy** wizard. To add an entity, click **Add** and complete the **Add Projects and Folders** wizard.

For a service account to be displayed in the list of available accounts, it must be added to Veeam Backup for Google Cloud as described in section [Adding Service Accounts](#), and must be assigned the *VM Instances Snapshot* and *Backup* operational roles as described in section [Adding Projects and Folders](#).



Step 4. Configure Backup Source Settings

At the **Resources** step of the wizard, specify the following backup source settings:

1. [Choose regions in which VM instances that you plan to back up reside.](#)
2. [Select VM instances to back up.](#)

Step 4a. Choose Regions

In the **Region** section of the **Resources** step of the wizard, choose regions in which VM instances that you want to protect reside.

1. Click **Choose regions**.
2. In the **Choose regions** window, select the necessary regions, click **Add** to include them in the backup policy, and then click **Apply**.

The screenshot shows the Veeam Backup for Google Cloud interface during the 'Add VM Policy' wizard. The top navigation bar includes the Veeam logo, the product name, server time (Nov 13, 2023 4:00 PM), user information (administrator), and configuration options. The main area is titled 'Add VM Policy' and shows a 'Cost: \$0.00' indicator. A sidebar on the left contains navigation links: Policy Info, Sources, Resources (selected), Targets, Schedule, Labels, Settings, Cost Estimation, Permissions, and Summary. The 'Specify resources to protect' section is active, showing the 'Region' step with options to 'Choose regions...' and 'Exclude resources...'. The 'Choose regions' dialog box is open, displaying a list of 'Available Regions (35)' and a list of 'Policy Regions (4)'. The 'Available Regions' list includes: europe-west10 (Berlin), europe-west12 (Turin), me-central1 (Doha), me-central2 (Dammam), me-west1 (Tel Aviv), northamerica-northeast1 (Montréal), northamerica-northeast2 (Toronto), southamerica-east1 (São Paulo), southamerica-west1 (Santiago), us-central1 (Iowa), us-east1 (South Carolina), us-east4 (Northern Virginia), us-east5 (Columbus), and us-south1 (Dallas). The 'Policy Regions' list includes: us-west1 (Oregon), us-west2 (Los Angeles), us-west3 (Salt Lake City), and us-west4 (Las Vegas). 'Add' and 'Remove' buttons are visible between the lists, and 'Apply' and 'Cancel' buttons are at the bottom of the dialog.

Step 4b. Select VM Instances

In the **Resources** section of the **Resources** step of the wizard, specify the backup scope — select VM instances that Veeam Backup for Google Cloud will back up:

1. Click **Choose resources to protect**.
2. In the **Choose resources** window, choose whether you want to back up all VM instances from the regions selected at [step 4a](#), or only specific VM instances.

If you select the **All resources** option, Veeam Backup for Google Cloud will regularly check for new VM instances launched in the selected regions and automatically update the backup policy settings to include these instances in the backup scope.

If you select the **Specific resources** option, you must also specify the instances explicitly:

- a. Use the **Resource type** drop-down list to choose whether you want to add individual VM instances or Google Cloud labels to the backup scope.

If you select the **Label** option, Veeam Backup for Google Cloud will back up only those VM instances that reside in the selected regions under specific labels.

- b. Use the **Instance\Label** list to find the necessary resource, and then click **Add to Protected** to add the resource to the backup scope.

For a resource to be displayed in the list of available resources, it must reside in a region that has ever been specified in any backup policy. Otherwise, the only option to discover available resources is to click **Browse** and wait for Veeam Backup for Google Cloud to populate the resource list.

TIP

You can simultaneously add multiple resources to the backup scope. To do that, click **Browse**, select check boxes next to the necessary VM instances or labels 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 Google Cloud will update the resource list.

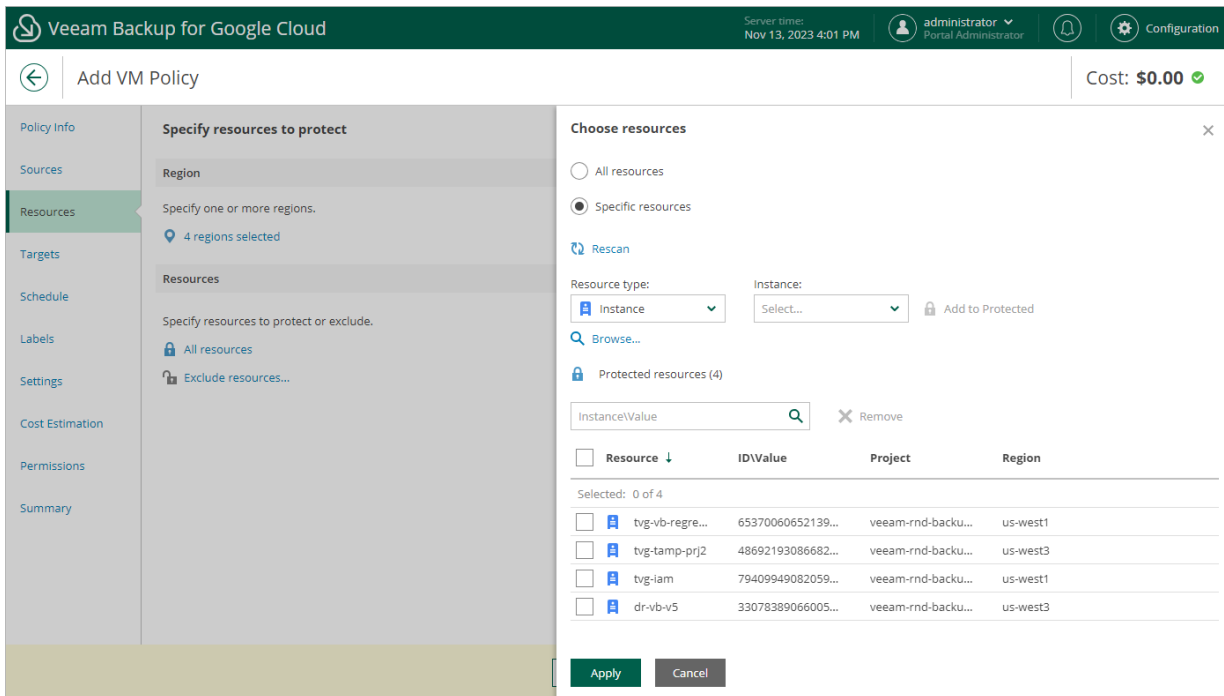
If you add a label to the backup scope, Veeam Backup for Google Cloud will regularly check for new VM instances assigned the added label and automatically update the backup policy settings to include these instances in the scope. However, this applies only to VM instances from the regions selected at [step 4a](#). If you select a label assigned to VM instances from other regions, these instances will not be protected by the backup policy. To work around the issue, either go back to step 4a 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 **Specific 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 **Exclude resources** and specify the VM instances or labels that you do not want to back up – the procedure is the same as described for including resources in the backup scope.

Consider that if a resource appears both in the list of included and excluded resources, Veeam Backup for Google Cloud will still not process the resource because the list of excluded resources has a higher priority.



The screenshot shows the Veeam Backup for Google Cloud interface. The main window is titled "Add VM Policy" and has a "Cost: \$0.00" indicator. The left sidebar contains navigation options: Policy Info, Sources, Resources (selected), Targets, Schedule, Labels, Settings, Cost Estimation, Permissions, and Summary. The "Specify resources to protect" section is active, showing "Region" and "Resources" options. The "Resources" section is expanded, showing "Specify resources to protect or exclude." with options for "All resources" and "Exclude resources...".

The "Choose resources" dialog is open, showing the "Specific resources" option selected. The "Resource type" is set to "Instance". The "Instance" dropdown is set to "Select...". The "Add to Protected" button is visible. Below the dialog, there is a search bar for "InstanceValue" and a "Remove" button. A table lists four resources:

<input type="checkbox"/>	Resource ↓	IDValue	Project	Region
<input type="checkbox"/>	tv-g-vb-regre...	65370060652139...	veeam-rnd-backu...	us-west1
<input type="checkbox"/>	tv-g-tamp-prj2	48692193086682...	veeam-rnd-backu...	us-west3
<input type="checkbox"/>	tv-g-iam	79409949082059...	veeam-rnd-backu...	us-west1
<input type="checkbox"/>	dr-vb-v5	33078389066005...	veeam-rnd-backu...	us-west3

Selected: 0 of 4

Buttons: Apply, Cancel

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 specify a location for the created snapshots and enable additional data protection scenarios.

Configuring Snapshot Settings

To choose whether you want to store the created cloud-native snapshots in multi-regional locations closest to the regions of the source VMs or in the same regions where the source VMs reside, select either the **Multi-regional** or **Regional** option in the **Snapshot settings** section.

Configuring Backup Settings

To instruct Veeam Backup for Google Cloud to create image-level backups of the selected VM instances, do the following:

1. In the **Backup settings** section, set the **Enable backups** toggle to *On*.
2. Click **Choose repository**.
3. In the **Choose repository** window, select a backup repository where the created image-level backups will be stored.

For a backup repository to be displayed in the **Repository** list, it must be added to Veeam Backup for Google Cloud as described in section [Adding Backup Repositories](#). The **Repository** list shows only backup repositories of the *Standard* and *Nearline* storage classes.

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

Configuring Archive Settings

To instruct Veeam Backup for Google Cloud to store backed-up data in a low-cost, long-term archive storage, do the following:

1. In the **Backup settings** section, select the **Enable backup archiving** check box.
2. Click **Choose repository**.
3. In the **Choose repository** window, select a backup repository where the archived data will be stored.

For a backup repository to be displayed in the **Repository** list, it must be added to Veeam Backup for Google Cloud as described in section [Adding Backup Repositories](#). The **Repository** list shows only backup repositories of the *Archive* storage class.

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

For more information on the backup archiving mechanism, see [Enabling Backup Archiving](#).

Specify target settings

Snapshot settings

Choose a location to store snapshots created by the policy.

Multi-regional
Snapshots will be stored in multiple regions

Regional
Snapshots will be stored in one region (same as the source VM location)

Backup settings

Choose a location to store backups created by the policy.

Enable backups: On

Backups will be stored in: Choose repository...

Enable backup archiving: Choose repository...

Choose repository for backups

Repository Refresh

Repository ↑	Folder	Storage Class	Description
backup-readd	std-encryped	Standard	
custom-buildrepo	custom-backup	Standard	
from-v4	fromv4-tov5	Nearline	

Apply **Cancel**

Step 6. Specify Policy Scheduling Options

At the **Schedule** step of the wizard, you can instruct Veeam Backup for Google Cloud to start the backup policy automatically according to a specific backup schedule. The backup schedule defines how often data of the VM instances added to the backup policy will be backed up.

To help you implement a comprehensive backup strategy, Veeam Backup for Google Cloud 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](#).

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** section, select hours when the backup policy will create cloud-native snapshots and image-level backups. Use the **Run at** drop-down list to choose whether you want the backup policy to run every day, on weekdays (Monday through Friday) or on specific days.

If you want to protect VM 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

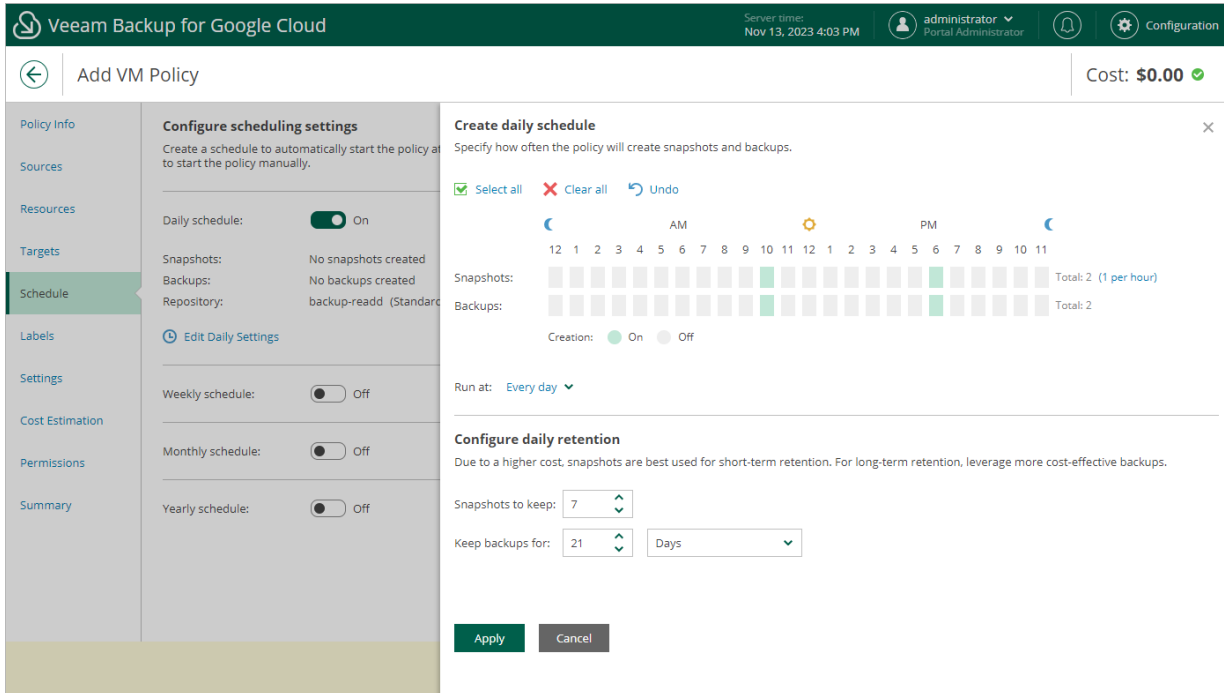
Veeam Backup for Google Cloud does not create image-level backups independently from cloud-native snapshots. That is why when you select hours for image-level backups, the same hours are automatically selected for cloud-native snapshots. To learn how Veeam Backup for Google Cloud performs backup, see [VM Backup](#).

3. In the **Configure daily retention** section, configure retention policy settings for the daily schedule:
 - For cloud-native snapshots, specify the number of restore points that you want to keep in a snapshot chain.

If the restore point limit is exceeded, Veeam Backup for Google Cloud removes the earliest restore point from the chain. For more information, see [Retention Policy for Snapshots](#).
 - 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 Google Cloud removes the restore point from the chain. For more information, see [Retention Policy for Backups](#).

4. 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** section, select days when the backup policy will create cloud-native snapshots and image-level backups. Use the **Create restore points** at drop-down list to schedule a specific time for the backup policy to run.

NOTE

Veeam Backup for Google Cloud does not create image-level backups independently from cloud-native snapshots. That is why when you select days for image-level backups, the same days are automatically selected for cloud-native snapshots. To learn how Veeam Backup for Google Cloud performs backup, see [VM Backup](#).

3. In the **Configure weekly retention** section, configure retention policy settings for the weekly schedule:
 - For cloud-native snapshots, specify the number of restore points that you want to keep in a snapshot chain.
If the restore point limit is exceeded, Veeam Backup for Google Cloud removes the earliest restore point from the chain. For more information, see [Retention Policy for Snapshots](#).
 - 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 Google Cloud removes the restore point from the chain. For more information, see [Retention Policy for Backups](#).

4. 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 **Choose monthly backup target** section, select months when the backup policy will create cloud-native snapshots and image-level backups. Use the **Create restore points at** and **Run on** drop-down lists to schedule a specific time and day for the backup policy to run.

NOTE

Veeam Backup for Google Cloud does not create image-level backups independently from cloud-native snapshots. That is why when you select months for image-level backups, the same months are automatically selected for cloud-native snapshots. To learn how Veeam Backup for Google Cloud performs backup, see [VM Backup](#).

3. In the **Configure monthly retention** section, configure retention policy settings for the monthly schedule:
 - For cloud-native snapshots, specify the number of restore points that you want to keep in a snapshot chain.

If the restore point limit is exceeded, Veeam Backup for Google Cloud removes the earliest restore point from the chain. For more information, see [Retention Policy for Snapshots](#).
 - 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 Google Cloud removes the restore point from the chain. For more information, see [Retention Policy for Backups](#).
4. To save changes made to the backup policy settings, click **Apply**.

TIP

If you have enabled backup archiving at the **Targets** step of the wizard, and want to store monthly backups in an archive backup repository, set the **Send backups to archive** toggle to *On*, and follow the instructions provided in section [Enabling Backup Archiving](#).

The screenshot displays the 'Add VM Policy' wizard in the Veeam Backup for Google Cloud interface. The 'Schedule' step is active, showing settings for both daily and monthly schedules. The monthly schedule is currently set to 'On' with a restore point at 10:00 AM. A dialog box titled 'Choose monthly backup target' is open, allowing the user to select a target archive repository for monthly backups. The dialog includes a calendar grid for selecting a month and year, and options to configure monthly retention (Snapshots to keep: 7, Keep backups for: 30 Days). The top of the wizard shows the cost of the policy as \$17.06.

Specifying Yearly Schedule

[This step applies only if you have instructed Veeam Backup for Google Cloud 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. In the **Create 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.

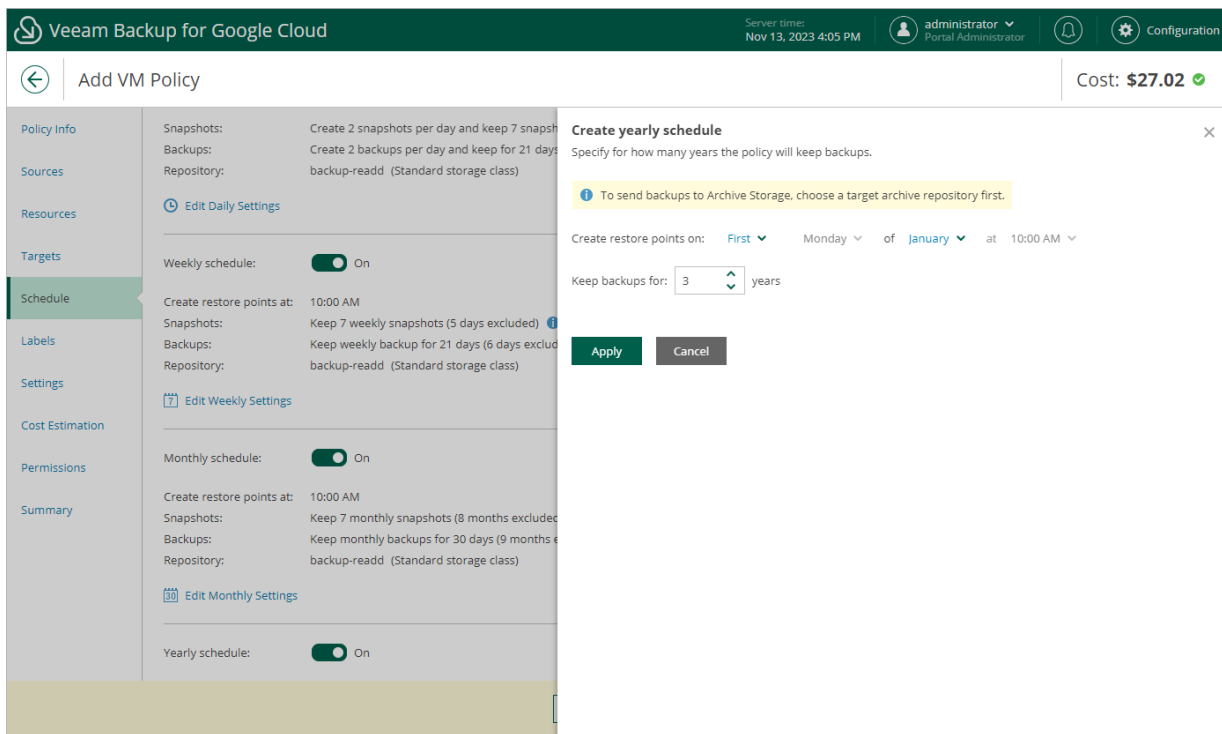
3. In the **Keep backups for** field, specify the number of years for which you want to keep restore points in a backup chain.

If a restore point is older than the specified time limit, Veeam Backup for Google Cloud removes the restore point from the chain. For more information, see [Retention Policy for Backups](#).

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

TIP

If you have enabled backup archiving at the **Targets** step of the wizard, and want to store yearly backups in an archive backup repository, set the **Send backups to archive** toggle to *On*, and follow the instructions provided in section [Enabling Backup Archiving](#).



Enabling Harmonized Scheduling

When you combine multiple types of schedules, Veeam Backup for Google Cloud 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 Google Cloud can keep restore points created according to a daily, weekly or monthly schedule for longer periods of time:

- Cloud-native snapshots can be kept for weeks and months.
- Image-level backups can be kept for weeks, months and years.

For Veeam Backup for Google Cloud 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 retaining restore points. In terms of harmonized scheduling, Veeam Backup for Google Cloud 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 (Daily) flag is used to mark restore points created daily, (Weekly) – weekly, (Monthly) – monthly, and (Yearly) – yearly. Veeam Backup for Google Cloud 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 by retention – 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.

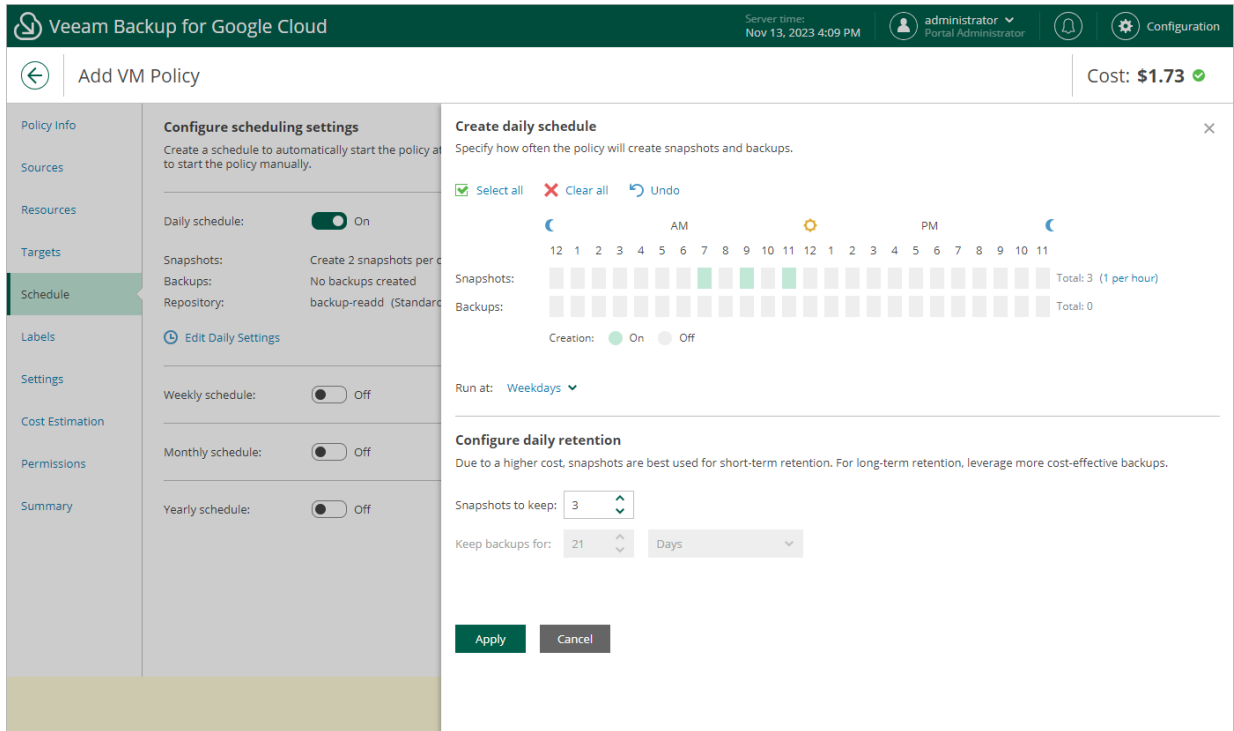
NOTE

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 Google Cloud 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 retain one of the created snapshots for 2 weeks. In this case, you create 2 schedules when configuring the backup policy settings – daily and weekly:

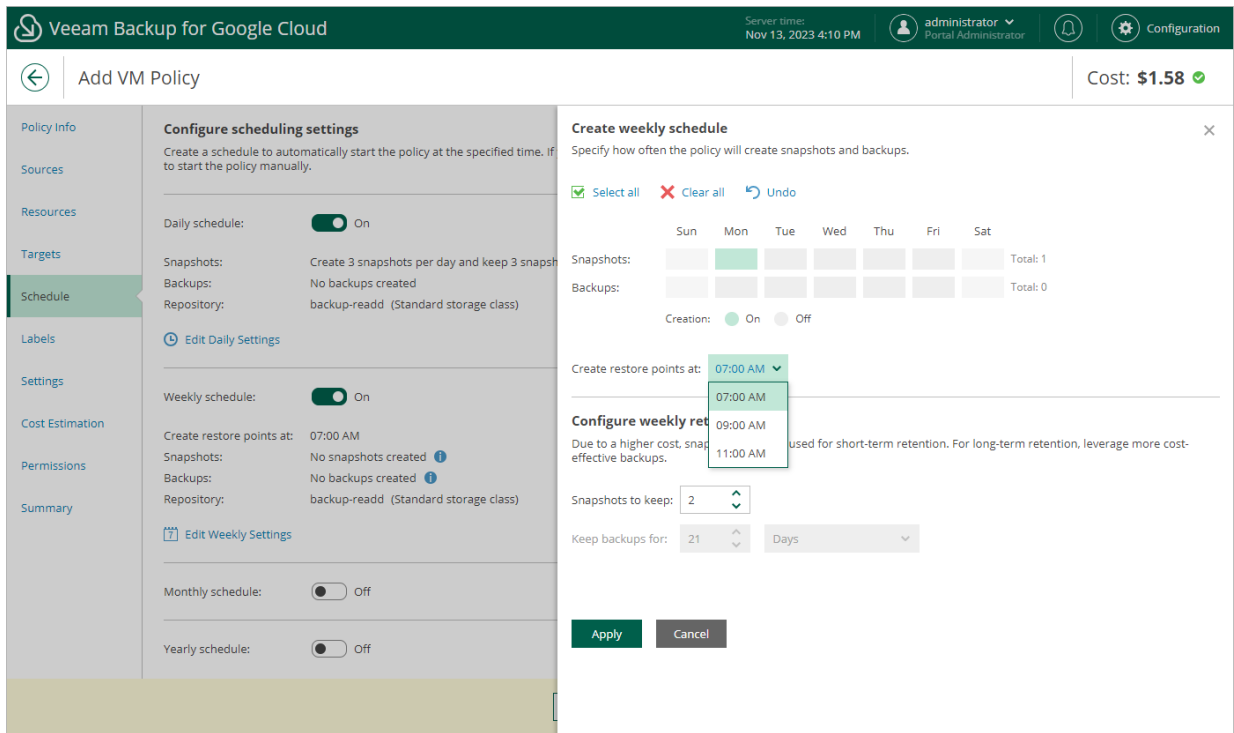
1. In the daily scheduling settings, you select hours and days when snapshots will be created (for example, *7:00 AM, 9:00 AM, and 11:00 AM; Weekdays*), and specify a number of daily restore points to retain (for example, *3*).

Veeam Backup for Google Cloud will propagate these settings to the schedule with 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 kept 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 Google Cloud will create cloud-native snapshots in the following way:

- On the first weekday (Monday), a backup session will start at 7:00 AM to create the first restore point. The restore point will be marked with the (Daily) flag as it was created according to the daily schedule.


Since *7:00 AM, Monday* is specified in the weekly scheduling settings, Veeam Backup for Google Cloud will assign the (Weekly) flag to this 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 (Daily) flag.




- On the next weekday (Tuesday), after a backup session runs at 7:00 AM, the created restore point will be marked with the (Daily) flag.


By the moment the backup session completes, the number of restore points with the (Daily) flag will exceed the retention limit specified in the daily scheduling settings. However, Veeam Backup for Google Cloud will not remove the earliest restore point (7:00 AM, Monday) with the (Daily) flag from the snapshot chain as this restore point is also marked with a flag of a less-frequent schedule. Instead, Veeam Backup for Google Cloud will unassign the (Daily) 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 (Daily) flag will exceed the retention limit once again. Veeam Backup for Google Cloud 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 Google Cloud 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 Google Cloud will unassign the (Weekly) flag from the earliest weekly restore point. Since no other flags are assigned to this restore point, Veeam Backup for Google Cloud 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 Google Cloud to store backed-up data in the low-cost, long-term Google Cloud archival storage. 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 Google Cloud standard storage.


With backup archiving, Veeam Backup for Google Cloud can retain backups 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 backup schedule (or both).
- To enable yearly archiving, you must configure a daily, a weekly or a monthly backup schedule (or all three).

For Veeam Backup for Google Cloud to use the archiving mechanism, there must be specified at least 2 different schedules: one schedule will control the regular creation of backups, while another schedule will control the process of copying backups to an archive repository. Backup chains created according to these two schedules will be completely different – for more information, see [Backup Chain](#) and [Archive Backup Chain](#).

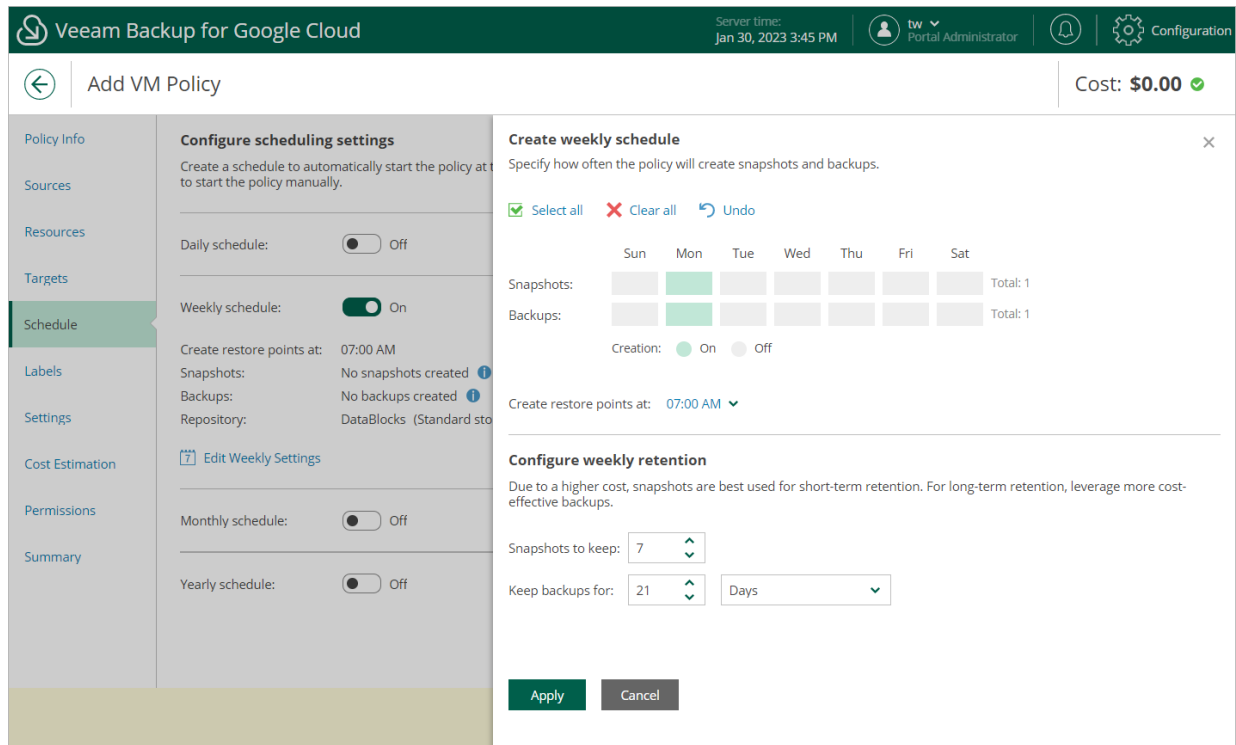
Consider the following example. You want a backup policy to create image-level backups of your critical workloads once a week, to keep the backed-up data in a standard repository for 3 weeks, and also to keep backups created once in 2 months in an archive repository for a year. In this case, you create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

1. In the policy target settings, you set the **Enable backups** toggle to *On*, select a standard repository that will store regular backups, and select an archive 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 Google Cloud will retain backups (for example, *21 days*).

Veeam Backup for Google Cloud will propagate these settings to the archive schedule (which is the monthly schedule in our example).

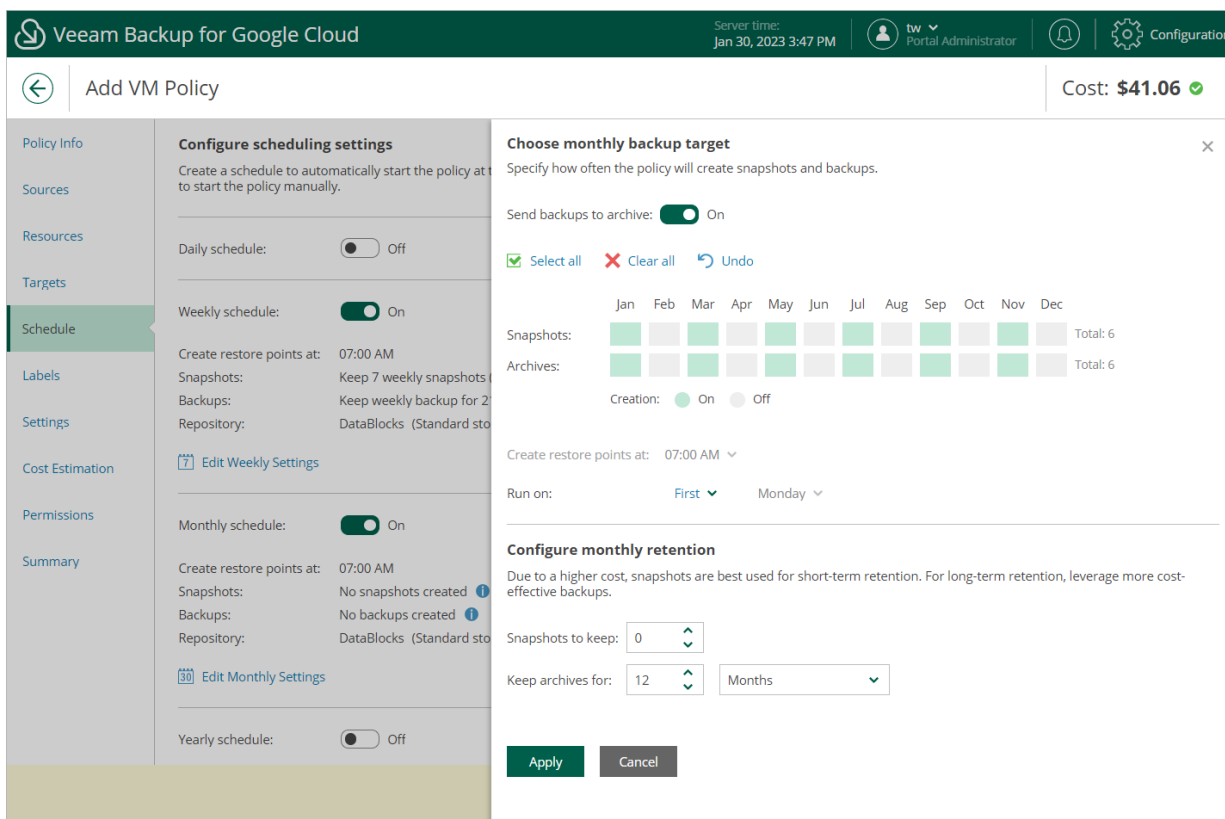


- In the monthly scheduling settings, you enable the archiving mechanism by setting the **Send backups to archive** toggle to *On*, specify when Veeam Backup for Google Cloud will create archive backups, and choose for how long you want to keep the created backups in the archive repository.

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


IMPORTANT

- When you enable backup archiving, you become no longer able to create a schedule of the same frequency for regular 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 *12 months (or 365 days)*, since the minimum storage duration of the Google Cloud archival storage is 365 days.
- If you select the **On day** option, **harmonized scheduling** cannot be guaranteed. Plus, to support the **On day** option, Veeam Backup for Google Cloud 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 Google Cloud Storage during approximately 24 hours, to reduce unexpected infrastructure charges.




According to the specified scheduling settings, Veeam Backup for Google Cloud 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 regular backup chain. Veeam Backup for Google Cloud will store this restore point as a full backup in the standard repository.
2. On the second and third Mondays of February, Veeam Backup for Google Cloud will create restore points at 7:00 AM and add them to the regular backup chain as incremental backups in the standard repository.



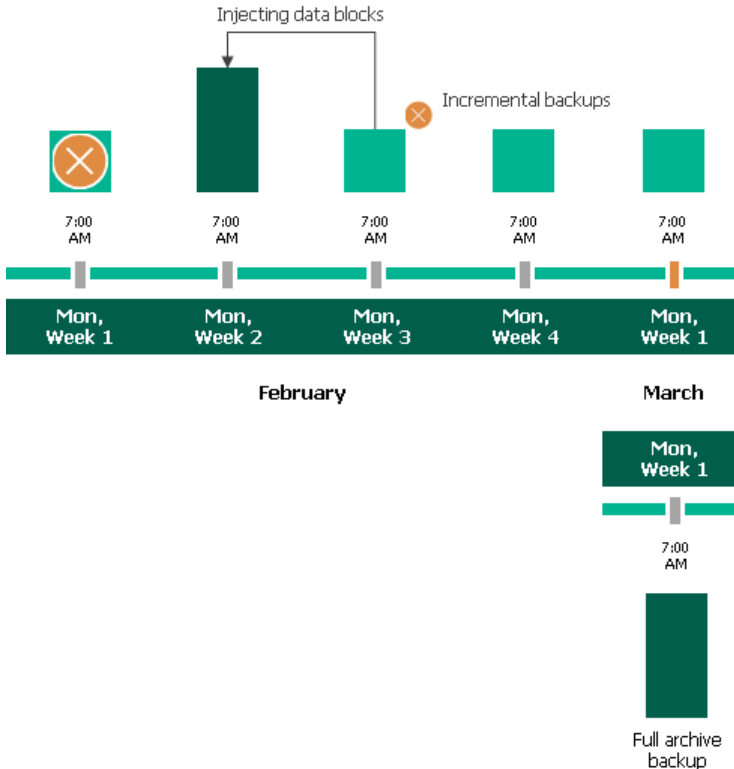
- On the fourth Monday of February, Veeam Backup for Google Cloud will create a new restore point at 7:00 AM. By the moment the backup session completes, the earliest restore point in the regular backup chain will get older than the specified retention limit. That is why Veeam Backup for Google Cloud will rebuild the full backup and remove from the chain the restore point created on the first Monday.

For more information on how Veeam Backup for Google Cloud transforms regular backup chains, see [Retention Policy for Backups](#).




- On the first Monday of March, a backup session will start at 7:00 AM to create another restore point in the regular backup chain. At the same time, the earliest restore point in the regular backup chain will get older than the specified retention limit again. That is why Veeam Backup for Google Cloud will rebuild the full backup 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 the data from the regular backup chain. Veeam Backup for Google Cloud will copy this restore point as a full archive backup to the archive repository.



- Up to May, Veeam Backup for Google Cloud will continue adding new restore points to the regular backup chain and deleting outdated backups from the standard 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 Google Cloud will copy this restore point as an incremental archive backup to the archive repository.



- Up to the first Monday of March of the next year, Veeam Backup for Google Cloud will continue adding new restore points to the regular backup chain and deleting outdated backups from the standard repository, according to the specified weekly scheduling settings. Veeam Backup for Google Cloud 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 Google Cloud will rebuild the full archive backup 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 Google Cloud transforms archive backup chains, see [Retention Policy for Archived Backups](#).



Step 7. Enable Label Assignment

At the **Labels** step, you can instruct Veeam Backup for Google Cloud to assign labels to cloud-native snapshots created by the backup policy:

1. Click the **Edit settings** link.
2. In the **Choose labels to assign** window, choose whether you want to assign to snapshots of the selected VM instances already existing labels from source persistent disks and your own custom labels.

If you set the **Assign custom labels** toggle to *On*, you must also specify the labels explicitly. To do that, use the **Name** and **Value** fields to specify a name and a value for the new custom label, and then click **Add**. Note that you cannot add more than 5 custom labels.

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

The screenshot displays the Veeam Backup for Google Cloud configuration interface. At the top, the header shows 'Veeam Backup for Google Cloud' on the left, 'Server time: Nov 13, 2023 4:16 PM' in the center, and 'administrator Portal Administrator' on the right. Below the header, the main area is titled 'Add VM Policy' with a 'Cost: \$29.21' indicator. A left sidebar contains navigation options: Policy Info, Sources, Resources, Targets, Schedule, Labels (highlighted), Settings, Cost Estimation, Permissions, and Summary. The main content area is divided into two panels. The left panel, 'Specify label settings', includes a sub-header and instructions: 'You can copy labels from source disks and additionally assign up to 5 custom labels to snapshots.' It shows 'Source labels: Will be copied' and 'Custom labels: 1 custom label will be assigned', with an 'Edit settings...' link. The right panel, 'Choose labels to assign', features a checkbox for 'Copy labels from source disks' (checked), a toggle for 'Assign custom labels: On', and input fields for 'Name' and 'Value'. An 'Add' button is next to the 'Value' field. A list below shows a custom label 'us-west-loc: west-location' with a close icon. A note states 'A maximum of 5 custom labels is allowed'. At the bottom of the right panel are 'Apply' and 'Cancel' buttons. At the bottom of the entire interface are 'Previous', 'Next', and 'Cancel' buttons.

Step 8. Configure General Settings

At the **Settings** 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 Google Cloud to run the backup policy again if it fails on the first try, do the following:

1. In the **Retries** 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 15 minutes.

When retrying backup policies, Veeam Backup for Google Cloud processes only those VM instances that failed to be backed up during the previous attempt.

NOTE

The automatic retry settings apply only to backup policies that run according to specific schedules — these settings do not apply to policies [started manually](#).

Health Check Settings

If you have enabled creation of image-level backups at [step 5](#), you can instruct Veeam Backup for Google Cloud to periodically perform a health check for backup restore points created by the backup policy. During the health check, Veeam Backup for Google Cloud performs an availability check for data blocks in the whole regular backup chain, and a cyclic redundancy check (CRC) for 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 Google Cloud does not verify archived restore points created by the policy.

To instruct Veeam Backup for Google Cloud to perform a monthly health check, 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 Google Cloud performs the health check during the last 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 Google Cloud 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 last policy session on Saturday.

Notification Settings

To instruct Veeam Backup for Google Cloud to send email notifications for the backup policy, do the following:

1. In the **Notifications** section of the step, set the **Enable notifications** toggle to *On*.
If you set the toggle to *Off*, Veeam Backup for Google Cloud 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.
3. Use the **Notify on** list to choose whether you want Veeam Backup for Google Cloud 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 Google Cloud 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 Google Cloud will send each notification to this recipient twice.

The screenshot shows the 'Add VM Policy' configuration page in Veeam Backup for Google Cloud. The 'Settings' tab is selected, displaying the following configuration options:

- Retries:** Automatically retry failed policy: 3 times
- Health check:** Enable health check: Off
- Notifications:** Enable notifications: On
Email: john.smith@veeam.com
 Success
 Warning
 Failure
 Suppress notifications until the last retry

Navigation buttons at the bottom: Previous, Next, Cancel.

How Health Check Works

When Veeam Backup for Google Cloud saves a new 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 VM instance

data. When performing a health check, Veeam Backup for Google Cloud verifies availability of data blocks for each restore point and uses the saved values to ensure that the restore points being verified are consistent.

NOTE

Veeam Backup for Google Cloud performs the health check during the last 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 Google Cloud 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 last policy session on Saturday.

Veeam Backup for Google Cloud performs the health check in the following way:

1. As soon as the backup policy session completes successfully, Veeam Backup for Google Cloud starts the health check as a new session. For each restore point in the regular backup chain, Veeam Backup for Google Cloud calculates CRC values for backup metadata and compares them to the CRC values that were previously saved to the restore point. Veeam Backup for Google Cloud 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 Google Cloud 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 Google Cloud starts the health check.

2. If Veeam Backup for Google Cloud 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 Google Cloud performs the following operations:

- If the health check detects corrupted metadata in a full or an incremental restore point, Veeam Backup for Google Cloud marks the regular backup chain as corrupted in the configuration database. During the next backup policy session, Veeam Backup for Google Cloud copies the full VM instance image, creates a new full restore point and starts a new backup chain in the backup repository.

NOTE

Veeam Backup for Google Cloud supports metadata check for encrypted backup chains unless the metadata is corrupted.

- If the health check detects corrupted data blocks in a full or an incremental restore point, Veeam Backup for Google Cloud marks the restore point that includes the corrupted data blocks and all subsequent incremental restore points as incomplete in the configuration database. During the next backup policy session, Veeam Backup for Google Cloud copies not only those data blocks that have changed since the previous backup session but also data blocks that have been corrupted. Veeam Backup for Google Cloud then 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 approximate monthly cost of Google Cloud services that Veeam Backup for Google Cloud will require to protect the VM instances added to the backup policy. The total estimated cost includes the following:

- The cost of creating and maintaining cloud-native snapshots of the VM instances.
For each VM instance included in the backup policy, Veeam Backup for Google Cloud takes into account the total size and the number of persistent disks 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 VM instances.
For each VM instance included in the backup policy, Veeam Backup for Google Cloud takes into account the total size and the number of persistent disks attached, the period of time during which restore points will be kept in the backup chain, and the configured scheduling and health check settings.
- The cost of creating and storing in backup repositories archived backups of the VM instances.
For each VM instance included in the backup policy, Veeam Backup for Google Cloud takes into account the total size and the number of persistent disks attached, the period of time during which restore points will be kept in the archive backup chain, and the configured scheduling settings.
- The cost of transferring the VM instance data between Google Cloud regions during data protection operations (for example, if a protected VM instance and the target storage bucket reside in different regions).
- The cost of sending API requests to Google Cloud during data protection operations.
- The cost of deploying worker instances for backup operations.

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, either select a backup repository that resides in the same region as VM instances that you plan to back up, or select an archive repository that resides in the same region as the nearline or standard repository used to store regular backups.
- 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, modify the scheduling settings to run the backup policy less frequently, or specify an archive repository for long-term retention of restore points.

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 Google Cloud makes predefined assumptions to calculate the cost, which means that the results should be used only as an approximation. For more information on how Veeam Backup for Google Cloud calculates the cost, see this [Veeam KB article](#).

\$12.13	\$16.34	\$0.00	\$0.36	\$0.38
---------	---------	--------	--------	--------

Estimated monthly cost: \$29.21

Instance

Instance	Snapshots	Backups	Archives	Traffic	Transactions	Total
dr-vb-v5	\$4.96	\$4.95	\$0.00	\$0.00	\$0.15	\$10.06
tv-g-iam	\$1.38	\$3.26	\$0.00	\$0.11	\$0.05	\$4.80
tv-g-ta...	\$1.65	\$3.84	\$0.00	\$0.00	\$0.05	\$5.55
tv-g-vb...	\$4.14	\$4.30	\$0.00	\$0.25	\$0.13	\$8.81

Step 10. Check Required Permissions

At the **Permissions** step of the wizard, Veeam Backup for Google Cloud will verify whether the specified service account has all the necessary permissions required to perform data protection tasks for the selected project or folder. For more information on the required permissions, see [Service Account Permissions](#).

To see the list of missing permissions that must be granted to the service account in order to perform an operation, click the link in the **Details** column. You can grant the missing permissions to the service account [using the Google Cloud console](#) or instruct Veeam Backup for Google Cloud to do it:

- To grant the missing permissions manually, click **Download Script**. Veeam Backup for Google Cloud will generate a gcloud script that you can run in the Google Cloud console to assign all the necessary permissions to the service account.

The account under which you run the script must have the permissions required both to get and set project IAM policies and to create custom IAM roles (for example, it can have the *iam.securityAdmin* and *iam.roleAdmin* roles assigned). To learn what permissions and roles are required to create custom roles in IAM, see [Google Cloud documentation](#).

- To let Veeam Backup for Google Cloud grant the missing permissions automatically, click **Grant** and then click **Sign in with Google** in the **Grant permissions** window. You will be redirected to the OAuth consent screen authorization page. Sign in using credentials of a Google account that will be used to grant the permissions.

The account under which you sign in to Google Cloud must have the permissions required both to get and set project IAM policies and to create custom IAM roles (for example, it can have the *iam.securityAdmin* and *iam.roleAdmin* roles assigned). To learn what permissions and roles are required to create service account, see [Google Cloud documentation](#).

NOTE

For Veeam Backup for Google Cloud to be able to authorize in Google Cloud, the OAuth consent screen must be configured as described in section [Registering Applications](#). Note that Veeam Backup for Google Cloud does not store in the configuration database the provided Google account credentials and access tokens received during authorization.

To make sure that the missing permissions have been successfully granted, click **Recheck**.

The screenshot shows the 'Add VM Policy' configuration page in Veeam Backup for Google Cloud. The top navigation bar includes the product name, server time (Nov 13, 2023 4:18 PM), user (administrator), and configuration options. The main content area is titled 'Add VM Policy' and shows a cost of \$29.21. A sidebar on the left lists various configuration steps, with 'Permissions' currently selected. The 'Check permissions' section displays a table with the following data:

Download Script	Result	Details
VM Snapshot	Passed	All the required permissions are ...
VM Backup	Passed	All the required permissions are ...
Repository	Passed	All the required permissions are ...
Worker	Passed	All the required permissions are ...

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

Step 11. Finish Working with Wizard

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

The screenshot shows the 'Add VM Policy' wizard in the 'Summary' step. The interface includes a top navigation bar with the Veeam logo, 'Veeam Backup for Google Cloud', server time (Nov 13, 2023 4:20 PM), user (administrator), and configuration options. A left sidebar lists steps: Policy Info, Sources, Resources, Targets, Schedule, Labels, Settings, Cost Estimation, Permissions, and Summary (highlighted). The main content area is titled 'Review configured settings' and contains a 'Copy to Clipboard' button. It is divided into several sections: 'General settings' (Name: us-west-policy, Description: protecting instances in us-west regions, Service account: veeam-1649186685-sa@rnd-backup-2.iam.gserviceaccount.com, Project: veeam-rnd-backup-2, Folder: —), 'Protected resources' (Regions: 4 regions, Instances: 4 instances, Labels: —, Exclusions: —), 'Snapshot settings' (Snapshots enabled: Yes, Region: Regional, Copy labels from source disks: Yes, Add custom labels: Yes, Custom labels: 1 custom label, Daily retention: Create 3 snapshots per day and keep 7 most recent snapshots, Weekly retention: Keep 7 weekly snapshots, Monthly retention: Keep 7 monthly snapshots), 'Backup settings' (Backups enabled: Yes, Archives enabled: Yes, Daily retention: Create 3 backups per day and keep for 21 days, Weekly retention: Keep weekly backup for 21 days (6 backups excluded), Monthly retention: Keep monthly backup for 30 days (9 backups excluded), Yearly retention: 3 years), and 'Other settings' (Automatic retries enabled: Yes, Notifications enabled: Yes, Health check enabled: No). At the bottom, there are 'Previous', 'Finish', and 'Cancel' buttons.

Creating Snapshots Manually

Veeam Backup for Google Cloud allows you to manually create snapshots of VM instances. Each snapshot is saved to the multi-regional location closest to the region in which the original VM instance resides.

NOTE

Veeam Backup for Google Cloud 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 Google Cloud Storage unless you remove them manually, as described in section [Removing Backups and Snapshots](#).

To manually create a cloud-native snapshot of a VM instance, do the following:

1. Navigate to **Resources > VM**.

2. Select the necessary instance and click **Take Snapshot Now**.

For a VM instance to be displayed in the list of available instances, it must reside in any of the regions added to a backup policy as described in section [Creating Backup Policies](#).

3. In the **Take Snapshot Now** window:

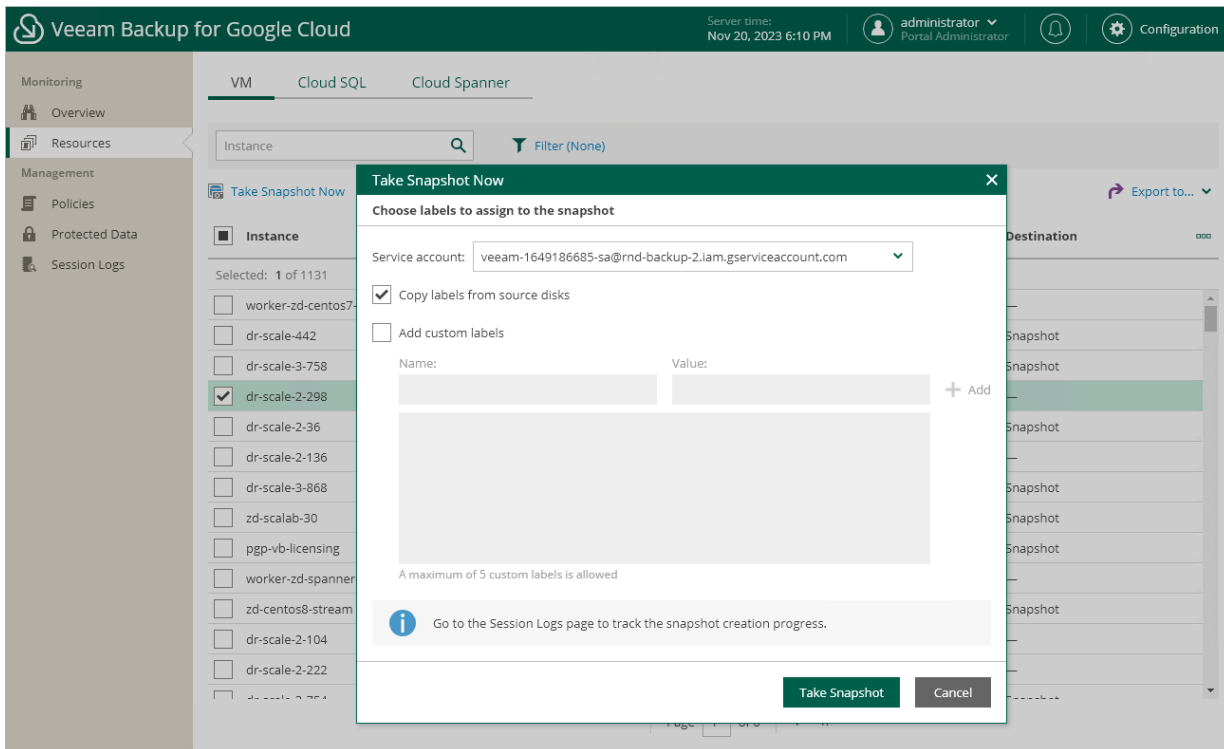
- a. Specify a service account whose permissions Veeam Backup for Google Cloud will use to create the snapshot.

For a service account to be displayed in the list of available accounts, it must be added to Veeam Backup for Google Cloud as described in section [Adding Service Accounts](#), and must be assigned the *VM Instances Snapshot* operational role as described in section [Adding Projects and Folders](#).

- b. Choose whether you want to assign labels to the created snapshot:

- § To assign already existing labels from the source persistent disk attached to the selected VM instance, select the **Copy labels from source disks** check box.
- § To assign your own custom labels, select the **Add custom labels** check box and specify the labels explicitly. To do that, use the **Name** and **Value** fields to specify a key and a value for the new custom label, and then click **Add**.

- c. Click **Take Snapshot**.



Performing SQL Backup

One backup policy can be used to process one or more Cloud SQL instances within one Google Cloud project or folder. The scope of data that you can protect in a project or folder is limited by permissions of a service account that is specified in the backup policy settings.

Before you create a Cloud SQL backup policy, check the following prerequisites:

- If you plan to create image-level backups of Cloud SQL 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 the backup policy results, configure SMTP server settings first. For more information, see [Configuring Global Notification Settings](#).

NOTE

Veeam Backup for Google Cloud allows you to protect MySQL and PostgreSQL instances. SQL Server instances are not supported. For more information on types of Cloud SQL instances, see [Google Cloud documentation](#).

To schedule data protection tasks to run automatically, [create backup policies](#). For each protected Cloud SQL instance, you can also [take a cloud-native snapshot manually](#) when needed.

Creating Backup Policies

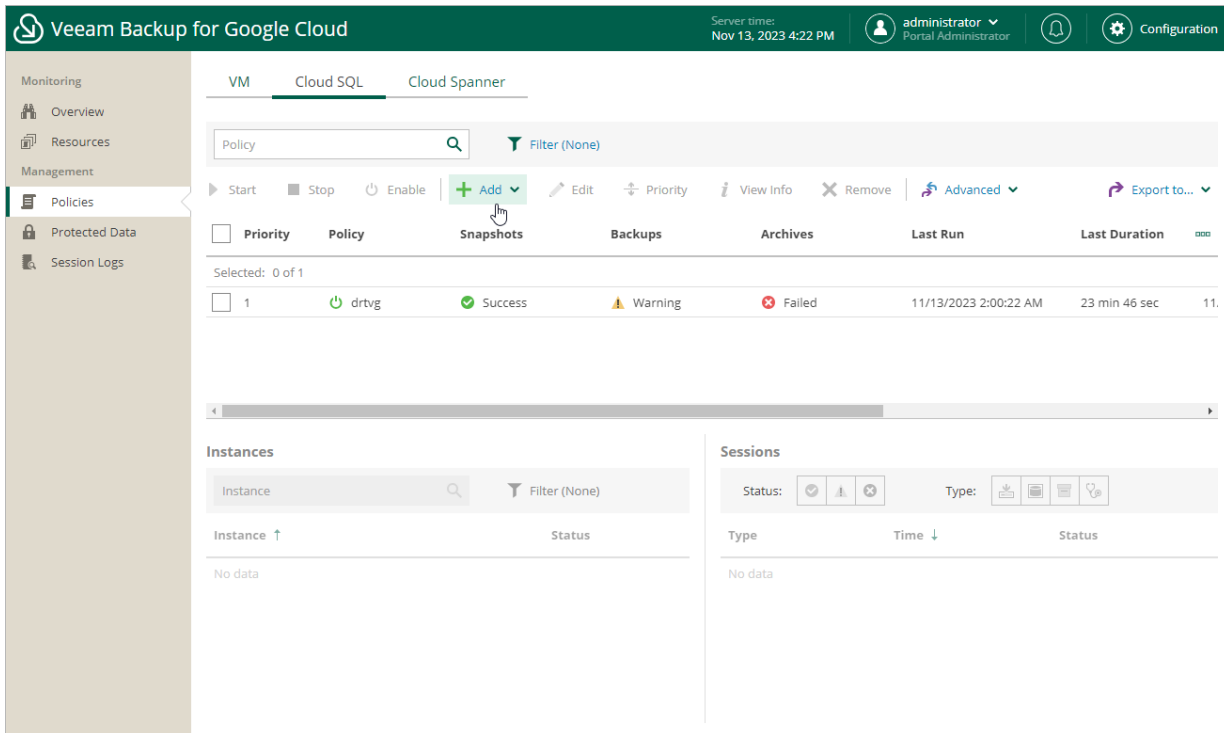
To create a backup policy, do the following:

1. [Launch the Add Cloud SQL Policy wizard](#).
2. [Specify a backup policy name and description](#).
3. [Choose a project to which Cloud SQL instances that you plan to back up belong](#).
4. [Configure backup source settings](#).
5. [Configure backup target settings](#).
6. [Create a schedule for the backup policy](#).
7. [Choose whether you want to use a staging server to perform backup](#).
8. [Specify automatic retry, health check and notification settings for the backup policy](#).
9. [Review the estimated cost of protecting the selected Cloud SQL instances](#).
10. [Check the required permissions](#).
11. [Finish working with the wizard](#).

Step 1. Launch Add Cloud SQL Policy Wizard

To launch the **Add Cloud SQL Policy** wizard, do the following:

1. Navigate to **Policies > Cloud SQL**.
2. Click **Add** and select either of the following options:
 - **MySQL** – to create a backup policy that will protect MySQL instances.
 - **PostgreSQL** – to create a backup policy that will protect PostgreSQL instances.



The screenshot shows the Veeam Backup for Google Cloud interface. The top navigation bar includes the Veeam logo, the product name, server time (Nov 13, 2023 4:22 PM), user information (administrator), and a Configuration icon. The left sidebar shows navigation options: Monitoring (Overview, Resources) and Management (Policies, Protected Data, Session Logs). The main content area is titled 'Cloud SQL' and contains a table of policies. A search bar and a 'Filter (None)' dropdown are at the top. Below the search bar are action buttons: Start, Stop, Enable, Add (highlighted with a mouse cursor), Edit, Priority, View Info, Remove, Advanced, and Export to... The table has columns for Priority, Policy, Snapshots, Backups, Archives, Last Run, and Last Duration. One policy is listed with ID 1, policy name 'drtvg', status 'Success', a warning icon, a failed icon, last run '11/13/2023 2:00:22 AM', and duration '23 min 46 sec'. Below the table are two sections: 'Instances' and 'Sessions', both showing 'No data'.

Step 2. Specify Backup Policy Name and Description

At the **Policy Info** step of the wizard, use the **Name** and **Description** fields to enter a name for the new backup policy and to provide a description for future reference. The policy name can contain only uppercase Latin letters, lowercase Latin letters, numeric characters and hyphens; the maximum length of the name is 127 characters.

NOTE

When Veeam Backup for Google Cloud runs a backup policy, it adds the word *Veeam* to the descriptions of snapshots created by the policy.

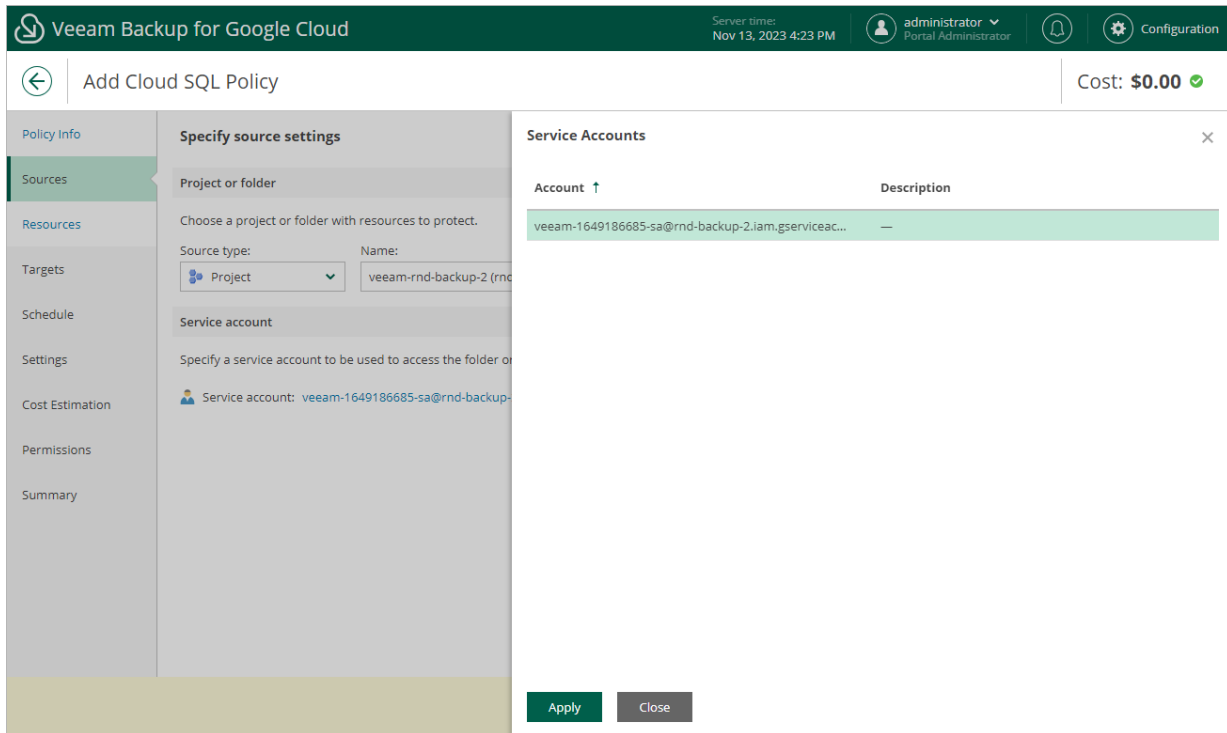
The screenshot shows the 'Add Cloud SQL Policy' wizard in the Veeam Backup for Google Cloud interface. The top navigation bar includes the Veeam logo, the title 'Veeam Backup for Google Cloud', the server time 'Nov 13, 2023 4:23 PM', the user 'administrator Portal Administrator', and a 'Configuration' link. The main header shows a back arrow, the title 'Add Cloud SQL Policy', and the cost '\$0.00'. A left sidebar lists the wizard steps: Policy Info (selected), Sources, Resources, Targets, Schedule, Settings, Cost Estimation, Permissions, 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 text input fields: 'Name:' with the value 'europe-policy' and 'Description:' with the value 'protecting instances in other EU regions'. At the bottom of the form are 'Next' and 'Cancel' buttons.

Step 3. Specify Project

At the **Sources** step of the wizard, choose a project or a folder with a project that manages resources that you want to protect, and specify a service account that will be used to access the project or folder.

For a project or folder to be displayed in the list of available entities, it must be added to Veeam Backup for Google Cloud as described in section [Adding Projects and Folders](#). If you have not added the necessary entity to Veeam Backup for Google Cloud beforehand, you can do it without closing the **Add Cloud SQL Policy** wizard. To add an entity, click **Add** and complete the **Add Projects and Folders** wizard.

For a service account to be displayed in the list of available accounts, it must be added to Veeam Backup for Google Cloud as described in section [Adding Service Accounts](#), and must be assigned the *Cloud SQL Instances Snapshot* and *Backup* operational roles as described in section [Adding Projects and Folders](#).



Step 4. Configure Backup Source Settings

At the **Resources** step of the wizard, specify the following backup source settings:

1. [Choose regions in which Cloud SQL instances that you plan to back up reside.](#)
2. [Select Cloud SQL instances to back up.](#)

Step 4a. Choose Regions

In the **Regions** section of the **Resources** step of the wizard, choose regions in which Cloud SQL instances that you want to protect reside.

1. Click **Choose regions**.
2. In the **Choose regions** window, select the necessary 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 Cloud SQL Policy' wizard in the Veeam Backup for Google Cloud interface. The 'Resources' step is active, and the 'Choose regions' window is open. The 'Available Regions (36)' list includes: europe-west2 (London), europe-west3 (Frankfurt), europe-west4 (Netherlands), europe-west10 (Berlin), europe-west12 (Turin) (highlighted), me-central1 (Doha), me-central2 (Dammam), me-west1 (Tel Aviv), northamerica-northeast1 (Montréal), northamerica-northeast2 (Toronto), southamerica-east1 (São Paulo), southamerica-west1 (Santiago), us-central1 (Iowa), us-east1 (South Carolina), and us-east4 (Northern Virginia). The 'Add' button is highlighted with a mouse cursor. The 'Policy Regions (3)' list contains: europe-west6 (Zürich), europe-west8 (Milan), and europe-west9 (Paris). The 'Apply' button is visible at the bottom of the window. The top navigation bar shows 'Veeam Backup for Google Cloud', 'Server time: Nov 13, 2023 4:24 PM', 'administrator Portal Administrator', and 'Configuration'. The left sidebar shows 'Policy Info', 'Sources', 'Resources', 'Targets', 'Schedule', 'Settings', 'Cost Estimation', 'Permissions', and 'Summary'. The top right corner shows 'Cost: \$0.00' with a green checkmark.

Step 4b. Select Cloud SQL Instances

In the **Resources** section of the **Resources** step of the wizard, specify the backup scope — select Cloud SQL instances that Veeam Backup for Google Cloud will back up:

1. Click **Choose resources to protect**.
2. In the **Choose resources** window, choose whether you want to back up all Cloud SQL instances from the regions selected at [step 4a](#), or only specific Cloud SQL instances.

If you select the **All resources** option, Veeam Backup for Google Cloud will regularly check for new Cloud SQL instances launched in the selected regions and automatically update the backup policy settings to include these instances in the backup scope.

If you select the **Specific resources** option, you must also specify the instances explicitly:

- a. Use the **Resource type** drop-down list to choose whether you want to add individual Cloud SQL instances or Google Cloud labels to the backup scope.

If you select the **Label** option, Veeam Backup for Google Cloud will back up only those Cloud SQL instances that reside in the selected regions under specific labels.

- b. Use the **Instance\Label** list to find the necessary resource, and then click **Add to Protected** to add the resource to the backup scope.

For a resource to be displayed in the list of available resources, it must reside in a region that has ever been specified in any backup policy. Otherwise, the only option to discover available resources is to click **Browse** and wait for Veeam Backup for Google Cloud to populate the resource list. Note that Veeam Backup for Google Cloud automatically filters the resource list to show only either MySQL or PostgreSQL instances — depending on the option that you have selected while running the wizard.

TIP

You can simultaneously add multiple resources to the backup scope. To do that, click **Browse**, select check boxes next to the necessary Cloud SQL instances or labels 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 Google Cloud will update the resource list.

If you add a label to the backup scope, Veeam Backup for Google Cloud will regularly check for new Cloud SQL instances assigned the added label and automatically update the backup policy settings to include these instances in the scope. However, this applies only to Cloud SQL instances from the regions selected at [step 4a](#). If you select a label assigned to Cloud SQL instances from other regions, these instances will not be protected by the backup policy. To work around the issue, either go back to step 4a 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 **Specific 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 **Exclude resources** and specify the Cloud SQL instances or labels that you do not want to backup — the procedure is the same as described for including resources in the backup scope.

Consider that if a resource appears both in the list of included and excluded resources, Veeam Backup for Google Cloud will still not process the resource because the list of excluded resources has a higher priority.

The screenshot displays the Veeam Backup for Google Cloud interface for adding a Cloud SQL policy. The main window is titled "Add Cloud SQL Policy" and shows a sidebar with navigation options: Policy Info, Sources, Resources, Targets, Schedule, Settings, Cost Estimation, Permissions, and Summary. The "Resources" section is active, showing "Specify resources to protect or exclude." with a lock icon and "1 resource will be protected" and an "Exclude resources..." option. A "Choose regions" dialog is open, showing a list of "Available Regions (34)" and a list of "Policy Regions (5)". The "Available Regions" list includes: asia-east1 (Taiwan), asia-east2 (Hong Kong), asia-northeast1 (Tokyo), asia-northeast2 (Osaka), asia-northeast3 (Seoul), asia-south1 (Mumbai), asia-south2 (Delhi), asia-southeast1 (Singapore), asia-southeast2 (Jakarta), australia-southeast1 (Sydney), australia-southeast2 (Melbourne), europe-central2 (Warsaw), europe-north1 (Finland), and europe-southwest1 (Madrid). The "Policy Regions" list includes: europe-west1 (Belgium), europe-west3 (Frankfurt), europe-west6 (Zürich), europe-west8 (Milan), and europe-west9 (Paris). There are "Add" and "Remove" buttons between the lists, and "Apply" and "Cancel" buttons at the bottom of the dialog. The top right of the interface shows the user "administrator" and the cost "\$0.00".

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 instruct Veeam Backup for Google Cloud to create image-level backups of the selected Cloud SQL instances:

1. Set the **Enable backups** toggle to *On*.
2. Click **Choose repository**.
3. In the **Choose repository** window, select a backup repository where the created image-level backups will be stored.

For a backup repository to be displayed in the **Repository** list, it must be added to Veeam Backup for Google Cloud as described in section [Adding Backup Repositories](#). The **Repository** list shows only backup repositories of the *Standard* and *Nearline* storage classes.

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

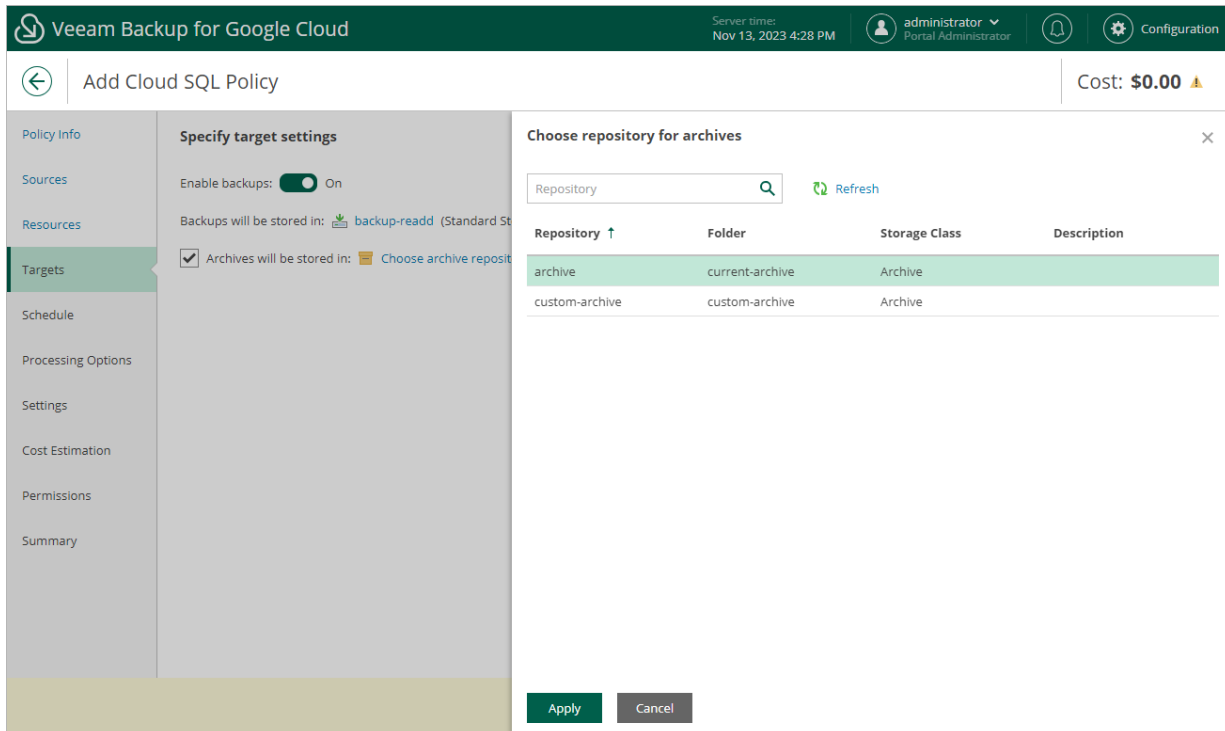
You can also enable the backup archiving mechanism to instruct Veeam Backup for Google Cloud to store backed-up data in a low-cost, long-term archive storage:

1. Select the **Archives will be stored in** check box.
2. Click **Choose repository**.
3. In the **Choose repository** window, select a backup repository where the archived data will be stored.

For a backup repository to be displayed in the **Repository** list, it must be added to Veeam Backup for Google Cloud as described in section [Adding Backup Repositories](#). The **Repository** list shows only backup repositories of the *Archive* storage class.

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

For more information on the backup archiving mechanism, see [Enabling Backup Archiving](#).



The screenshot shows the Veeam Backup for Google Cloud interface. The top navigation bar includes the Veeam logo, the text 'Veeam Backup for Google Cloud', the server time 'Nov 13, 2023 4:28 PM', the user 'administrator Portal Administrator', and a 'Configuration' button. The main content area is titled 'Add Cloud SQL Policy' and shows a 'Cost: \$0.00' warning. The 'Targets' step is selected in the left sidebar. The 'Specify target settings' section is active, showing 'Enable backups: On' and 'Backups will be stored in: backup-readd (Standard Storage Class)'. The 'Archives will be stored in' checkbox is checked, and a 'Choose archive repository' button is visible. A 'Choose repository for archives' dialog is open, displaying a search bar and a table of repositories.

Repository	Folder	Storage Class	Description
archive	current-archive	Archive	
custom-archive	custom-archive	Archive	

Step 6. Specify Policy Scheduling Options

At the **Schedule** step of the wizard, you can instruct Veeam Backup for Google Cloud to start the backup policy automatically according to a specific backup schedule. The backup schedule defines how often data of the Cloud SQL instances added to the backup policy will be backed up.

To help you implement a comprehensive backup strategy, Veeam Backup for Google Cloud 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](#).

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** section, select hours when the backup policy will create cloud-native snapshots and image-level backups. Use the **Run at** drop-down list to choose whether you want the backup policy to run every day, on weekdays (Monday through Friday) or on specific days.

If you want to protect Cloud SQL 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 Google Cloud does not create image-level backups independently from cloud-native snapshots. That is why when you select hours for image-level backups, the same hours are automatically selected for cloud-native snapshots. To learn how Veeam Backup for Google Cloud performs backup, see [How Backup Works](#).

3. In the **Configure daily retention** section, configure retention policy settings for the daily schedule:
 - For cloud-native snapshots, specify the number of restore points that you want to keep in a snapshot chain.

If the restore point limit is exceeded, Veeam Backup for Google Cloud removes the earliest restore point from the chain. For more information, see [Retention Policy for Snapshots](#).

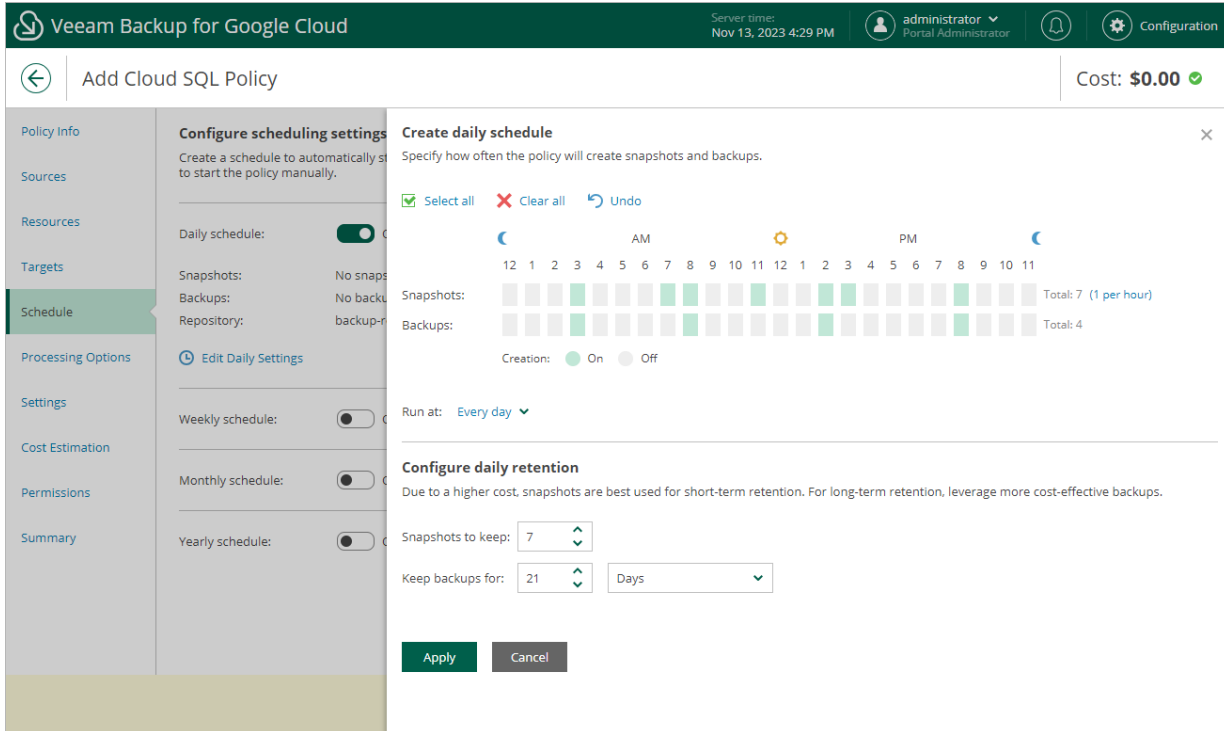
IMPORTANT

Snapshots of Cloud SQL instances are taken using [native Google Cloud capabilities](#), and therefore, if you delete a Cloud SQL instance from Google Cloud, all cloud-native snapshots created by the backup policy for the removed instance will be automatically deleted from Google Cloud Storage as well, despite the retention settings configured at the **Schedule** step of the wizard.

- o 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 Google Cloud removes the restore point from the chain. For more information, see [Retention Policy for Backups](#).

4. 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** section, select days when the backup policy will create cloud-native snapshots and image-level backups. Use the **Create restore points at** drop-down list to schedule a specific time for the backup policy to run.

NOTE

Veeam Backup for Google Cloud does not create image-level backups independently from cloud-native snapshots. That is why when you select days for image-level backups, the same days are automatically selected for cloud-native snapshots. To learn how Veeam Backup for Google Cloud performs backup, see [How Backup Works](#).

3. In the **Configure weekly retention** section, configure retention policy settings for the weekly schedule:
 - o For cloud-native snapshots, specify the number of restore points that you want to keep in a snapshot chain.

If the restore point limit is exceeded, Veeam Backup for Google Cloud removes the earliest restore point from the chain. For more information, see [Retention Policy for Snapshots](#).

IMPORTANT

Snapshots of Cloud SQL instances are taken using [native Google Cloud capabilities](#), and therefore, if you delete a Cloud SQL instance from Google Cloud, all cloud-native snapshots created by the backup policy for the removed instance will be automatically deleted from Google Cloud Storage as well, despite the retention settings configured at the **Schedule** step of the wizard.

- 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 Google Cloud removes the restore point from the chain. For more information, see [Retention Policy for Backups](#).

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

The screenshot displays the 'Add Cloud SQL Policy' wizard in the Veeam Backup for Google Cloud interface. The 'Schedule' step is active, showing configuration options for scheduling settings. A 'Create weekly schedule' dialog box is open, allowing the user to specify how often the policy will create snapshots and backups. The dialog includes a calendar grid for selecting days, a 'Create restore points at' dropdown set to 03:00 AM, and a 'Configure weekly retention' section with input fields for 'Snapshots to keep' (7) and 'Keep backups for' (21 Days). The 'Apply' button is highlighted.

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 **Choose monthly backup target** section, select months when the backup policy will create cloud-native snapshots and image-level backups. Use the **Create restore points at** and **Run on** drop-down lists to schedule a specific time and day for the backup policy to run.

NOTE

Veeam Backup for Google Cloud does not create image-level backups independently from cloud-native snapshots. That is why when you select months for image-level backups, the same months are automatically selected for cloud-native snapshots. To learn how Veeam Backup for Google Cloud performs backup, see [How Backup Works](#).

3. In the **Configure monthly retention** section, configure retention policy settings for the monthly schedule:
 - o For cloud-native snapshots, specify the number of restore points that you want to keep in a snapshot chain.

If the restore point limit is exceeded, Veeam Backup for Google Cloud removes the earliest restore point from the chain. For more information, see [Retention Policy for Snapshots](#).

IMPORTANT

Snapshots of Cloud SQL instances are taken using [native Google Cloud capabilities](#), and therefore, if you delete a Cloud SQL instance from Google Cloud, all cloud-native snapshots created by the backup policy for the removed instance will be automatically deleted from Google Cloud Storage as well, despite the retention settings configured at the **Schedule** step of the wizard.

- o 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 Google Cloud removes the restore point from the chain. For more information, see [Retention Policy for Backups](#).

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

TIP

If you have enabled backup archiving at the **Targets** step of the wizard, and want to store monthly backups in an archive backup repository, set the **Send backups to archive** toggle to *On*, and follow the instructions provided in section [Enabling Backup Archiving](#).

The screenshot shows the 'Add Cloud SQL Policy' wizard in Veeam Backup for Google Cloud. The 'Schedule' step is selected in the left-hand navigation pane. The main area displays settings for three schedules: Daily, Weekly, and Monthly. The 'Choose monthly backup target' dialog is open, showing a calendar view for selecting months for snapshots and backups. The 'Configure monthly retention' section is also visible, showing options to keep 7 snapshots and 30 backups for a specified number of days. The 'Apply' button is highlighted.

Specifying Yearly Schedule

[This step applies only if you have instructed Veeam Backup for Google Cloud 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. In the **Create 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.

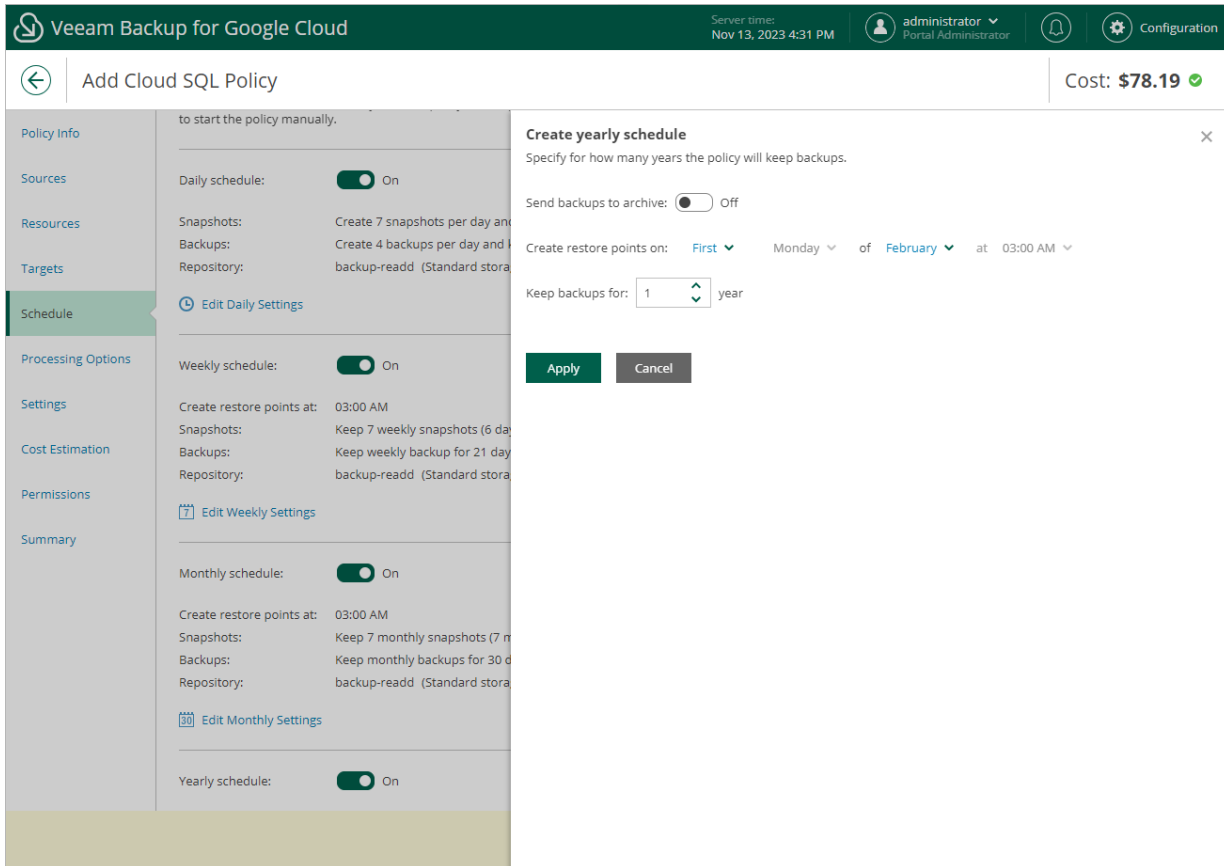
3. In the **Keep backups for** field, specify the number of years for which you want to keep restore points in a backup chain.

If a restore point is older than the specified time limit, Veeam Backup for Google Cloud removes the restore point from the chain. For more information, see [Retention Policy for Backups](#).

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

TIP

If you have enabled backup archiving at the **Targets** step of the wizard, and want to store yearly backups in an archive backup repository, set the **Send backups to archive** toggle to *On*, and follow the instructions provided in section [Enabling Backup Archiving](#).



Enabling Harmonized Scheduling

When you combine multiple types of schedules, Veeam Backup for Google Cloud 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 Google Cloud can keep restore points created according to a daily, weekly or monthly schedule for longer periods of time:

- Cloud-native snapshots can be kept for weeks and months.
- Image-level backups can be kept for weeks, months and years.

For Veeam Backup for Google Cloud 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 retaining restore points. In terms of harmonized scheduling, Veeam Backup for Google Cloud 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 (Daily) flag is used to mark restore points created daily, (Weekly) – weekly, (Monthly) – monthly, and (Yearly) – yearly. Veeam Backup for Google Cloud 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 by retention – 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.

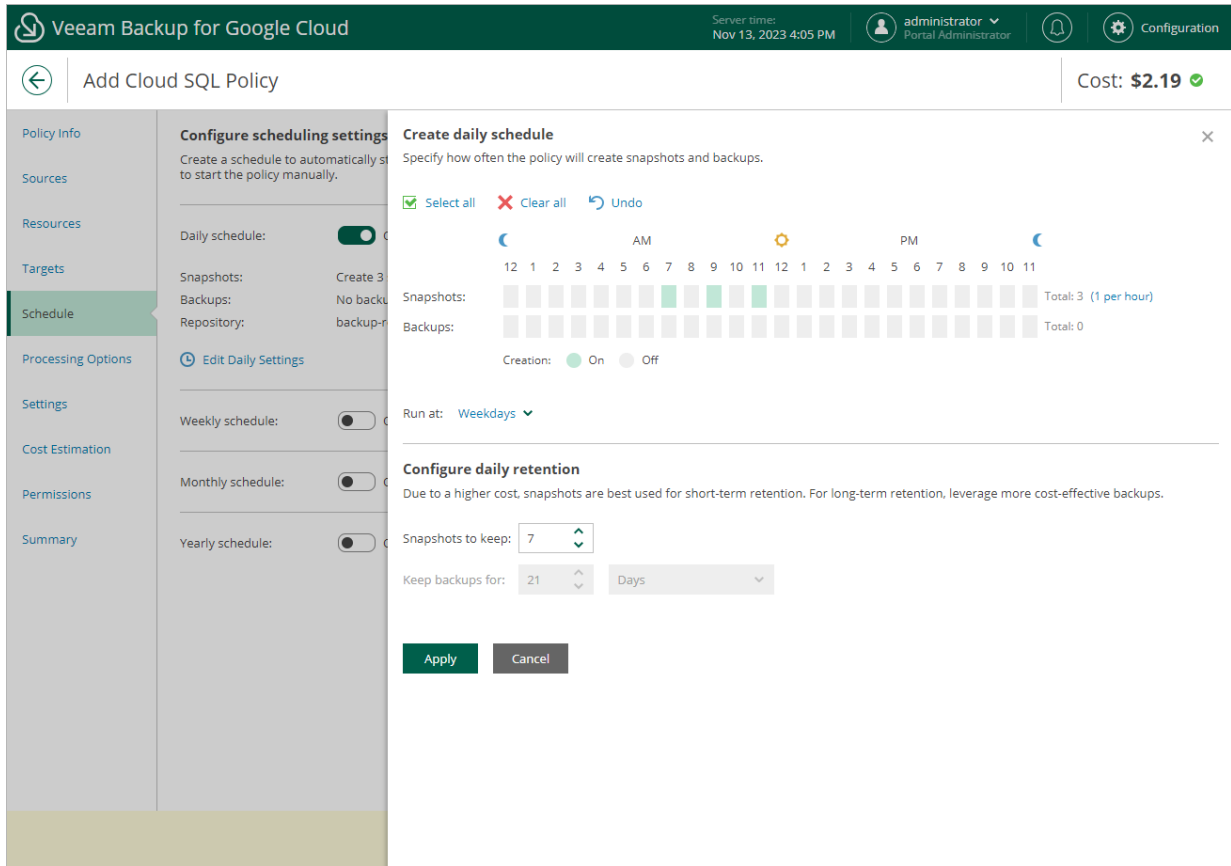
NOTE

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 Google Cloud 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 retain one of the created snapshots for 2 weeks. In this case, you create 2 schedules when configuring the backup policy settings – daily and weekly:

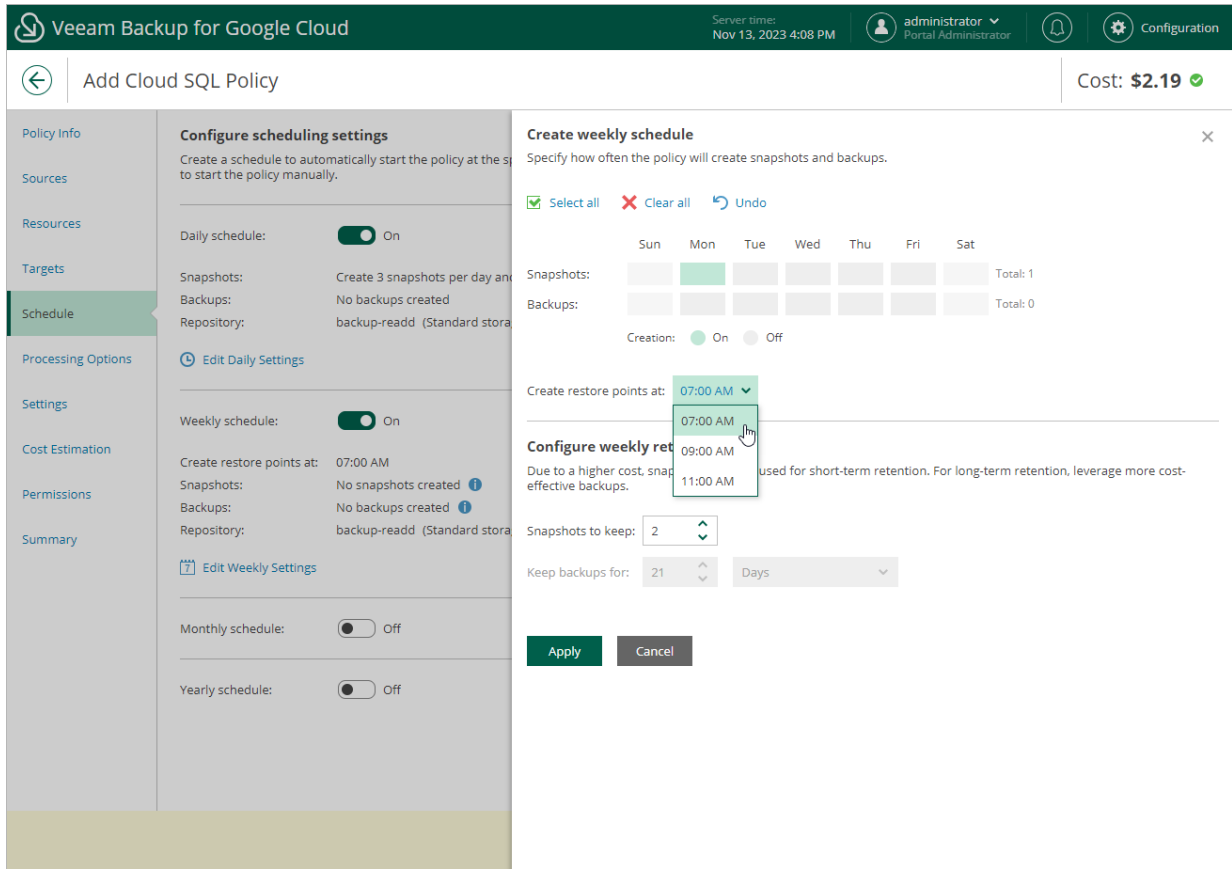
1. In the daily scheduling settings, you select hours and days when snapshots will be created (for example, *7:00 AM, 9:00 AM, and 11:00 AM; Weekdays*), and specify a number of daily restore points to retain (for example, *3*).

Veeam Backup for Google Cloud will propagate these settings to the schedule with 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 kept 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 Google Cloud will create cloud-native snapshots in the following way:

- On the first weekday (Monday), a backup session will start at 7:00 AM to create the first restore point. The restore point will be marked with the (Daily) flag as it was created according to the daily schedule.


Since *7:00 AM, Monday* is specified in the weekly scheduling settings, Veeam Backup for Google Cloud will assign the (Weekly) flag to this 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 (Daily) flag.




- On the next weekday (Tuesday), after a backup session runs at 7:00 AM, the created restore point will be marked with the (Daily) flag.


By the moment the backup session completes, the number of restore points with the (Daily) flag will exceed the retention limit specified in the daily scheduling settings. However, Veeam Backup for Google Cloud will not remove the earliest restore point (7:00 AM, Monday) with the (Daily) flag from the snapshot chain as this restore point is also marked with a flag of a less-frequent schedule. Instead, Veeam Backup for Google Cloud will unassign the (Daily) 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 (Daily) flag will exceed the retention limit once again. Veeam Backup for Google Cloud 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 Google Cloud 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 Google Cloud will unassign the (Weekly) flag from the earliest weekly restore point. Since no other flags are assigned to this restore point, Veeam Backup for Google Cloud 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 Google Cloud to store backed-up data in the low-cost, long-term Google Cloud archival storage. 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 Google Cloud standard storage.

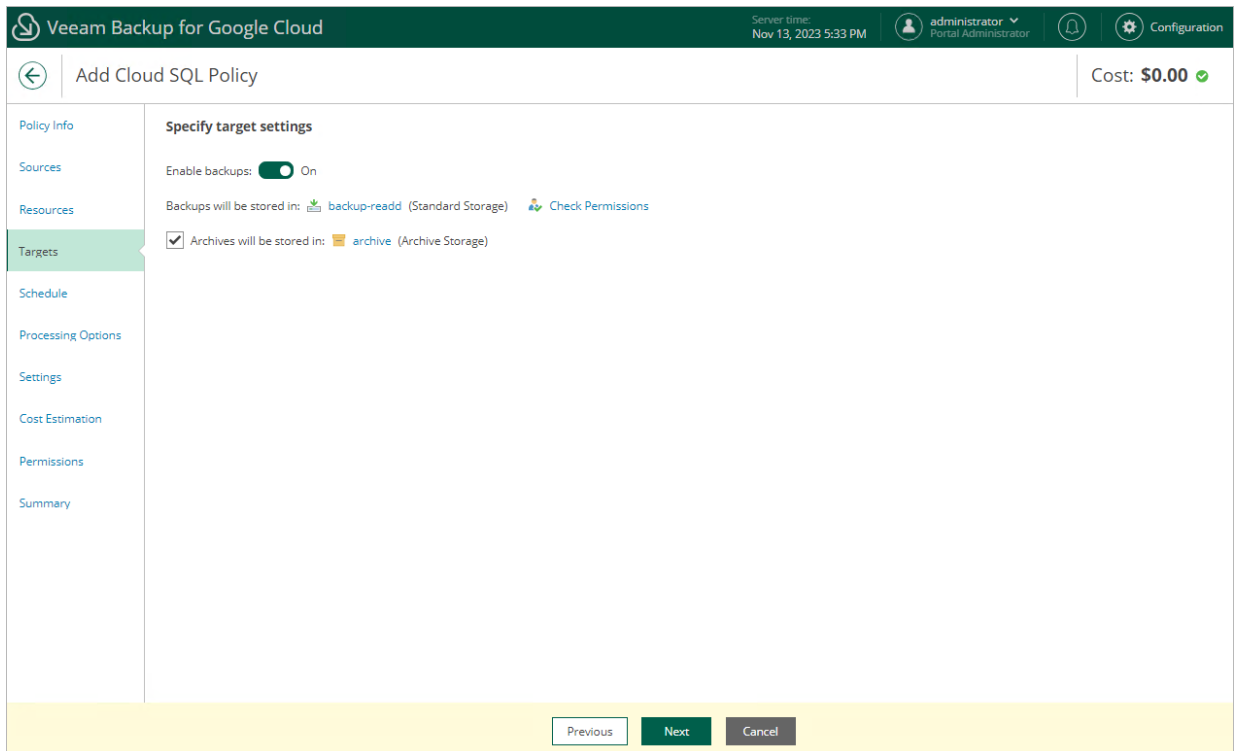
With backup archiving, Veeam Backup for Google Cloud can retain backups created according to a daily, weekly or monthly schedule for longer periods of time:

- To enable monthly archiving, you must configure a daily or a weekly backup schedule (or both).
- To enable yearly archiving, you must configure a daily, a weekly or a monthly backup schedule (or all three).

For Veeam Backup for Google Cloud to use the archiving mechanism, there must be specified at least 2 different schedules: one schedule will control the regular creation of backups, while another schedule will control the process of copying backups to an archive repository. Backup chains created according to these two schedules will be completely different – for more information, see [Backup Chain](#) and [Archive Backup Chain](#).

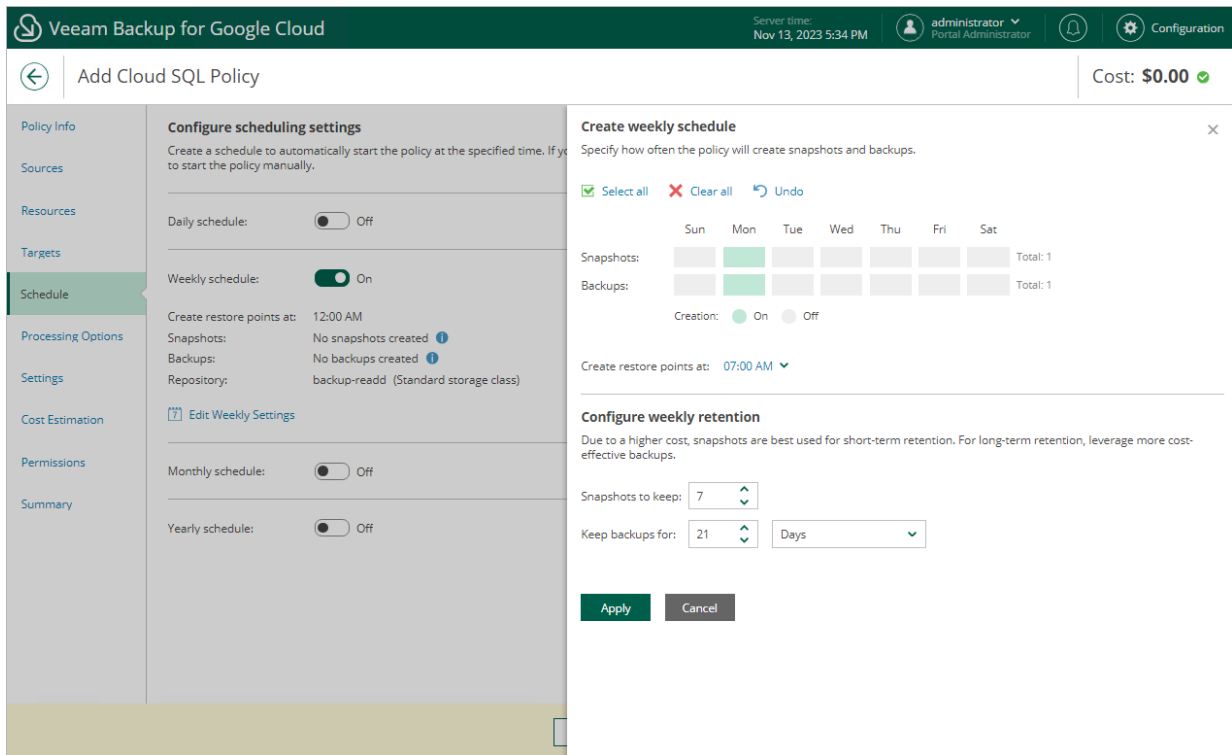
Consider the following example. You want a backup policy to create image-level backups of your critical workloads once a week, to keep the backed-up data in a standard repository for 3 weeks, and also to keep backups created once in 2 months in an archive repository for a year. In this case, you create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

1. In the policy target settings, you set the **Enable backups** toggle to *On*, select a standard repository that will store regular backups, and select an archive 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 Google Cloud will retain backups (for example, *21 days*).

Veeam Backup for Google Cloud will propagate these settings to the archive schedule (which is the monthly schedule in our example).

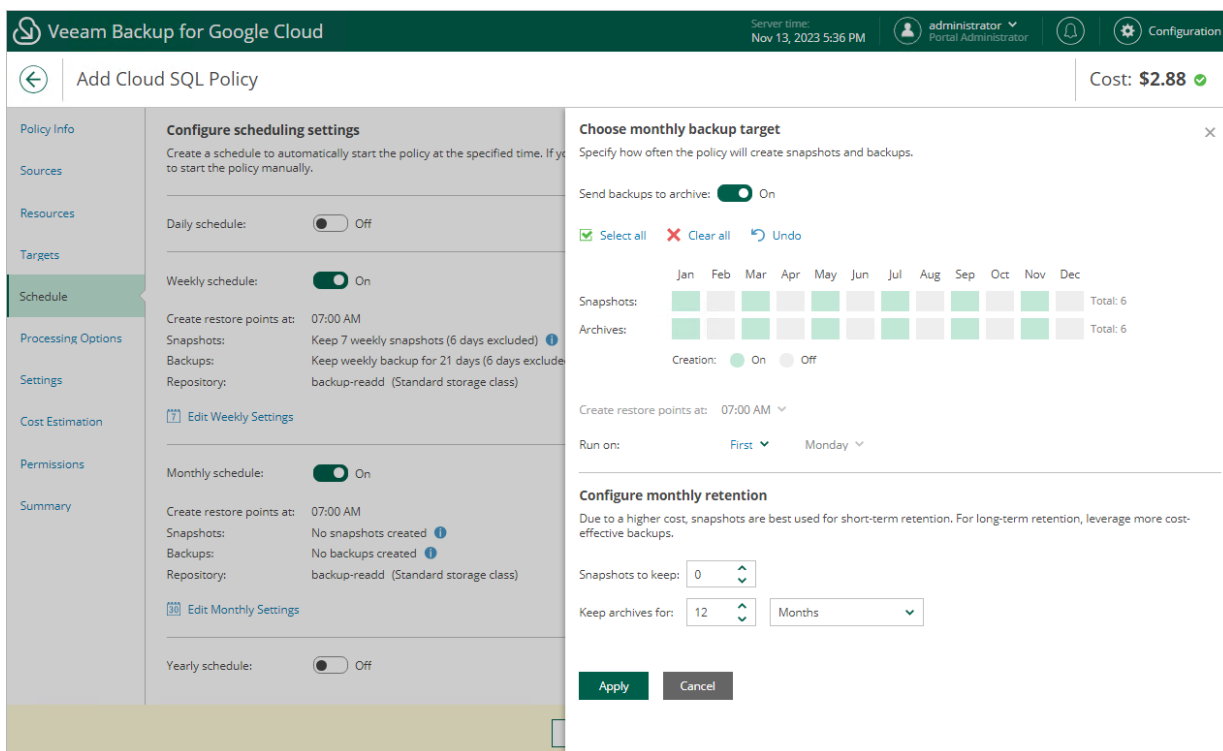


- In the monthly scheduling settings, you enable the archiving mechanism by setting the **Send backups to archive** toggle to *On*, specify when Veeam Backup for Google Cloud will create archive backups, and choose for how long you want to keep the created backups in the archive repository.

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


IMPORTANT

- When you enable backup archiving, you become no longer able to create a schedule of the same frequency for regular 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 *12 months (or 365 days)*, since the minimum storage duration of the Google Cloud archival storage is 365 days.
- If you select the **On day** option, **harmonized scheduling** cannot be guaranteed. Plus, to support the **On day** option, Veeam Backup for Google Cloud 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 Google Cloud Storage during approximately 24 hours, to reduce unexpected infrastructure charges.




According to the specified scheduling settings, Veeam Backup for Google Cloud 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 regular backup chain. Veeam Backup for Google Cloud will store this restore point as a full backup in the standard repository.
2. On the second and third Mondays of February, Veeam Backup for Google Cloud will create restore points at 7:00 AM and add them to the regular backup chain as incremental backups in the standard repository.




- On the fourth Monday of February, Veeam Backup for Google Cloud will create a new restore point at 7:00 AM. By the moment the backup session completes, the earliest restore point in the regular backup chain will get older than the specified retention limit. That is why Veeam Backup for Google Cloud will rebuild the full backup and remove from the chain the restore point created on the first Monday.

For more information on how Veeam Backup for Google Cloud transforms regular backup chains, see [Retention Policy for Backups](#).




- On the first Monday of March, a backup session will start at 7:00 AM to create another restore point in the regular backup chain. At the same time, the earliest restore point in the regular backup chain will get older than the specified retention limit again. That is why Veeam Backup for Google Cloud will rebuild the full backup 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 the data from the regular backup chain. Veeam Backup for Google Cloud will copy this restore point as a full archive backup to the archive repository.



- Up to May, Veeam Backup for Google Cloud will continue adding new restore points to the regular backup chain and deleting outdated backups from the standard repository, according to the specified weekly scheduling settings.

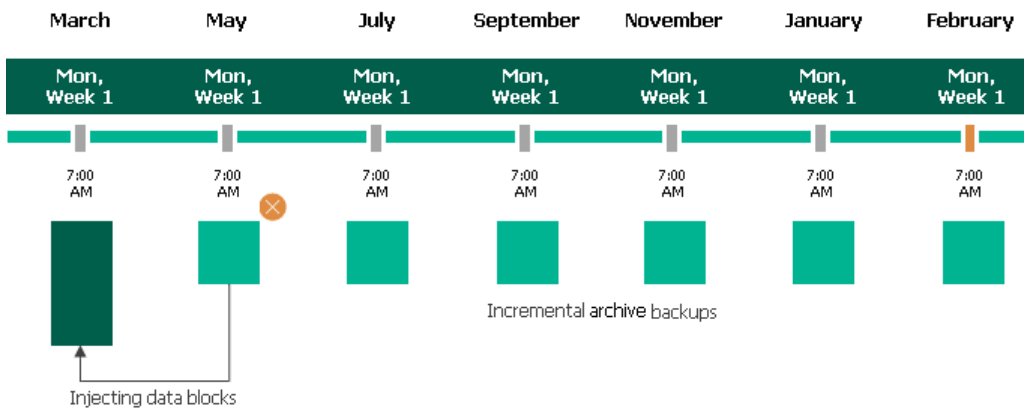
On the first Monday of May, an archive session will create a restore point with only that data that has changed since the previous archive session in March. Veeam Backup for Google Cloud will copy this restore point as an incremental archive backup to the archive repository.



- Up to the first Monday of March of the next year, Veeam Backup for Google Cloud will continue adding new restore points to the regular backup chain and deleting outdated backups from the standard repository, according to the specified weekly scheduling settings. Veeam Backup for Google Cloud 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 Google Cloud will rebuild the full archive backup 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 Google Cloud transforms archive backup chains, see [Retention Policy for Archived Backups](#).



Step 7. Specify Processing Options

At the **Processing Options** step of the wizard, choose whether you want to use a staging server to perform backup operations. To learn how Veeam Backup for Google Cloud uses staging servers to protect Cloud SQL instances, see [SQL Backup](#).

IMPORTANT

If a Cloud SQL instance is configured to accept SSL connections, you will be able to back up the instance using a staging server only.

Protecting Instances Without Staging Server

To back up the selected Cloud SQL instances without a staging server, do the following:

1. In the **Staging server** section, select the **Use production instance** option.
2. In the **Authentication** section, click **Choose**.
3. In the **Choose an account** window, select a service account whose credentials will be used to authenticate against the production Cloud SQL instances. For an account to be displayed in the list of available accounts, it must be added to Veeam Backup for Google Cloud as described in section [Managing Cloud SQL Accounts](#).

If you select the [default IAM account](#), Veeam Backup for Google Cloud will access the instances using the service account that has been specified at [step 3](#).

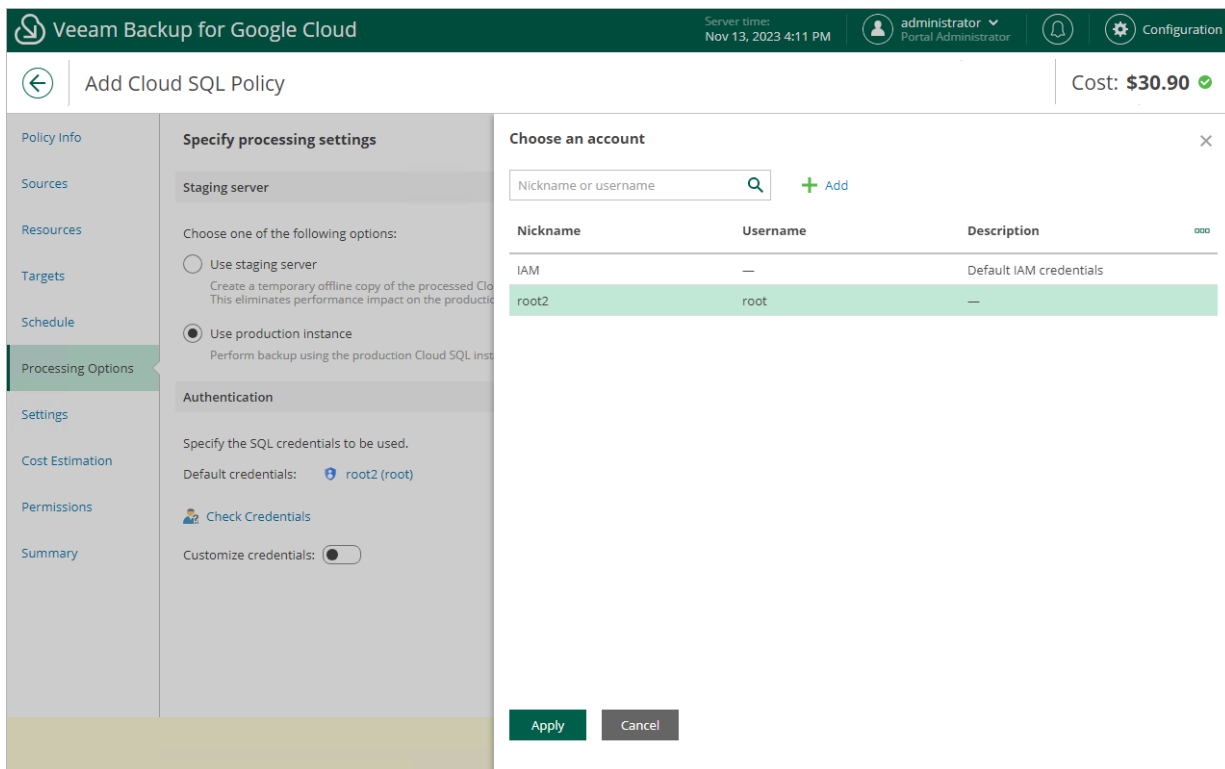
IMPORTANT

- If you select the default IAM account, the service account that will be used to access MySQL instances included into the backup scope must be granted the `cloudsql.instances.login` permission. For more information on Cloud SQL roles, see [Google Cloud documentation](#).
- Regardless of the option you choose, make sure that the selected account has either the permissions required to perform database dumping operations or the `cloudsqlsuperuser` role assigned.

By default, the selected account will be used to access all the instances added to the backup policy. You can also granularly specify credentials that Veeam Backup for Google Cloud will use to access specific production Cloud SQL instances. To do that, set the **Customize credentials** toggle to *On*, choose a resource for which you want to specify the credentials and click **Edit Credentials**.

TIP

It is recommended that you check whether the selected account exists on the protected instances. To do that, click **Check Credentials**.



Protecting Instances With Staging Server

To back up the selected Cloud SQL instances using a staging server, select the **Use staging server** option in the **Staging server** section.

When performing backup with a staging server, Veeam Backup for Google Cloud uses the default administrator account to send REST API requests to MySQL instances processed by the backup policy — that is why there is no need to specify credentials for authentication against the processed instances. However, Veeam Backup for Google Cloud is unable to use the default administrator account for PostgreSQL instances due to technical limitations — that is why you must also do the following if the backup policy protects PostgreSQL instances:

1. In the **Authentication** section, click **Choose**.
2. In the **Choose an account** window, select a service account that exists on all PostgreSQL instances processed by the policy. Veeam Backup for Google Cloud will create a user with the specified name and one-time password on the staging server to get access to the instance databases and to perform the backup operation.

By default, the selected account will be used to access all the instances added to the backup policy. You can also granularly specify credentials that Veeam Backup for Google Cloud will use to access specific PostgreSQL instances. To do that, set the **Customize credentials** toggle to *On*, choose a resource for which you want to specify the credentials and click **Edit Credentials**.

TIP

It is recommended that you check whether the selected account exists on the protected instances. To do that, click **Check Credentials**.

Veeam Backup for Google Cloud Server time: Nov 13, 2023 4:11 PM administrator Portal Administrator Configuration

← Add Cloud SQL Policy Cost: \$30.90 ✓

Policy Info

Sources

Resources

Targets

Schedule

Processing Options

Settings

Cost Estimation

Permissions

Summary

Specify processing settings

Staging server

Choose one of the following options:

Use staging server
Create a temporary offline copy of the processed Cloud SQL instance. This eliminates performance impact on the production instance.

Use production instance
Perform backup using the production Cloud SQL instance.

Authentication

Specify the SQL credentials to be used.

Default credentials: Choose...

Check Credentials

Customize credentials:

Choose an account

i When the Use staging server option is selected, Veeam Backup for Google Cloud requires only the username to perform backup.

+ Add

Nickname	Username	Description	☰
pgr	postgres	—	
pgr2	postgres	—	

Apply
Cancel

Step 8. Configure General Settings

At the **Settings** 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 Google Cloud to run the backup policy again if it fails on the first try, do the following:

1. In the **Retries** 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 15 minutes.

When retrying backup policies, Veeam Backup for Google Cloud processes only those Cloud SQL instances that failed to be backed up during the previous attempt.

NOTE

The automatic retry settings apply only to backup policies that run according to specific schedules — these settings do not apply to policies [started manually](#).

Health Check Settings

If you have enabled creation of image-level backups at [step 5](#), you can instruct Veeam Backup for Google Cloud to periodically perform a health check for backup restore points created by the backup policy. During the health check, Veeam Backup for Google Cloud performs an availability check for data blocks in the whole regular backup chain, and a cyclic redundancy check (CRC) for 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 Google Cloud does not verify archived restore points created by the policy.

To instruct Veeam Backup for Google Cloud to perform a monthly health check, 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 Google Cloud performs the health check during the last 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 Google Cloud 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 last policy session on Saturday.

Notification Settings

To instruct Veeam Backup for Google Cloud to send email notifications for the backup policy, do the following:

1. In the **Notifications** section of the step, set the **Enable notifications** toggle to *On*.
If you set the toggle to *Off*, Veeam Backup for Google Cloud 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.
3. Use the **Notify on** list to choose whether you want Veeam Backup for Google Cloud 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 Google Cloud 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 Google Cloud will send each notification to this recipient twice.

The screenshot shows the 'Add Cloud SQL Policy' configuration page in Veeam Backup for Google Cloud. The page is titled 'Add Cloud SQL Policy' and has a cost of '\$30.90'. The 'Settings' tab is selected in the left-hand navigation menu. The main content area is titled 'Configure retry and notification settings, and enable health check'. It includes a 'Retries' section with a checkbox for 'Automatically retry failed policy' set to '3' times. A 'Health check' section has a checkbox for 'Enable health check' set to 'On' and a 'Run on' dropdown set to 'First Monday of every month'. A 'Notifications' section has a checkbox for 'Enable notifications' set to 'On' and an email field containing 'john.smith@veeam.com'. There are checkboxes for 'Notify on' for 'Success', 'Warning', and 'Failure', and a checkbox for 'Suppress notifications until the last retry'. At the bottom, there are 'Previous', 'Next', and 'Cancel' buttons.

How Health Check Works

When Veeam Backup for Google Cloud saves a new 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 Cloud SQL instance data. When performing a health check, Veeam Backup for Google Cloud verifies availability of data blocks for each restore point and uses the saved values to ensure that the restore points being verified are consistent.

NOTE

Veeam Backup for Google Cloud performs the health check during the last 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 Google Cloud 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 last policy session on Saturday.

Veeam Backup for Google Cloud performs the health check in the following way:

1. As soon as the backup policy session completes successfully, Veeam Backup for Google Cloud starts the health check as a new session. For each restore point in the regular backup chain, Veeam Backup for Google Cloud calculates CRC values for backup metadata and compares them to the CRC values that were previously saved to the restore point. Veeam Backup for Google Cloud 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 Google Cloud 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 Google Cloud starts the health check.

2. If Veeam Backup for Google Cloud 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 Google Cloud performs the following operations:

- If the health check detects corrupted metadata in a full or an incremental restore point, Veeam Backup for Google Cloud marks the regular backup chain as corrupted in the configuration database. During the next backup policy session, Veeam Backup for Google Cloud copies the full Cloud SQL instance image, creates a new full restore point and starts a new backup chain in the backup repository.

NOTE

Veeam Backup for Google Cloud supports metadata check for encrypted backup chains unless the metadata is corrupted.

- If the health check detects corrupted data blocks in a full or an incremental restore point, Veeam Backup for Google Cloud marks the restore point that includes the corrupted data blocks and all subsequent incremental restore points as incomplete in the configuration database. During the next backup policy session, Veeam Backup for Google Cloud copies not only those data blocks that have changed since the previous backup session but also data blocks that have been corrupted. Veeam Backup for Google Cloud then 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 approximate monthly cost of Google Cloud services that Veeam Backup for Google Cloud will require to protect the Cloud SQL instances added to the backup policy. The total estimated cost includes the following:

- The cost of creating and maintaining cloud-native snapshots of the Cloud SQL instances.
For each Cloud SQL instance included in the backup policy, Veeam Backup for Google Cloud takes into account the amount of storage provisioned for the instance, 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 Cloud SQL instances.
For each Cloud SQL instance included in the backup policy, Veeam Backup for Google Cloud takes into account the amount of storage provisioned for the instance, the period of time during which restore points will be kept in the backup chain, the specified processing options, and the configured scheduling and health check settings.
- The cost of creating and storing in backup repositories archived backups of the Cloud SQL instances.
For each Cloud SQL instance included in the backup policy, Veeam Backup for Google Cloud takes into account the amount of storage provisioned for the instance, the period of time during which restore points will be kept in the archive backup chain, and the configured scheduling settings.
- The cost of transferring the Cloud SQL instance data between Google Cloud regions during data protection operations (for example, if a protected Cloud SQL instance and the target storage bucket reside in different regions).
- The cost of sending API requests to Google Cloud during data protection operations.

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, either select a backup repository that resides in the same region as Cloud SQL instances that you plan to back up, or select an archive repository that resides in the same region as the nearline or standard repository used to store regular backups.
- 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, modify the scheduling settings to run the backup policy less frequently, or specify an archive repository for long-term retention of restore points.

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 Cloud SQL Policy' configuration page in Veeam Backup for Google Cloud. The top navigation bar includes the Veeam logo, the product name, the server time (Nov 13, 2023 4:16 PM), the user (administrator), and a Configuration icon. The main header shows a back arrow, the page title 'Add Cloud SQL Policy', and the total cost 'Cost: \$30.90' with a green checkmark.

The left sidebar contains the following navigation items: Policy Info, Sources, Resources, Targets, Schedule, Processing Options, Settings, Cost Estimation (highlighted), Permissions, and Summary.

The main content area is titled 'Review cost estimation'. It contains the following text: 'The estimated cost takes into account the configured target settings, the specified scheduling options and the number of resources to protect.' and 'Note that Veeam Backup for Google Cloud makes predefined assumptions to calculate the cost, which means that the results should be used only as an approximation. For more information on how Veeam Backup for Google Cloud calculates the cost, see this Veeam KB article.'

Below the text are five cost breakdown items:

- Snapshots: \$17.46
- Backups: \$10.79
- Archives: \$0.00
- Traffic: \$2.49
- Transactions: \$0.16

A dark green banner displays the 'Estimated monthly cost: \$30.90' with a dollar sign icon.

Below the banner is a search bar for 'Instance' and an 'Export to...' dropdown menu.

Instance ↑	Snapshots	Backups	Archives	Traffic	Transactions	Total
mmfm...	\$5.82	\$3.60	\$0.00	\$0.83	\$0.05	\$10.30
mmfm...	\$5.82	\$3.60	\$0.00	\$0.83	\$0.05	\$10.30
prkr-sq...	\$5.82	\$3.60	\$0.00	\$0.83	\$0.05	\$10.30

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

Step 10. Check Required Permissions

At the **Permissions** step of the wizard, Veeam Backup for Google Cloud will verify whether the specified service account has all the necessary permissions required to perform data protection tasks for the selected project or folder. For more information on the required permissions, see [Service Account Permissions](#).

To see the list of missing permissions that must be granted to the service account in order to perform an operation, click the link in the **Details** column. You can grant the missing permissions to the service account [using the Google Cloud console](#) or instruct Veeam Backup for Google Cloud to do it:

- To grant the missing permissions manually, click **Download Script**. Veeam Backup for Google Cloud will generate a gcloud script that you can run in the Google Cloud console to assign all the necessary permissions to the service account.

The account under which you run the script must have the permissions required both to get and set project IAM policies and to create custom IAM roles (for example, it can have the *iam.securityAdmin* and *iam.roleAdmin* roles assigned). To learn what permissions and roles are required to create custom roles in IAM, see [Google Cloud documentation](#).

- To let Veeam Backup for Google Cloud grant the missing permissions automatically, click **Grant** and then click **Sign in with Google** in the **Grant permissions** window. You will be redirected to the OAuth consent screen authorization page. Sign in using credentials of a Google account that will be used to grant the permissions.

The account under which you sign in to Google Cloud must have the permissions required both to get and set project IAM policies and to create custom IAM roles (for example, it can have the *iam.securityAdmin* and *iam.roleAdmin* roles assigned). To learn what permissions and roles are required to create service account, see [Google Cloud documentation](#).

NOTE

For Veeam Backup for Google Cloud to be able to authorize in Google Cloud, the OAuth consent screen must be configured as described in section [Registering Applications](#). Note that Veeam Backup for Google Cloud does not store in the configuration database the provided Google account credentials and access tokens received during authorization.

To make sure that the missing permissions have been successfully granted, click **Recheck**.

Check permissions
Verify whether all the required permissions are granted.

[Recheck](#) [Download Script](#)

Check	Result	Details
Cloud SQL Snapshot	Passed	All the required permissions are ...
Cloud SQL Backup	Passed	All the required permissions are ...
Repository	Passed	All the required permissions are ...
Worker	Passed	All the required permissions are ...

Previous Next Cancel

Step 11. Finish Working with Wizard

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

The screenshot shows the 'Add Cloud SQL Policy' wizard in the 'Summary' step. The interface includes a top navigation bar with the Veeam logo, server time (Nov 13, 2023 4:17 PM), user (administrator), and configuration icons. The main content area is titled 'Review configured settings' and contains a 'Copy to Clipboard' button. The settings are organized into several sections:

- General settings:** Name: europe-policy; Description: protecting instances in other EU regions; Service account: veeam-1649186685-sa@rmd-backup-2.iam.gserviceaccount.com; Project: veeam-rnd-backup-2; Folder: —
- Protected resources:** SQL Engine: MySQL; Regions: 5 regions; Instances: 3 instances; Labels: —; Exclusions: —
- Backup settings:** Backups enabled: Yes; Backup repository: backup-readd; Archive repository: archive; Storage bucket: dr-us-west3; Daily retention: Create 4 backups per day and keep for 21 days; Weekly retention: Keep weekly backup for 21 days (6 backups excluded); Monthly retention: Keep monthly backup for 30 days (10 backups excluded); Yearly retention: 1 year
- Snapshot settings:** Snapshots enabled: Yes; Daily retention: Create 7 snapshots per day and keep 7 most recent snapshots; Weekly retention: Keep 7 weekly snapshots; Monthly retention: Keep 7 monthly snapshots
- Other settings:** Automatic retries enabled: Yes; Notifications enabled: Yes; Health check enabled: Yes

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

Creating Snapshots Manually

Veeam Backup for Google Cloud allows you to manually create snapshots of Cloud SQL instances. Each snapshot is saved to the multi-regional location closest to the region in which the original Cloud SQL instance resides.

NOTE

Veeam Backup for Google Cloud 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 Google Cloud Storage unless you remove them manually, as described in section [Removing Backups and Snapshots](#).

To manually create a cloud-native snapshot of a Cloud SQL instance, do the following:


1. Navigate to **Resources > Cloud SQL**.

2. Select the necessary instance and click **Take Snapshot Now**.

For a Cloud SQL instance to be displayed in the list of available instances, it must reside in any of the regions added to a backup policy as described in section [Creating Backup Policies](#).

3. In the **Take Snapshot Now** window, select a service account whose permissions Veeam Backup for Google Cloud will use to create the snapshot, and click **Take Snapshot**.

For a service account to be displayed in the **Service account** drop-down list, it must be added to Veeam Backup for Google Cloud as described in section [Adding Service Accounts](#), and must be assigned the *Cloud SQL Instances Snapshot* operational role as described in section [Adding Projects and Folders](#).



Performing Spanner Backup

One backup policy can be used to process one or more Cloud Spanner instances within one Google Cloud project or folder. The scope of data that you can protect in a project or folder is limited by permissions of a service account that is specified in the backup policy settings.

Before you create a Cloud Spanner backup policy, check the following prerequisites:

- If you plan to create image-level backups of Cloud Spanner 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 the backup policy results, configure SMTP server settings first. For more information, see [Configuring Global Notification Settings](#).

To schedule data protection tasks to run automatically, [create backup policies](#). For each protected Cloud Spanner instance, you can also [take a cloud-native snapshot manually](#) when needed.

Creating Backup Policies

To create a backup policy, do the following:

1. [Launch the Add Cloud Spanner Policy wizard](#).
2. [Specify a backup policy name and description](#).
3. [Choose a project to which Cloud Spanner instances that you plan to back up belong](#).
4. [Configure backup source settings](#).
5. [Configure backup target settings](#).
6. [Create a schedule for the backup policy](#).
7. [Specify automatic retry, health check and notification settings for the backup policy](#).
8. [Review the estimated cost of protecting the selected Cloud Spanner instances](#).
9. [Check the required permissions](#).
10. [Finish working with the wizard](#).

Step 1. Launch Add Cloud Spanner Policy Wizard

To launch the **Add Cloud Spanner Policy** wizard, do the following:

1. Navigate to **Policies > Cloud Spanner**.
2. Click **Add**.

The screenshot displays the Veeam Backup for Google Cloud interface. The top navigation bar includes the Veeam logo, the text "Veeam Backup for Google Cloud", the server time "Nov 13, 2023 6:41 PM", and the user "administrator Portal Administrator". The left sidebar shows a navigation menu with "Monitoring" (Overview, Resources) and "Management" (Policies, Protected Data, Session Logs). The main content area is titled "Cloud Spanner" and features a search bar, a filter dropdown, and a toolbar with actions like Start, Stop, Enable, Add, Edit, Priority, View Info, Remove, Advanced, and Export to... Below this is a table of policies with columns for Priority, Policy, Snapshots, Backups, Archives, Last Run, and Last Duration. The table shows four policies, with the last one being in a "Running" state. Below the table are two empty sections: "Instances" and "Sessions".

Priority	Policy	Snapshots	Backups	Archives	Last Run	Last Duration
1	spanner-pol...	Success	Success	Not Configured	12/01/2023 12:04:32 PM	40 min 19 sec
2	spanner-pol...	Success	Not Configured	Not Configured	11/21/2023 1:05:45 PM	53 min 34 sec
3	spanner-cm...	Success	Not Configured	Not Configured	11/21/2023 11:00:14 AM	46 min 18 sec
4	spanner-pol...	Running	Running	Never Executed	12/06/2023 10:25:01 AM	—

Step 2. Specify Backup Policy Name and Description

At the **Policy Info** step of the wizard, use the **Name** and **Description** fields to enter a name for the new backup policy and to provide a description for future reference. The policy name can contain only uppercase Latin letters, lowercase Latin letters, numeric characters and hyphens; the maximum length of the name is 127 characters.

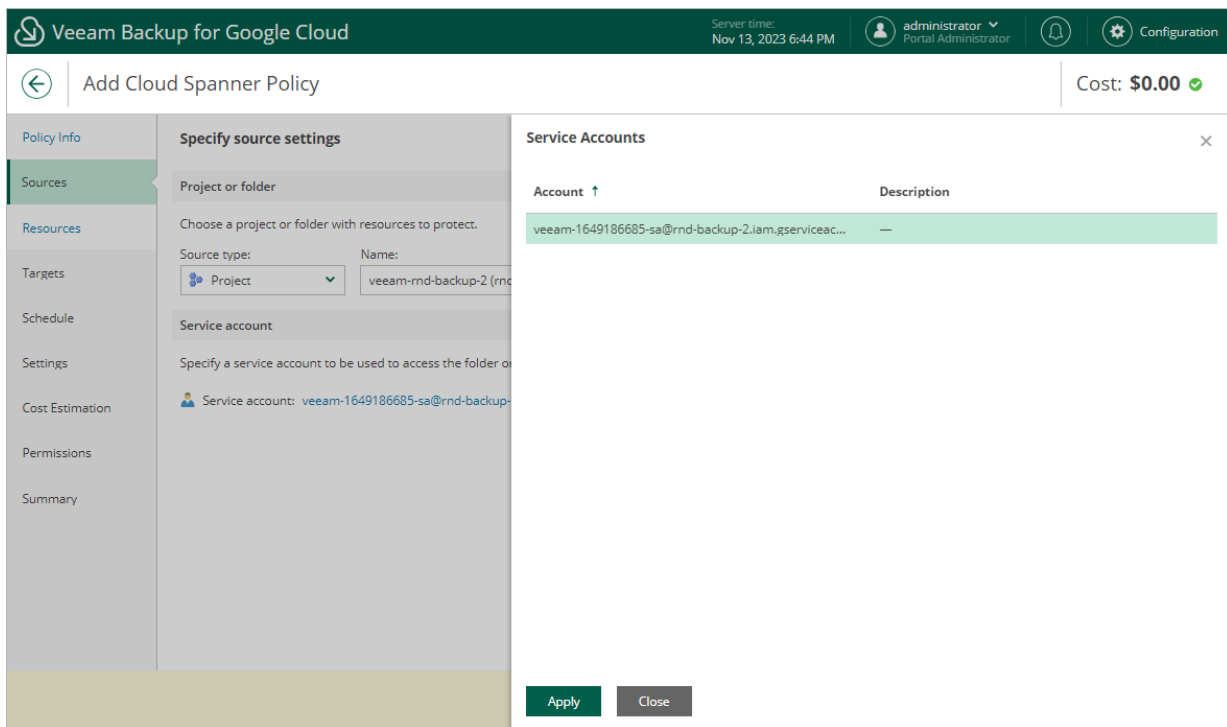
The screenshot shows the 'Add Cloud Spanner Policy' wizard in the Veeam Backup for Google Cloud interface. The current step is 'Specify policy name and description'. The interface includes a top navigation bar with the Veeam logo, server time (Nov 13, 2023 6:44 PM), user information (administrator, Portal Administrator), and a Configuration icon. A breadcrumb trail shows a back arrow and the title 'Add Cloud Spanner Policy'. On the right, the cost is displayed as '\$0.00' with a green checkmark. A left-hand sidebar lists the wizard steps: Policy Info (selected), Sources, Resources, Targets, Schedule, Settings, Cost Estimation, Permissions, 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, there are two input fields: 'Name:' with the value 'eu-policy' and 'Description:' with the value 'protecting Spanner instances in EU regions'. At the bottom of the wizard, there are 'Next' and 'Cancel' buttons.

Step 3. Specify Project

At the **Sources** step of the wizard, choose a project or a folder with a project that manages resources that you want to protect, and specify a service account that will be used to access the project or folder.

For a project or folder to be displayed in the list of available entities, it must be added to Veeam Backup for Google Cloud as described in section [Adding Projects and Folders](#). If you have not added the necessary entity to Veeam Backup for Google Cloud beforehand, you can do it without closing the **Add Cloud Spanner Policy** wizard. To add an entity, click **Add** and complete the **Add Projects and Folders** wizard.

For a service account to be displayed in the list of available accounts, it must be added to Veeam Backup for Google Cloud as described in section [Adding Service Accounts](#), and must be assigned the *Cloud Spanner Instances Snapshot* and *Backup* operational roles as described in section [Adding Projects and Folders](#).



Step 4. Configure Backup Source Settings

At the **Resources** step of the wizard, specify the following backup source settings:

1. [Choose regions in which Cloud Spanner instances that you plan to back up reside.](#)
2. [Select Cloud Spanner instances to back up.](#)

Step 4a. Choose Regions

In the **Regions** section of the **Resources** step of the wizard, choose regions in which Cloud Spanner instances that you want to protect reside.

1. Click **Choose regions**.
2. In the **Choose regions** window, select the necessary regions from the **Available Regions** list, and click **Add**.
3. To save changes made to the backup policy settings, click **Apply**.

IMPORTANT

If you want to protect a multi-regional Cloud Spanner instance, you must choose regions where its read-write or read-only replicas are located; witness replicas do not participate in the backup process due to [Google Cloud limitations](#).

The screenshot shows the 'Add Cloud Spanner Policy' wizard in the Veeam Backup for Google Cloud interface. The 'Resources' step is active, and the 'Choose regions' sub-step is displayed. The interface is divided into three main sections:

- Left Panel (Navigation):** A vertical sidebar with options: Policy Info, Sources, Resources (selected), Targets, Schedule, Settings, Cost Estimation, Permissions, and Summary.
- Center Panel (Specify resources to protect):** Under the 'Region' section, it says 'Specify one or more regions.' and has a 'Choose regions...' button. Under the 'Resources' section, it says 'Specify resources to protect or exclude.' and has 'Choose resources to protect...' and 'Exclude resources...' buttons.
- Right Panel (Choose regions):** A modal window titled 'Choose regions' with a close button (X). It contains:
 - Available Regions (36):** A scrollable list of regions including asia-southeast2 (Jakarta), australia-southeast1 (Sydney), europe-central2 (Warsaw), europe-west8 (Milan) (highlighted), and others.
 - Action Buttons:** 'Add' and 'Remove' buttons.
 - Policy Regions (3):** A list of selected regions: europe-west2 (London), europe-west3 (Frankfurt), and europe-west4 (Netherlands).
 - Footer Buttons:** 'Apply' and 'Cancel' buttons.

At the top right of the wizard, the cost is shown as '\$0.00' with a green checkmark. The top navigation bar includes the Veeam logo, server time (Nov 13, 2023 6:45 PM), and user information (administrator, Portal Administrator).

Step 4b. Select Cloud Spanner Instances

In the **Resources** section of the **Resources** step of the wizard, specify the backup scope — select Cloud Spanner instances that Veeam Backup for Google Cloud will back up:

1. Click **Choose resources to protect**.
2. In the **Choose resources** window, choose whether you want to back up all Cloud Spanner instances from the regions selected at [step 4a](#), or only specific Cloud Spanner instances.

If you select the **All resources** option, Veeam Backup for Google Cloud will regularly check for new Cloud Spanner instances launched in the selected regions and automatically update the backup policy settings to include these instances in the backup scope.

If you select the **Specific resources** option, you must also specify the instances explicitly:

- a. Use the **Resource type** drop-down list to choose whether you want to add individual Cloud Spanner instances or Google Cloud labels to the backup scope.

If you select the **Label** option, Veeam Backup for Google Cloud will back up only those Cloud Spanner instances that reside in the selected regions under specific labels.

- b. Use the **Instance\Label** list to find the necessary resource, and then click **Add to Protected** to add the resource to the backup scope.

For a resource to be displayed in the list of available resources, it must reside in a region that has ever been specified in any backup policy. Otherwise, the only option to discover available resources is to click **Browse** and wait for Veeam Backup for Google Cloud to populate the resource list.

TIP

You can simultaneously add multiple resources to the backup scope. To do that, click **Browse**, select check boxes next to the necessary Cloud Spanner instances or labels 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 Google Cloud will update the resource list.

If you add a label to the backup scope, Veeam Backup for Google Cloud will regularly check for new Cloud Spanner instances assigned the added label and automatically update the backup policy settings to include these instances in the scope. However, this applies only to Cloud Spanner instances from the regions selected at [step 4a](#). If you select a label assigned to Cloud Spanner instances from other regions, these instances will not be protected by the backup policy. To work around the issue, either go back to step 4a 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 **Specific 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 **Exclude resources** and specify the Cloud Spanner instances or labels that you do not want to back up — the procedure is the same as described for including resources in the backup scope.

Consider that if a resource appears both in the list of included and excluded resources, Veeam Backup for Google Cloud will still not process the resource because the list of excluded resources has a higher priority.

The screenshot shows the 'Add Cloud Spanner Policy' dialog in Veeam Backup for Google Cloud. The interface is divided into several sections:

- Header:** Veeam Backup for Google Cloud, Server time: Nov 13, 2023 6:45 PM, administrator Portal Administrator, Configuration.
- Navigation:** Policy Info, Sources, Resources (selected), Targets, Schedule, Settings, Cost Estimation, Permissions, Summary.
- Specify resources to protect:**
 - Region:** Specify one or more regions. 3 regions selected.
 - Resources:** Specify resources to protect or exclude. Options: Choose resources to protect..., Exclude resources...
- Choose resources:**
 - Radio buttons for All resources and Specific resources (selected).
 - Buttons: Rescan, Browse...
 - Resource type: Instance (dropdown), Instance: Select... (dropdown), Add to Protected.
 - Protected resources (1): InstanceValue (input), Remove (button).
 - Table with columns: Resource, Project, Region\Value.
 - Selected: 0 of 1.
 - Table row: InstanceValue, prkr-spanne..., veeam-rnd-backu..., europe-west3 (Frankfurt).
 - Buttons: Apply, Cancel.

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 instruct Veeam Backup for Google Cloud to create image-level backups of the selected Cloud Spanner instances:

1. Set the **Enable backups** toggle to *On*.
2. Click **Choose repository**.
3. In the **Choose repository** window, select a backup repository where the created image-level backups will be stored.

For a backup repository to be displayed in the **Repository** list, it must be added to Veeam Backup for Google Cloud as described in section [Adding Backup Repositories](#). The **Repository** list shows only backup repositories of the *Standard* and *Nearline* storage classes.

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

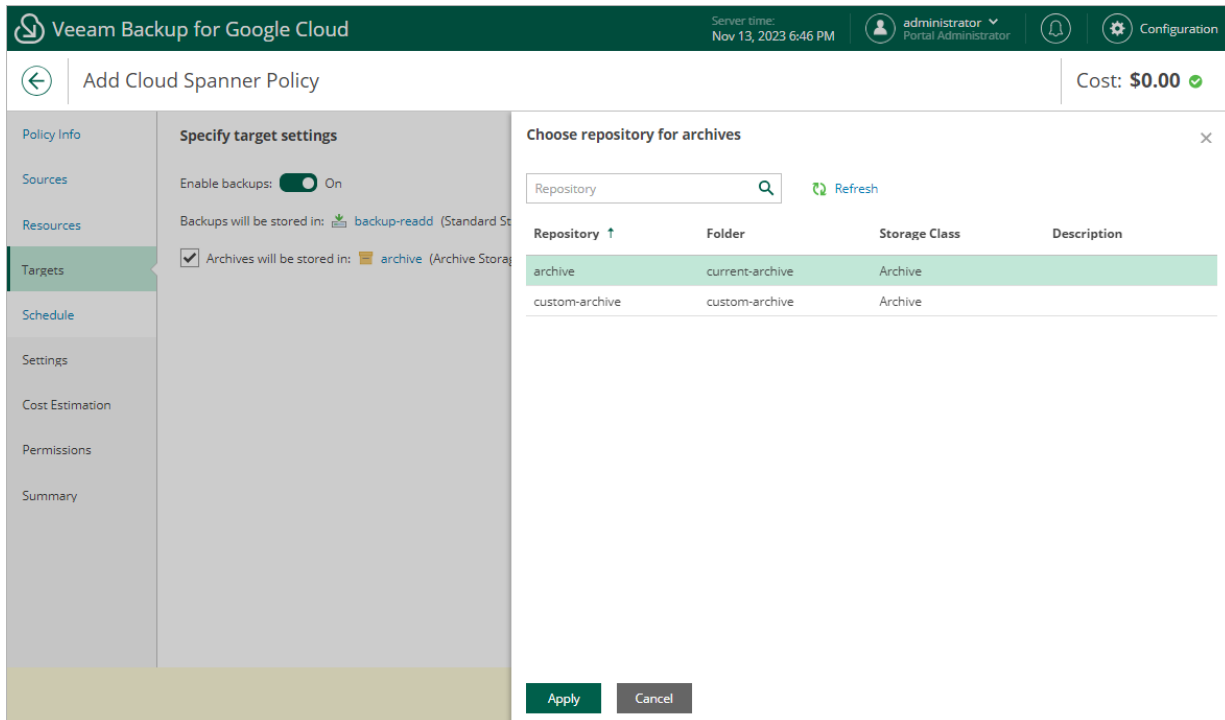
You can also enable the backup archiving mechanism to instruct Veeam Backup for Google Cloud to store backed-up data in a low-cost, long-term archive storage:

1. Select the **Archives will be stored in** check box.
2. Click **Choose repository**.
3. In the **Choose repository** window, select a backup repository where the archived data will be stored.

For a backup repository to be displayed in the **Repository** list, it must be added to Veeam Backup for Google Cloud as described in section [Adding Backup Repositories](#). The **Repository** list shows only backup repositories of the *Archive* storage class.

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

For more information on the backup archiving mechanism, see [Enabling Backup Archiving](#).



The screenshot shows the Veeam Backup for Google Cloud interface. The top navigation bar includes the Veeam logo, the product name, server time (Nov 13, 2023 6:46 PM), user (administrator), and configuration options. The main window is titled 'Add Cloud Spanner Policy' and shows a sidebar with navigation options: Policy Info, Sources, Resources, Targets (selected), Schedule, Settings, Cost Estimation, Permissions, and Summary. The 'Specify target settings' section is visible, showing 'Enable backups' set to 'On' and 'Archives will be stored in' checked. A 'Choose repository for archives' dialog is open, displaying a search bar and a table of repositories.

Repository	Folder	Storage Class	Description
archive	current-archive	Archive	
custom-archive	custom-archive	Archive	

Step 6. Specify Policy Scheduling Options

At the **Schedule** step of the wizard, you can instruct Veeam Backup for Google Cloud to start the backup policy automatically according to a specific backup schedule. The backup schedule defines how often data of the Cloud Spanner instances added to the backup policy will be backed up.

To help you implement a comprehensive backup strategy, Veeam Backup for Google Cloud 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](#).

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** section, select hours when the backup policy will create cloud-native snapshots and image-level backups. Use the **Run at** drop-down list to choose whether you want the backup policy to run every day, on weekdays (Monday through Friday) or on specific days.

If you want to protect Cloud Spanner 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.

NOTES


- Veeam Backup for Google Cloud does not create image-level backups independently from cloud-native snapshots. That is why when you select hours for image-level backups, the same hours are automatically selected for cloud-native snapshots. To learn how Veeam Backup for Google Cloud performs backup, see [Spanner Backup](#).
- In Google Cloud Storage, Cloud Spanner snapshots are stored for a period of up to one year. If you need to keep snapshots for a longer period of time, you can export the databases to a Cloud Storage bucket. To learn how to do this, see [Google Cloud documentation](#).

3. In the **Configure daily retention** section, configure retention policy settings for the daily schedule:
 - For cloud-native snapshots, specify the number of restore points that you want to keep in a snapshot chain.

If the restore point limit is exceeded, Veeam Backup for Google Cloud removes the earliest restore point from the chain. For more information, see [Spanner 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 Google Cloud removes the restore point from the chain. For more information, see [Spanner Backup Retention](#).

4. 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** section, select days when the backup policy will create cloud-native snapshots and image-level backups. Use the **Create restore points** at drop-down list to schedule a specific time for the backup policy to run.

NOTE

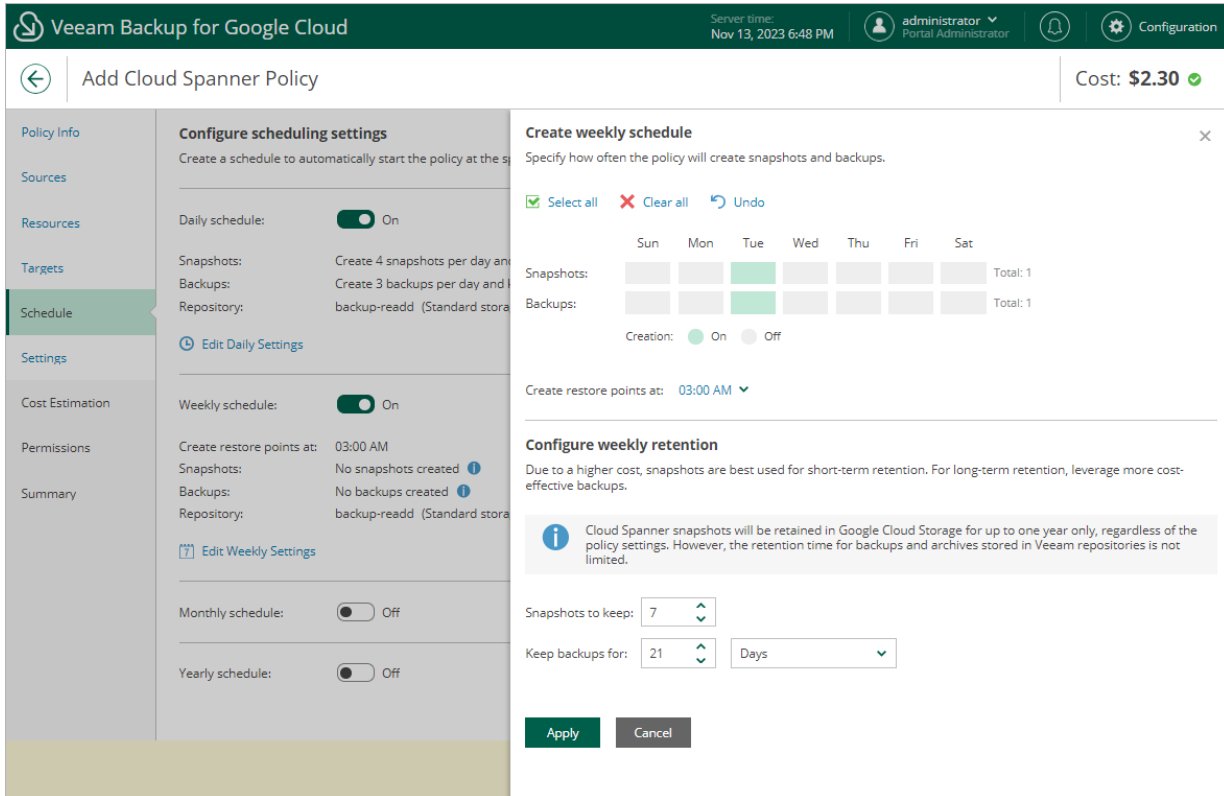
Veeam Backup for Google Cloud does not create image-level backups independently from cloud-native snapshots. That is why when you select days for image-level backups, the same days are automatically selected for cloud-native snapshots. To learn how Veeam Backup for Google Cloud performs backup, see [Spanner Backup](#).

3. In the **Configure weekly retention** section, configure retention policy settings for the weekly schedule:
 - For cloud-native snapshots, specify the number of restore points that you want to keep in a snapshot chain.

If the restore point limit is exceeded, Veeam Backup for Google Cloud removes the earliest restore point from the chain. For more information, see [Spanner 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 Google Cloud removes the restore point from the chain. For more information, see [Spanner Backup Retention](#).

4. 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 **Choose monthly backup target** section, select months when the backup policy will create cloud-native snapshots and image-level backups. Use the **Create restore points at** and **Run on** drop-down lists to schedule a specific time and day for the backup policy to run.

NOTE


Veeam Backup for Google Cloud does not create image-level backups independently from cloud-native snapshots. That is why when you select months for image-level backups, the same months are automatically selected for cloud-native snapshots. To learn how Veeam Backup for Google Cloud performs backup, see [Spanner Backup](#).

3. In the **Configure monthly retention** section, configure retention policy settings for the monthly schedule:
 - For cloud-native snapshots, specify the number of restore points that you want to keep in a snapshot chain.
If the restore point limit is exceeded, Veeam Backup for Google Cloud removes the earliest restore point from the chain. For more information, see [Retention Policy for Snapshots](#).
 - 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 Google Cloud removes the restore point from the chain. For more information, see [Retention Policy for Backups](#).

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

TIP

If you have enabled backup archiving at the **Targets** step of the wizard, and want to store monthly backups in an archive backup repository, set the **Send backups to archive** toggle to *On*, and follow the instructions provided in section [Enabling Backup Archiving](#).



Specifying Yearly Schedule

[This step applies only if you have instructed Veeam Backup for Google Cloud 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. In the **Create yearly schedule** section, specify a day, month and time when the backup policy will create image-level backups.

For example, if you select *First, Tuesday, March* and *03:00 AM*, the backup policy will run every first Tuesday of March at 03:00 AM.


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 Google Cloud removes the restore point from the chain. For more information, see [Retention Policy for Backups](#).

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

TIP

If you have enabled backup archiving at the **Targets** step of the wizard, and want to store yearly backups in an archive backup repository, set the **Send backups to archive** toggle to *On*, and follow the instructions provided in section [Enabling Backup Archiving](#).



Enabling Harmonized Scheduling

When you combine multiple types of schedules, Veeam Backup for Google Cloud 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 Google Cloud can keep restore points created according to a daily, weekly or monthly schedule for longer periods of time:

- Cloud-native snapshots can be kept for weeks and months.
- Image-level backups can be kept for weeks, months and years.

For Veeam Backup for Google Cloud 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 retaining restore points. In terms of harmonized scheduling, Veeam Backup for Google Cloud 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 (Daily) flag is used to mark restore points created daily, (Weekly) – weekly, (Monthly) – monthly, and (Yearly) – yearly. Veeam Backup for Google Cloud 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 Google Cloud 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 retain one of the created snapshots for 2 weeks. In this case, you create 2 schedules when configuring the backup policy settings — daily and weekly:

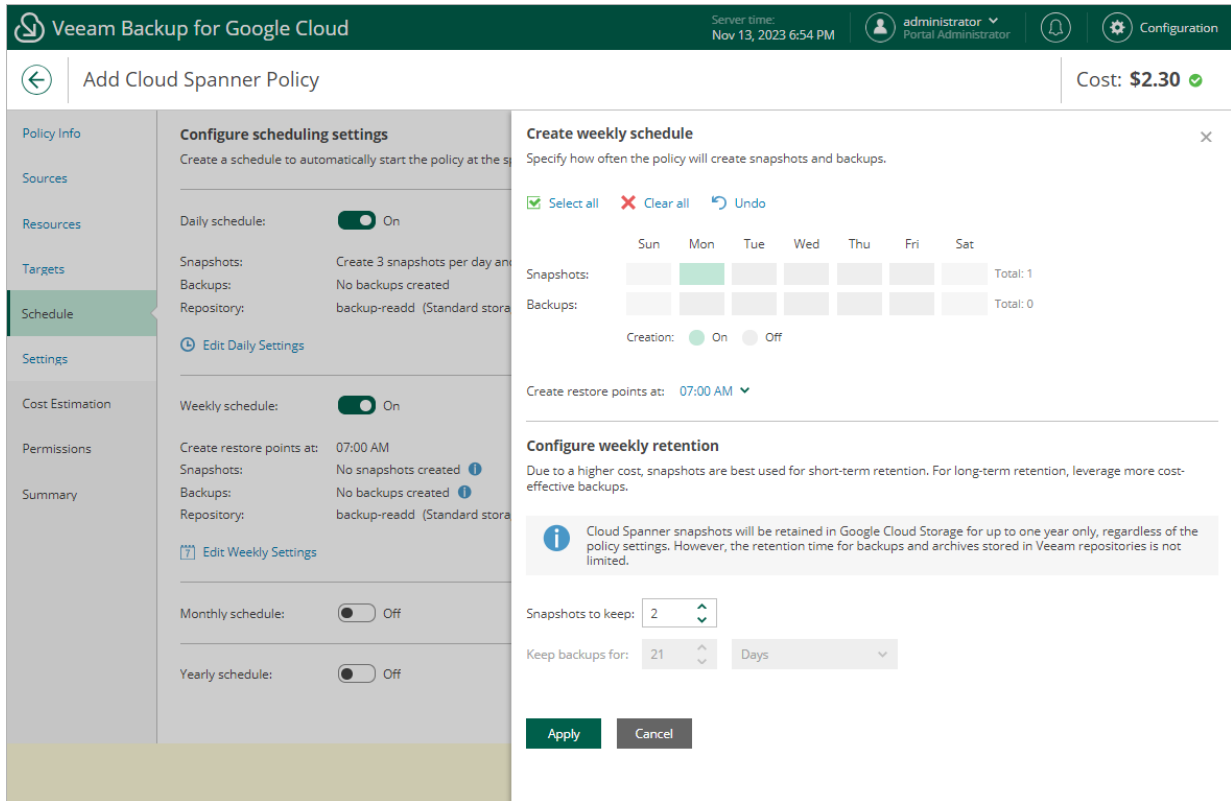
1. In the daily scheduling settings, you select hours and days when snapshots will be created (for example, *7:00 AM, 9:00 AM, and 11:00 AM; Weekdays*), and specify a number of daily restore points to retain (for example, *3*).

Veeam Backup for Google Cloud will propagate these settings to the schedule with a lower frequency (which is the weekly schedule in our example).

The screenshot displays the Veeam Backup for Google Cloud configuration interface. At the top, the server time is Nov 13, 2023 6:53 PM, and the user is administrator. The main window is titled 'Add Cloud Spanner Policy' with a cost of \$2.30. The left sidebar shows navigation options: Policy Info, Sources, Resources, Targets, Schedule (selected), Settings, Cost Estimation, Permissions, and Summary. The 'Configure scheduling settings' panel is active, showing options for Daily, Weekly, Monthly, and Yearly schedules, all currently disabled. The 'Create daily schedule' dialog is open, allowing the user to specify how often the policy will create snapshots and backups. It features a calendar grid for weekdays (Monday to Friday) with a time slot selector (AM/PM) and a grid for selecting snapshot and backup times. The 'Snapshots' row shows 3 snapshots scheduled at 7 AM, 9 AM, and 11 AM, with a total of 3 snapshots per hour. The 'Backups' row shows 0 backups. The 'Run at' dropdown is set to 'Weekdays'. Below the calendar, the 'Configure daily retention' section provides information about retention in Google Cloud Storage and allows setting 'Snapshots to keep' (3) and 'Keep backups for' (21 Days). 'Apply' and 'Cancel' buttons are at the bottom.


- In the weekly scheduling settings, you specify which one of the snapshots created by the daily schedule will be kept 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 Google Cloud will create cloud-native snapshots in the following way:

- On the first weekday (Monday), a backup session will start at 7:00 AM to create the first restore point. The restore point will be marked with the (Daily) flag as it was created according to the daily schedule. Since *7:00 AM, Monday* is specified in the weekly scheduling settings, Veeam Backup for Google Cloud will assign the (Weekly) flag to this 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 (Daily) flag.




- On the next weekday (Tuesday), after a backup session runs at 7:00 AM, the created restore point will be marked with the (Daily) flag.


By the moment the backup session completes, the number of restore points with the (Daily) flag will exceed the retention limit specified in the daily scheduling settings. However, Veeam Backup for Google Cloud will not remove the earliest restore point (7:00 AM, Monday) with the (Daily) flag from the snapshot chain as this restore point is also marked with a flag of a less-frequent schedule. Instead, Veeam Backup for Google Cloud will unassign the (Daily) 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 (Daily) flag will exceed the retention limit once again. Veeam Backup for Google Cloud 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 Google Cloud 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 Google Cloud will unassign the (Weekly) flag from the earliest weekly restore point. Since no other flags are assigned to this restore point, Veeam Backup for Google Cloud 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 Google Cloud to store backed-up data in the low-cost, long-term Google Cloud archival storage. 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 Google Cloud standard storage.

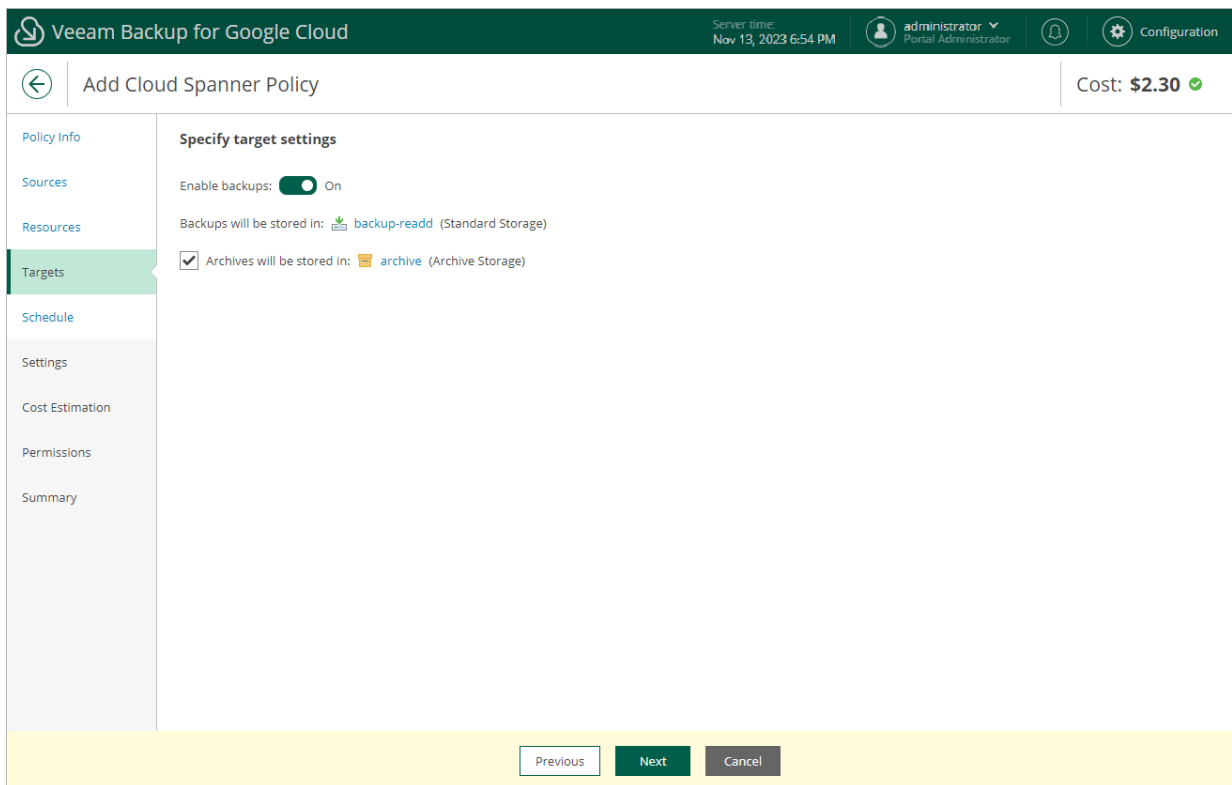
With backup archiving, Veeam Backup for Google Cloud can retain backups created according to a daily, weekly or monthly schedule for longer periods of time:

- To enable monthly archiving, you must configure a daily or a weekly backup schedule (or both).
- To enable yearly archiving, you must configure a daily, a weekly or a monthly backup schedule (or all three).

For Veeam Backup for Google Cloud to use the archiving mechanism, there must be specified at least 2 different schedules: one schedule will control the regular creation of backups, while another schedule will control the process of copying backups to an archive repository. Backup chains created according to these two schedules will be completely different – for more information, see [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 repository for 3 weeks, and also to keep backups created once in 2 months in an archive repository for a year. In this case, you create 2 schedules when configuring the backup policy settings – weekly and monthly:

1. In the policy target settings, you set the **Enable backups** toggle to *On*, select a standard repository that will store regular backups, and select an archive repository that will store archived data.



- In the weekly scheduling settings, you select hours and days when backups will be created (for example, *8:00 AM, Wednesday*), and specify a number of days for which Veeam Backup for Google Cloud will retain backups (for example, *21 days*).

Veeam Backup for Google Cloud 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 Google Cloud will create archive backups, and choose for how long you want to keep the created backups in the archive repository.

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

IMPORTANT

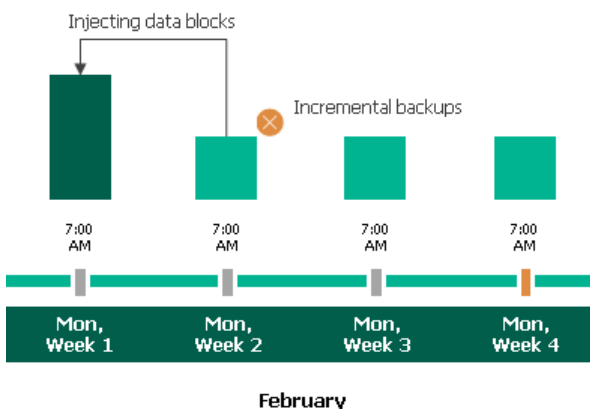
- When you enable backup archiving, you become no longer able to create a schedule of the same frequency for regular backups. By design, these two functionalities are mutually exclusive.
- If you enable backup archiving, it is recommended that you set the **Snapshots to keep** value to 0, to reduce unexpected snapshot charges.
- If you enable backup archiving, it is recommended that you set the **Keep archives for** value to at least *12 months* (or *365 days*), since the minimum storage duration of the Google Cloud archival storage is 365 days.
- If you select the **On day** option, **harmonized scheduling** cannot be guaranteed. Plus, to support the **On day** option, Veeam Backup for Google Cloud 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 Google Cloud Storage during approximately 24 hours, to reduce unexpected infrastructure charges.



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

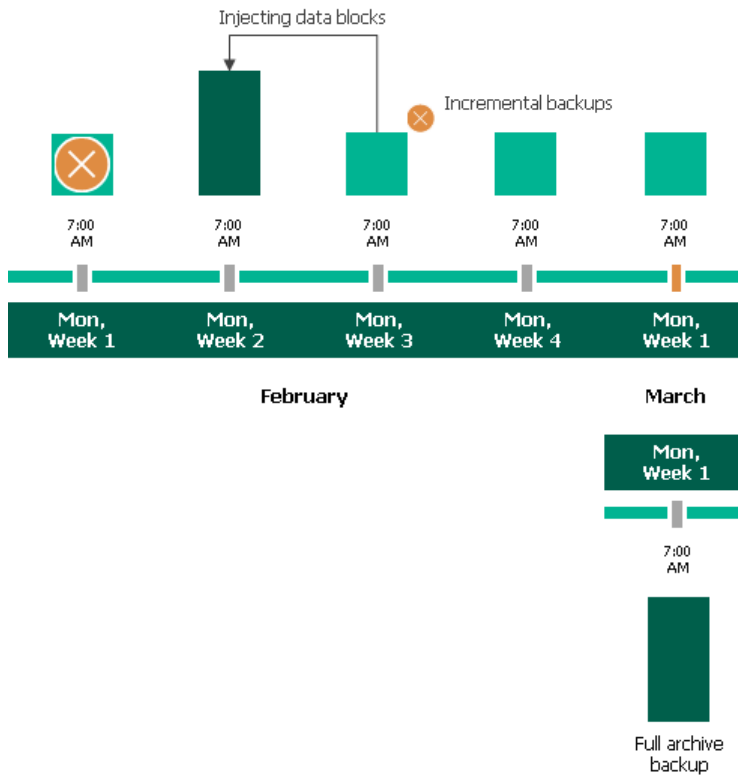
1. On the first Monday of February, a backup session will start at 7:00 AM to create the first restore point in the regular backup chain. Veeam Backup for Google Cloud will store this restore point as a full backup in the standard repository.
2. On the second and third Mondays of February, Veeam Backup for Google Cloud will create restore points at 7:00 AM and add them to the regular backup chain in the standard repository.
3. On the fourth Monday of February, Veeam Backup for Google Cloud will create a new restore point at 7:00 AM. By the moment the backup session completes, the earliest restore point in the regular backup chain will get older than the specified retention limit. That is why Veeam Backup for Google Cloud will rebuild the full backup and remove from the chain the restore point created on the first Monday.

For more information on how Veeam Backup for Google Cloud transforms regular backup chains, see [Retention Policy for Backups](#).




- On the first Monday of March, a backup session will start at 7:00 AM to create another restore point in the regular backup chain. At the same time, the earliest restore point in the regular backup chain will get older than the specified retention limit again. That is why Veeam Backup for Google Cloud will rebuild the full backup 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 the data from the regular backup chain. Veeam Backup for Google Cloud will copy this restore point as a full archive backup to the archive repository.



- Up to May, Veeam Backup for Google Cloud will continue adding new restore points to the regular backup chain and deleting outdated backups from the standard repository, according to the specified weekly scheduling settings.

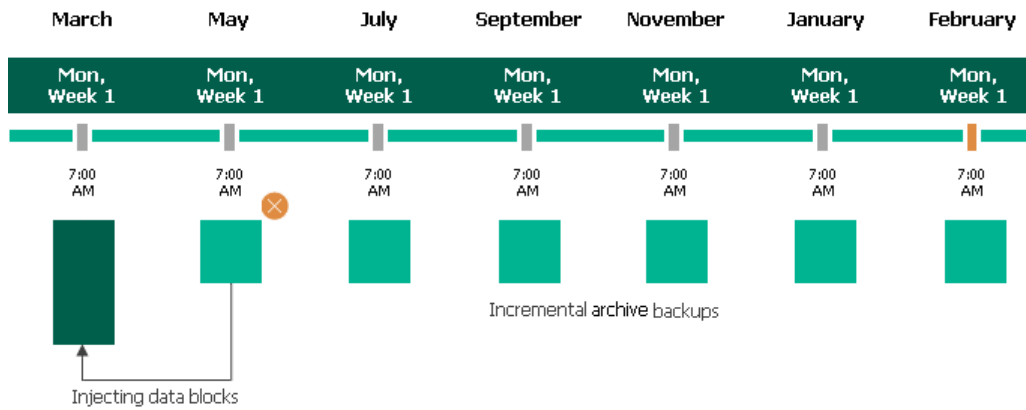
On the first Monday of May, an archive session will create a restore point with only that data that has changed since the previous archive session in March. Veeam Backup for Google Cloud will copy this restore point as an incremental archive backup to the archive repository.



- Up to the first Monday of March of the next year, Veeam Backup for Google Cloud will continue adding new restore points to the regular backup chain and deleting outdated backups from the standard repository, according to the specified weekly scheduling settings. Veeam Backup for Google Cloud 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 Google Cloud will rebuild the full archive backup 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 Google Cloud transforms archive backup chains, see [Retention Policy for Archived Backups](#).



Step 7. Configure General Settings

At the **Settings** 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 Google Cloud to run the backup policy again if it fails on the first try, do the following:

1. In the **Retries** 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 15 minutes.

When retrying backup policies, Veeam Backup for Google Cloud processes only those Cloud Spanner instances that failed to be backed up during the previous attempt.

NOTE

The automatic retry settings apply only to backup policies that run according to specific schedules — these settings do not apply to policies [started manually](#).

Health Check Settings

If you have enabled creation of image-level backups at [step 5](#), you can instruct Veeam Backup for Google Cloud to periodically perform a health check for backup restore points created by the backup policy. During the health check, Veeam Backup for Google Cloud performs an availability check for data blocks in the whole regular backup chain, and a cyclic redundancy check (CRC) for 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 Google Cloud does not verify archived restore points created by the policy.

To instruct Veeam Backup for Google Cloud to perform a monthly health check, 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 Google Cloud performs the health check during the last 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 Google Cloud 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 last policy session on Saturday.

Notification Settings

To instruct Veeam Backup for Google Cloud to send email notifications for the backup policy, do the following:

1. In the **Notifications** section of the step, set the **Enable notifications** toggle to *On*.
If you set the toggle to *Off*, Veeam Backup for Google Cloud 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.
3. Use the **Notify on** list to choose whether you want Veeam Backup for Google Cloud 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 Google Cloud 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 Google Cloud will send each notification to this recipient twice.

The screenshot shows the 'Add Cloud Spanner Policy' configuration page in Veeam Backup for Google Cloud. The page is titled 'Add Cloud Spanner Policy' and has a cost of '\$2.30'. The 'Settings' tab is selected, showing the following configuration options:

- Retries:** 'Automatically retry failed policy: 3 times' (checked).
- Health check:** 'Enable health check: On' (checked). 'Run on: First Wednesday of every month'.
- Notifications:** 'Enable notifications: Off' (unchecked).

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

How Health Check Works

When Veeam Backup for Google Cloud saves a new 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 Cloud Spanner instance data. When performing a health check, Veeam Backup for Google Cloud verifies availability of

data blocks for each restore point and uses the saved values to ensure that the restore points being verified are consistent.

NOTE

Veeam Backup for Google Cloud performs the health check during the last 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 Google Cloud 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 last policy session on Saturday.

Veeam Backup for Google Cloud performs the health check in the following way:

1. As soon as the backup policy session completes successfully, Veeam Backup for Google Cloud starts the health check as a new session. For each restore point in the regular backup chain, Veeam Backup for Google Cloud calculates CRC values for backup metadata and compares them to the CRC values that were previously saved to the restore point. Veeam Backup for Google Cloud 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 Google Cloud 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 Google Cloud starts the health check.

2. If Veeam Backup for Google Cloud 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 Google Cloud performs the following operations:

- If the health check detects corrupted metadata in a restore point, Veeam Backup for Google Cloud marks the regular backup chain as corrupted in the configuration database. During the next backup policy session, Veeam Backup for Google Cloud copies the full Cloud Spanner instance image, creates a new restore point and starts a new backup chain in the backup repository.

NOTE

Veeam Backup for Google Cloud supports metadata check for encrypted backup chains unless the metadata is corrupted.

- If the health check detects corrupted data blocks in a restore point, Veeam Backup for Google Cloud marks the restore point that includes the corrupted data blocks as incomplete in the configuration database. During the next backup policy session, Veeam Backup for Google Cloud copies the full Cloud Spanner instance image, creates a new restore point and starts a new backup chain in the backup repository.

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 approximate monthly cost of Google Cloud services that Veeam Backup for Google Cloud will require to protect the Cloud Spanner instances added to the backup policy. The total estimated cost includes the following:

- The cost of creating and maintaining cloud-native snapshots of the Cloud Spanner instances.
For each Cloud Spanner instance included in the backup policy, Veeam Backup for Google Cloud takes into account the amount of storage provisioned for the instance, 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 Cloud Spanner instances.
For each Cloud Spanner instance included in the backup policy, Veeam Backup for Google Cloud takes into account the amount of storage provisioned for the instance, the period of time during which restore points will be kept in the backup chain, and the configured scheduling and health check settings.
- The cost of creating and storing in backup repositories archived backups of the Cloud Spanner instances.
For each Cloud Spanner instance included in the backup policy, Veeam Backup for Google Cloud takes into account the amount of storage provisioned for the instance, the period of time during which restore points will be kept in the archive backup chain, and the configured scheduling settings.
- The cost of transferring the Cloud Spanner instance data between Google Cloud regions during data protection operations (for example, if a protected Cloud Spanner instance and the target storage bucket reside in different regions).
- The cost of sending API requests to Google Cloud during data protection operations.

During every backup session, Veeam Backup for Google Cloud creates a full backup of each Cloud Spanner instance included in the backup scope. This means that 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, either select a backup repository that resides in the same region as Cloud Spanner instances that you plan to back up, or select an archive repository that resides in the same region as the nearline or standard repository used to store regular backups.
- 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, modify the scheduling settings to run the backup policy less frequently, or specify an archive repository for long-term retention of restore points.

Veeam Backup for Google Cloud Server time: Nov 13, 2023 6:54 PM administrator Portal Administrator Configuration

← Add Cloud Spanner Policy Cost: \$2.30 ✓

Policy Info

Sources

Resources

Targets

Schedule

Settings

Cost Estimation

Permissions

Summary

Review cost estimation

The estimated cost takes into account the configured target settings, the specified scheduling options and the number of resources to protect.

Note that Veeam Backup for Google Cloud makes [predefined assumptions](#) to calculate the cost, which means that the results should be used only as an approximation.

For more information on how Veeam Backup for Google Cloud calculates the cost, see [this Veeam KB article](#).

\$0.00

\$0.11

\$0.01

\$0.00

\$0.00

Estimated monthly cost: \$0.12

Instance	Snapshots	Backups	Archives	Traffic	Transactions	Total
prkr-sp...	\$0.00	\$0.11	\$0.01	\$0.00	\$0.00	\$0.12

Step 9. Check Required Permissions

At the **Permissions** step of the wizard, Veeam Backup for Google Cloud will verify whether the specified service account has all the necessary permissions required to perform data protection tasks for the selected project or folder. For more information on the required permissions, see [Service Account Permissions](#).

To see the list of missing permissions that must be granted to the service account in order to perform an operation, click the link in the **Details** column. You can grant the missing permissions to the service account [using the Google Cloud console](#) or instruct Veeam Backup for Google Cloud to do it:

- To grant the missing permissions manually, click **Download Script**. Veeam Backup for Google Cloud will generate a gcloud script that you can run in the Google Cloud console to assign all the necessary permissions to the service account.

The account under which you run the script must have the permissions required both to get and set project IAM policies and to create custom IAM roles (for example, it can have the *iam.securityAdmin* and *iam.roleAdmin* roles assigned). To learn what permissions and roles are required to create custom roles in IAM, see [Google Cloud documentation](#).

- To let Veeam Backup for Google Cloud grant the missing permissions automatically, click **Grant** and then click **Sign in with Google** in the **Grant permissions** window. You will be redirected to the OAuth consent screen authorization page. Sign in using credentials of a Google account that will be used to grant the permissions.

The account under which you sign in to Google Cloud must have the permissions required both to get and set project IAM policies and to create custom IAM roles (for example, it can have the *iam.securityAdmin* and *iam.roleAdmin* roles assigned). To learn what permissions and roles are required to create service account, see [Google Cloud documentation](#).

NOTE

For Veeam Backup for Google Cloud to be able to authorize in Google Cloud, the OAuth consent screen must be configured as described in section [Registering Applications](#). Note that Veeam Backup for Google Cloud does not store in the configuration database the provided Google account credentials and access tokens received during authorization.

To make sure that the missing permissions have been successfully granted, click **Recheck**.

Veeam Backup for Google Cloud Server time: Nov 13, 2023 6:54 PM administrator Partial Administrator Configuration

← Add Cloud Spanner Policy Cost: \$2.30

Policy Info
Sources
Resources
Targets
Schedule
Settings
Cost Estimation
Permissions
Summary

Check permissions
Verify whether all the required permissions are granted.

[Recheck](#) [Download Script](#)

Check	Result	Details
Spanner Snapshot	Passed	All the required permissions are g...
Spanner Backup	Passed	All the required permissions are g...
Repository	Passed	All the required permissions are g...
Worker	Passed	All the required permissions are g...

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

Step 10. Finish Working with Wizard

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

The screenshot shows the 'Add Cloud Spanner Policy' wizard in the 'Summary' step. The interface includes a top navigation bar with the Veeam logo, server time (Nov 13, 2023 6:54 PM), user role (administrator), and configuration settings. The main content area is divided into a left sidebar with navigation options (Policy Info, Sources, Resources, Targets, Schedule, Settings, Cost Estimation, Permissions, Summary) and a main panel titled 'Review configured settings'. The main panel contains a 'Copy to Clipboard' button and several sections of configuration details: 'General settings' (Name: spannerdb2, Description: —, Service account: veeam-1691140381-sa@rnd-backup-2.iam.gserviceaccount.com), 'Protected resources' (SQL engine: Cloud Spanner, Folder: —, Project: RnD Backup 2, Regions: 1 region, Instances: 1 instance, Labels: —, Exclusions: —), 'Backup settings' (Backups enabled: Yes, Backup repository: prkr-spanner-rep, Archive repository: prkr-spanner-rep2, Weekly retention: Keep weekly backup for 21 days (6 backups excluded)), and 'Snapshot settings' (Snapshots enabled: Yes, Weekly retention: Keep 7 weekly snapshots). At the bottom, there are three buttons: 'Previous', 'Finish', and 'Cancel'.

Creating Snapshots Manually

Veeam Backup for Google Cloud allows you to manually create snapshots of Cloud Spanner instances. Each snapshot is stored in the location that depends on the [regional configuration](#) of the processed instance.

NOTE

Veeam Backup for Google Cloud 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 Google Cloud Storage unless you remove them manually, as described in section [Removing Backups and Snapshots](#).


To manually create a cloud-native snapshot of a Cloud Spanner instance, do the following:

1. Navigate to **Resources > Cloud Spanner**.
2. Select the necessary instance and click **Take Snapshot Now**.

For a Cloud Spanner instance to be displayed in the list of available instances, it must reside in any of the regions added to a backup policy as described in section [Creating Backup Policies](#).

3. In the **Take Snapshot Now** window, select a service account whose permissions Veeam Backup for Google Cloud will use to create the snapshot, and click **Take Snapshot**.

For a service account to be displayed in the **Service account** drop-down list, it must be added to Veeam Backup for Google Cloud as described in section [Adding Service Accounts](#), and must be assigned the *Cloud Spanner Instances Snapshot* operational role as described in section [Adding Projects and Folders](#).



Managing Backup Policies

You can manage and edit created VM, Cloud SQL and Cloud Spanner backup policies, and view each backup policy details in Veeam Backup for Google Cloud. You can also remove backup policies that you do not use anymore, export existing or import new backup policies.

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. Navigate to **Policies**.
2. Switch to the necessary tab and select the backup policy.
3. Click **Start** or **Stop**.

The screenshot displays the Veeam Backup for Google Cloud interface. The top navigation bar includes the Veeam logo, the product name, the server time (Nov 20, 2023 6:19 PM), the user (administrator), and a Configuration icon. The left sidebar shows navigation options: Monitoring (Overview, Resources), Management (Policies, Protected Data, Session Logs), and the current view is 'Policies'. The main content area is divided into three tabs: VM, Cloud SQL, and Cloud Spanner. The Cloud Spanner tab is active, showing a search bar for 'Policy' and a 'Filter (None)' button. Below the search bar are action buttons: Start, Stop, Disable, Add, Edit, Priority, View Info, Remove, Advanced, and Export to... A table lists backup policies with columns for Priority, Policy, Snapshots, Backups, Archives, Last Run, and Last Duration. Two policies are shown: 'sapaner' (Failed) and 'dry' (Cancelled). Below the table are two sections: 'Instances' and 'Sessions'. The 'Instances' section shows a search bar and a list of instances: 'dr-testrestore' and 'dr2k'. The 'Sessions' section shows a status bar (Success, Warning, Cancelled) and a table of sessions with columns for Type, Time, and Status. The sessions table lists various backup and snapshot operations with their respective times and statuses (Cancelled, Warning, Success).

Enabling and Disabling Backup Policies

By default, Veeam Backup for Google Cloud runs all created backup policies according to the specified schedules. However, you can temporarily disable a backup policy so that Veeam Backup for Google Cloud 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 **Enable** or **Disable**.

The screenshot displays the Veeam Backup for Google Cloud interface. The top navigation bar includes the Veeam logo, server time (Nov 20, 2023 6:28 PM), and user information (administrator, Portal Administrator). The left sidebar shows navigation options: Monitoring, Overview, Resources, Management, Policies, Protected Data, and Session Logs. The main area is divided into tabs for VM, Cloud SQL, and Cloud Spanner. A search bar and filter options are present above a table of backup policies. Below the table are sections for Instances and Sessions.

Priority	Policy	Snapshots	Backups	Archives	Last Run	Last Duration
<input checked="" type="checkbox"/>	1 vms	Warning	Warning	Failed	11/20/2023 6:00:06 PM	18 min 15 sec
<input type="checkbox"/>	2 empty	Warning	Not Configured	Not Configured	11/06/2023 11:53:16 PM	2 sec
<input type="checkbox"/>	3 abigone	Failed	Not Configured	Not Configured	11/09/2023 9:02:49 PM	4 min 23 sec
<input type="checkbox"/>	4 v4	Success	Not Configured	Not Configured	11/16/2023 7:18:41 PM	2 min 49 sec

Instance	Status
dr-import2	Success
dr-import2-1	Warning
dr-ub-renamed	Success

Type	Time	Status
Health Check	11/20/2023 6:18:24 ...	Success
Backup	11/20/2023 6:00:07 ...	Warning
Snapshot	11/20/2023 6:00:06 ...	Warning
Backup	11/20/2023 5:00:13 ...	Warning
Snapshot	11/20/2023 5:00:13 ...	Warning
Health Check	11/19/2023 6:18:19 ...	Success

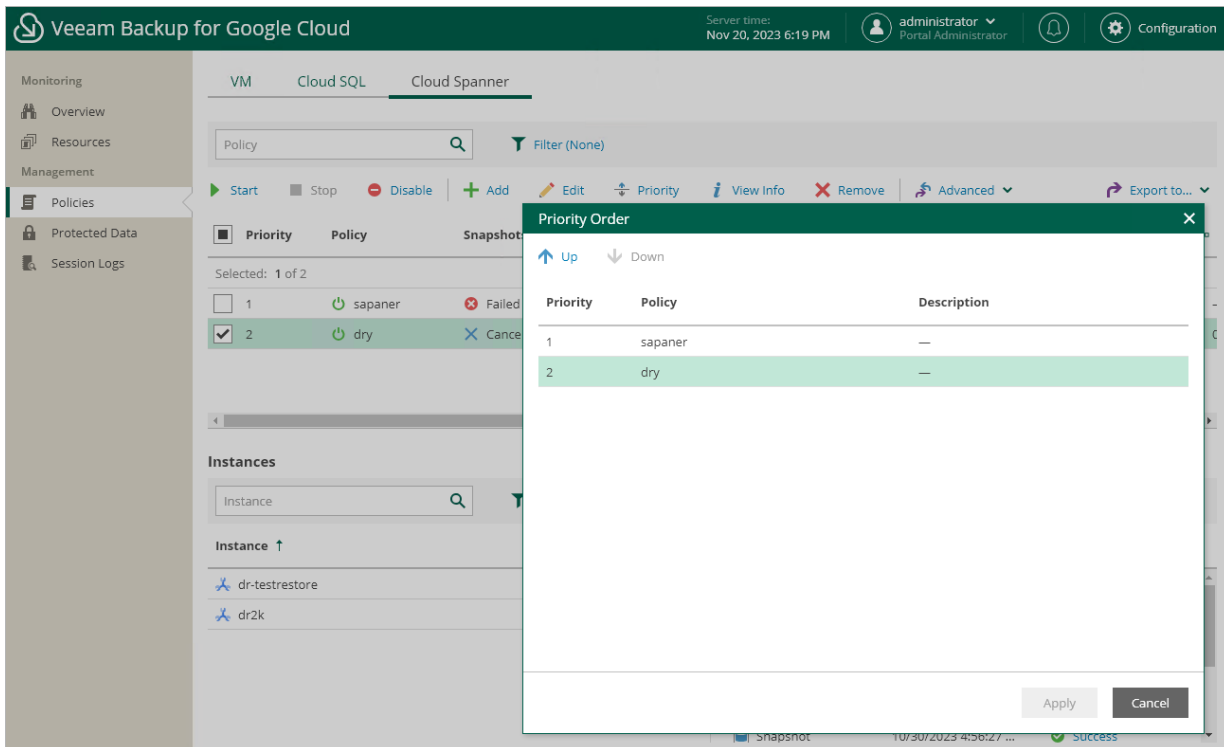
Setting Backup Policy Priority

By default, Veeam Backup for Google Cloud runs backup policies in the order you create them. However, you can set the backup policy priority manually:

1. Navigate to **Policies**.
2. Switch to the necessary tab and click **Priority**.
3. In the **Priority Order** window, do the following:
 - a. Select a backup policy in the list of existing policies.
 - b. To move the policy up or down the list, use the **Up** and **Down** arrows.
 - c. To save changes made to the priority order, click **Apply**.

NOTE

If an instance is included into multiple backup policies, it will be processed only by the backup policy that has the highest priority.



Editing Backup Policy Settings

For each backup policy, you can modify settings configured while creating the policy:

1. Navigate to **Policies**.
2. Switch to the necessary tab and select the backup policy.
3. Click **Edit**.
4. Complete the **Edit Policy** wizard:
 - a. To provide a new name and description for the policy, follow the instructions provided in section [Performing VM Backup](#) (step 2), [Performing SQL Backup](#) (step 2) or [Performing Spanner Backup](#) (step 2).
 - b. To choose another project or folder that manages resources that you want to protect, or change the service account whose permissions are used to perform backup operations, follow the instructions provided in section [Performing VM Backup](#) (step 3), [Performing SQL Backup](#) (step 3) or [Performing Spanner Backup](#) (step 3).

IMPORTANT

If you change the project, folder or service account, it is recommended that you check whether the selected service account has all the permissions required to perform data protection tasks in the specified entity. To do that, follow the instructions provided in section [Performing VM Backup](#) (step 10), [Performing SQL Backup](#) (step 10) or [Performing Spanner Backup](#) (step 10).

- c. To modify the list of regions in which instances that you plan to back up reside, or to add instances to the backup scope, follow the instructions provided in section [Performing VM Backup \(step 4a or step 4b\)](#), [Performing SQL Backup \(step 4a or step 4b\)](#) or [Performing Spanner Backup \(step 4a or step 4b\)](#).
- d. To instruct Veeam Backup for Google Cloud to create image-level backups, follow the instructions provided in section [Performing VM Backup \(step 5\)](#), [Performing SQL Backup \(step 5\)](#) or [Performing Spanner Backup \(step 5\)](#).
- e. To modify the schedule configured for the policy, follow the instructions provided in section [Performing VM Backup \(step 6\)](#), [Performing SQL Backup \(step 6\)](#) or [Performing Spanner Backup \(step 6\)](#).
- f. [This step applies only to VM backup policies] To assign labels to cloud-native snapshots, follow the instructions provided in section [Performing VM Backup \(step 7\)](#).
- g. [This step applies only to SQL backup policies] To choose whether you want to use a staging server to perform backup, follow the instructions provided in section [Performing SQL Backup \(step 7\)](#).
- h. To configure automatic retry, health check and notification settings, follow the instructions provided in section [Performing VM Backup \(step 8\)](#), [Performing SQL Backup \(step 8\)](#) or [Performing Spanner Backup \(step 8\)](#).

- i. At the **Summary** step of the wizard, review configuration information and click **Finish** to confirm the changes.

The screenshot shows the 'Edit Policy dry' wizard in Veeam Backup for Google Cloud. The interface is in the 'Summary' step, which is highlighted in the left-hand navigation menu. The main content area displays 'Review configured settings' with a 'Copy to Clipboard' button. The settings are organized into several sections:

- General settings:** Name: dry; Description: —; Service account: veeam-1649186685-sa@rnd-backup-2.iam.gserviceaccount.com
- Protected resources:** SQL engine: Cloud Spanner; Folder: folder-with-a-thousand-of-proj; Project: 200 projects; Regions: 1 region; Instances: 2 instances; Labels: —; Exclusions: —
- Backup settings:** Backups enabled: Yes; Backup repository: backup-readd; Archive repository: archive; Monthly retention: Keep monthly backup for 30 days (11 backups excluded)
- Snapshot settings:** Snapshots enabled: Yes; Monthly retention: Keep 7 monthly snapshots
- Other settings:** Automatic retries enabled: Yes; Notifications enabled: No; Health check enabled: No

At the bottom of the wizard, there are three buttons: 'Previous', 'Finish' (highlighted in green), and 'Cancel'. The top right corner shows the cost: '\$619.74' with a green checkmark. The top navigation bar includes the Veeam logo, server time (Nov 20, 2023 6:20 PM), user profile (administrator, Portal Administrator), and a 'Configuration' gear icon.

Exporting and Importing Backup Policies

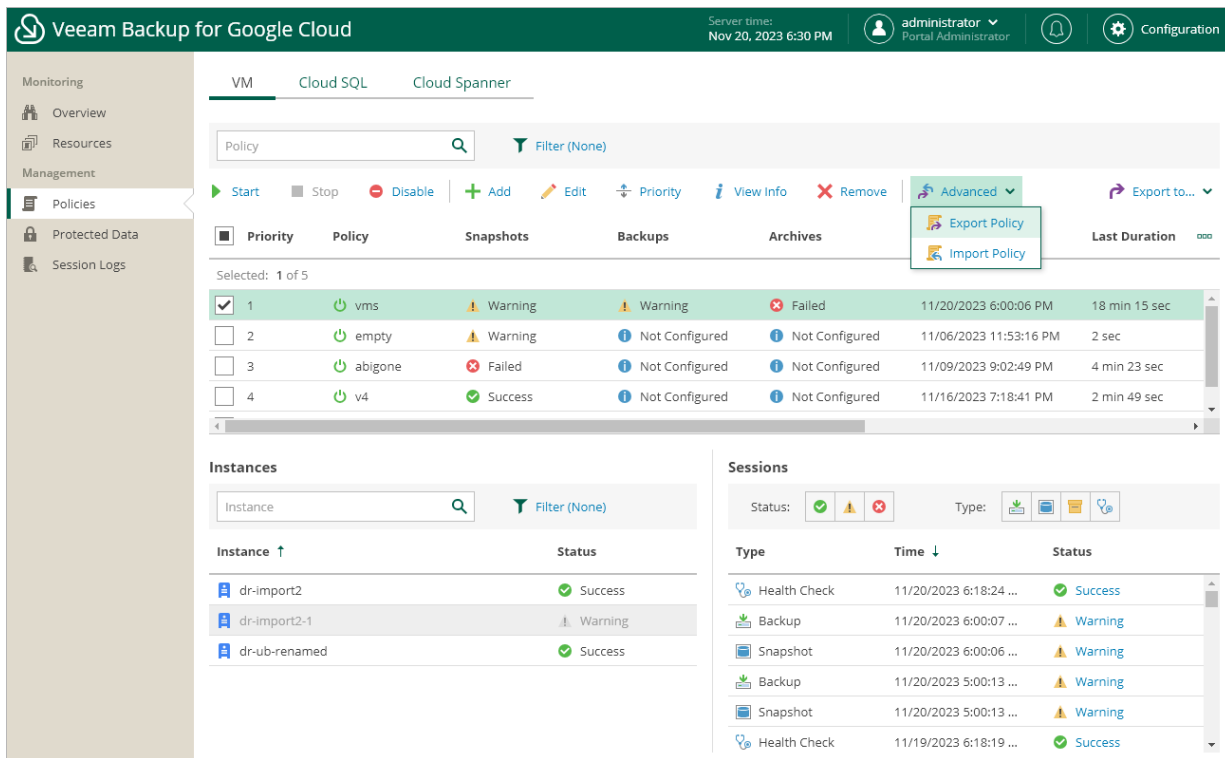
Veeam Backup for Google Cloud allows you to use settings of an existing backup policy as a template for creating other backup policies. You can export a backup policy to a .JSON file, modify the necessary settings in the file, and then import the policy to the same or a different backup appliance.

Exporting Backup Policies

To export a backup policy to a .JSON file, do the following:

1. Navigate to **Policies**.
2. Switch to the necessary tab and select the backup policy.
3. Click **Advanced** > **Export Policy**.

Veeam Backup for Google Cloud 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, do the following:

1. Click **Advanced > Import Policy**.
2. In the **Import Policy** window, specify a name for the imported backup policy, paste the content of the necessary .JSON file, and click **Import**.

Veeam Backup for Google Cloud

Server time: Dec 6, 2023 3:21 PM administrator Portal Administrator Configuration

Monitoring: Overview, Resources

Management: Policies, Protected Data, Session Logs

VM Cloud SQL Cloud Spanner

Policy Filter (None)

Start Stop Enable Add Edit Priority View Info Remove Advanced Export to...

Priority Selected: 0 of 6

Instances Instance Instance No data

Last Duration

1:00:06 PM 26 min

7:20:14 AM 2 min 12 sec

Status

Import Policy

Enter a policy name and paste a JSON expression that specifies the policy configuration

Name: asia-east-policy2

JSON:

```
{
  "backupType": "SelectedItems",
  "snapshotSettings": {
    "additionalTags": [],
    "copyTagsFromVolumeEnabled": false,
    "tryCreateVSSSnapshot": false
  },
  "backupSettings": {
    "enabled": true,
    "repositoryId": "97de d4dd-7fe8-4f57-a4ed-b7fa614984a8",
    "repositoryName": "",
    "archiveRepositoryId": "7eb30b48-562e-4cd4-9e8f-f309c7f117ed",
    "archiveRepositoryName": "",
    "healthCheckEnabled": false
  }
}
```

Import Cancel

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 Google Cloud 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 VM, Cloud SQL and Cloud Spanner instances, [open the Veeam Backup for Google Cloud Web UI](#) and follow the instructions provided in section [Removing Backups and Snapshots](#).

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 a VM instance.
	Indicates that the protected workload is a Cloud SQL instance.
	Indicates that the protected workload is a Cloud Spanner instance.

The **Backups** node contains 4 subnodes:

- The **Snapshots** subnode displays information on cloud-native snapshots of the protected VM, Cloud SQL and Cloud Spanner instances:
 - *<appliance_name>* nodes show snapshots created manually on backup appliances and snapshots imported to the appliances from Google Cloud regions specified in backup policy settings.
 - *<backup_policy_name>* nodes show snapshots created by backup policies.

To learn how Veeam Backup for Google Cloud creates cloud-native snapshots, see [VM Snapshot Chain](#), [Cloud SQL Snapshot Chain](#) and [Cloud Spanner Snapshot Chain](#).

- The **External Repository** subnode displays information on image-level backups of the protected VM, Cloud SQL and Cloud Spanner instances that are stored in standard repositories.

To learn how Veeam Backup for Google Cloud creates image-level backups, see [VM Backup Chain](#), [Cloud SQL Backup Chain](#) and [Cloud Spanner Backup Chain](#).

NOTE

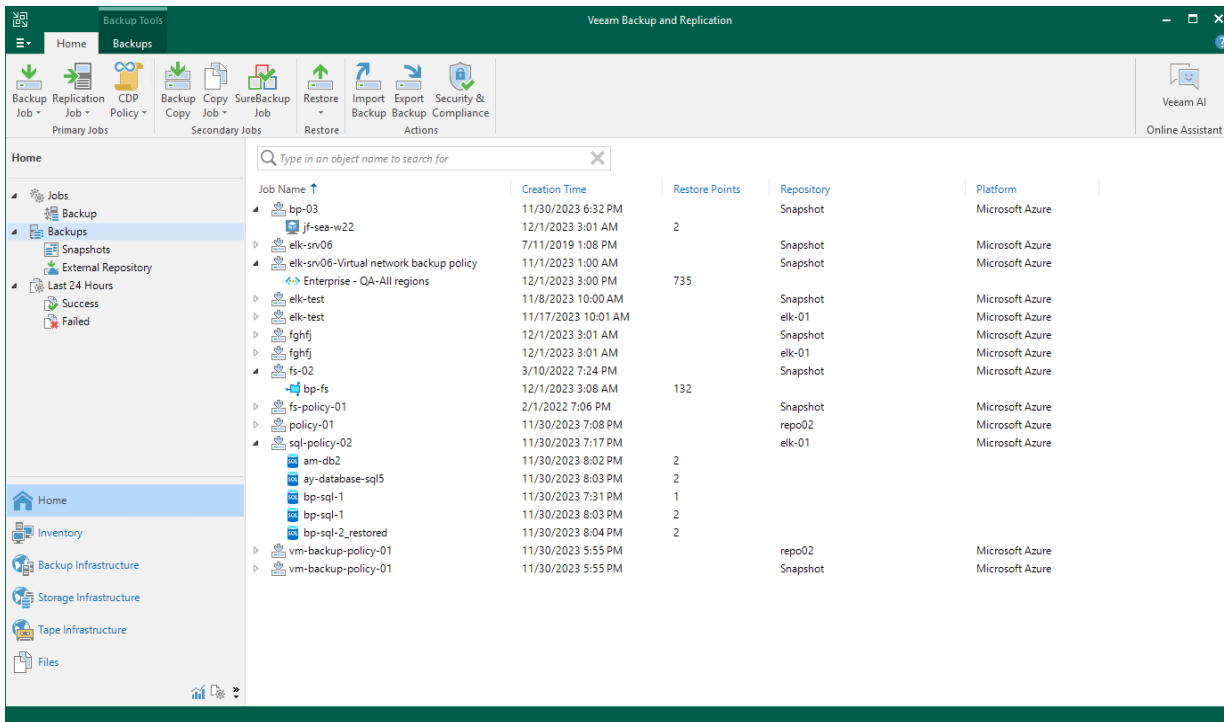
If a backup chain was originally encrypted and then got decrypted by Veeam Backup & Replication, the backup chain will be marked with the **Key** icon.

- The **External Repository (Encrypted)** subnode displays information on encrypted image-level backups of the protected VM, Cloud SQL and Cloud Spanner instances that are stored in standard 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 the protected VM, Cloud SQL and Cloud Spanner instances that are stored in archive repositories.

To learn how Veeam Backup for Google Cloud creates archive backups, see [VM Archive Backup Chain](#), [Cloud SQL Archive Backup Chain](#) and [Cloud Spanner Archive Backup Chain](#).



Decrypting Backups

Veeam Backup & Replication automatically decrypts backup files stored in repositories using passwords that you specify when [adding these repositories](#) to the backup infrastructure. If you do not specify decryption passwords, the backup files remain encrypted.

To decrypt backup files, do the following:

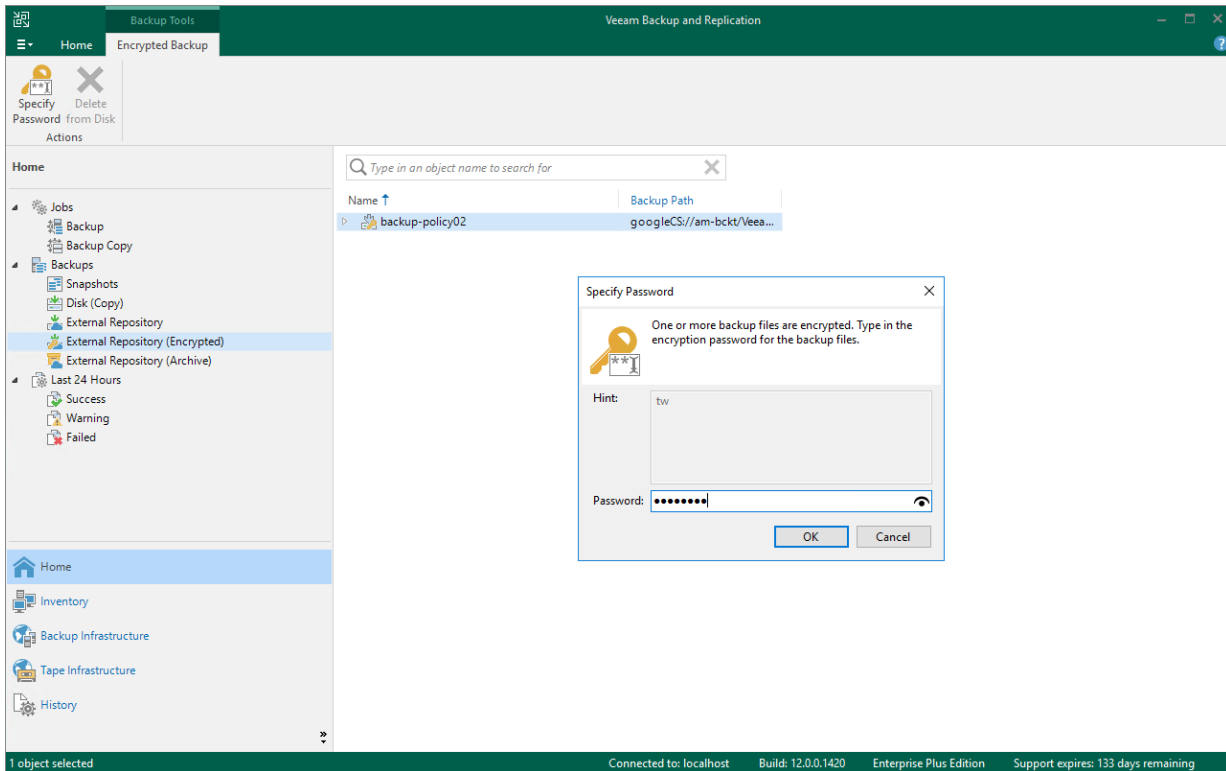
- In the Veeam Backup & Replication console, open the **Home** view.
- Navigate to **Backups > External Repository (Encrypted)**.
- Expand the backup policy that protects a VM instance whose image-level backups you want to decrypt, and select the backup chain that belongs to the instance. Click **Specify Password** on the ribbon.

Alternatively, you can right-click the necessary backup chain and select **Specify password**.

TIP

To decrypt all backups created by the policy, right-click the backup policy and select **Specify Password**.

4. 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

After a backup policy successfully creates a restore point for a Google Cloud resource, or after you create a snapshot of a resource manually using Veeam Backup for Google Cloud, the resource is automatically added to the resource list on the **Protected Data** tab.

The **Protected Data** tab displays Google Cloud resources that are already protected by Veeam Backup for Google Cloud. Each resource is represented with a set of properties, such as:


- **Instance** — the name of the resource.
- **Policy** — the name of the backup policy that protects the resource.
- **Restore Points** — the number of restore points created for the resource.
- **Latest Restore Point** — the date and time of the most recent restore point created for the resource.
- **Region** — the region in which the resource resides.
- **Configuration** — the instance configuration that defines the geographic location where the Cloud Spanner instance data is stored.
- **Engine** — the database engine and version installed on the Cloud SQL instance.
- **Operating System** — the operating system running on the VM instance.
- **File-level Recovery URL** — the link to the File-Level Restore browser.

The link appears when Veeam Backup for Google Cloud starts a restore session to perform [file-level recovery](#). The link contains a public DNS name of the worker instance hosting the File-Level Restore browser and authentication information used to access this worker instance.

On the **Protected Data** tab, you can perform the following actions:

- Remove restore points if you no longer need them. For more information, see [Removing Backups and Snapshots](#).

- Restore data of backed-up VM, Cloud SQL and Cloud Spanner instances. For more information, see sections [Performing VM Restore](#), [Performing SQL Restore](#) and [Performing Spanner Restore](#).



Removing Backups and Snapshots

Veeam Backup for Google Cloud stores information on all protected Google Cloud resources in the configuration database. Even if a resource is no longer protected by any backup policy, information on the backed-up data will not be deleted from the database until Veeam Backup for Google Cloud automatically removes all restore points associated with this resource according to the retention settings saved in the backup metadata. If necessary, you can also remove the restore points manually.

IMPORTANT

Do not delete backups from Google Cloud storage buckets in the Google Cloud console. If some backup in a backup chain is missing, you will not be able to roll back the resource data to the necessary state.

To remove restore points manually, do the following:

- Navigate to **Protected Data**.
- Switch to the necessary tab and select resources whose restore points you want to remove.
- Click **Remove** and select either of the following options:
 - Snapshots > All** — to remove all cloud-native snapshots created for the selected resources both by backup policies and manually.
 - Snapshots > Created by Policy** — to remove all cloud-native snapshots created for the selected resources by backup policies.
 - Snapshots > Created Manually** — to remove all cloud-native snapshots created for the selected resources manually.

- **Backups > All** – to remove all image-level backups created for the selected resources.
- **Backups > Standard and Nearline** – to remove all image-level backups created for the selected resources in backup repositories of the *Standard* and *Nearline* storage classes.
- **Backups > Archived** – to remove all image-level backups created for the selected resources in backup repositories of the *Archive* storage class.
- **Snapshots and Backups** – to remove both cloud-native snapshots and image-level backups created for the selected resources.

TIP

Cloud Spanner snapshots will be retained in Google Cloud Storage for up to one year only, regardless of the policy settings. However, the retention time for regular backups and archived backups stored in Veeam repositories is not limited.

The screenshot shows the Veeam Backup for Google Cloud interface. The top navigation bar includes the Veeam logo, the product name, the server time (Nov 24, 2023 9:32 PM), the user role (administrator), and a configuration icon. The left sidebar contains navigation options: Monitoring (Overview, Resources), Management (Policies, Protected Data, Session Logs), and Cloud Spanner. The main area displays a table of instances under the 'Cloud Spanner' tab. A context menu is open over the table, showing options: Remove, Snapshots, Backups, Snapshots and Backups, All, Standard and Nearline, and Archived. The table lists various instances with columns for Instance, Project, Policy, Restored, and File Level Recovery.

Instance	Project	Policy	Restored	File Level Recovery
azagnolko1	veeam-rnd-backup-3	—	5	europa-west3-c
azagnolko2	veeam-rnd-backup-3	—	2	europa-west3-c
azenkov-gen-vm	rnd-backup-scalability	—	1	us-central1-a
dr-current-vm	veeam-rnd-backup-2	core	2	us-central1-a
dr-fstabinstancer	veeam-rnd-backup-2	—	2	us-west3-a
<input checked="" type="checkbox"/> dr-import2	veeam-rnd-backup-2	dual-disks	430	us-central1-a
<input checked="" type="checkbox"/> dr-import2-1	veeam-rnd-backup-2	dual-disks	249	us-central1-a
dr-ub-renamed	veeam-rnd-backup-2	v2	12	us-central1-a
dr-veeam-backup-for...	veeam-rnd-backup-2	—	1	us-west3-c
dr-win-disks	veeam-rnd-backup-2	—	12	us-central1-a
general-purpose-clou...	veeam-rnd-backup-3	—	1	us-central1-a
kk-test-resmon1	veeam-rnd-backup-3	—	1	europa-west3-c
kl-flr-big-disk	veeam-rnd-backup-3	—	5	europa-north1-a
kl-flr-test-vm-vm	veeam-rnd-backup-3	—	1	europa-north1-a
kl-perf-lab-vm	veeam-rnd-backup-3	—	7	europa-north1-a
kl-rocky-test	veeam-rnd-backup-3	—	1	europa-north1-a
pim-1disk	veeam-rnd-backup-3	—	2	us-central1-a
pim-2disks	veeam-rnd-backup-3	—	7	us-central1-a

Performing Restore

In various disaster recovery scenarios, you can perform the following restore operations using backed-up data:

- [Restore of VM instances](#) — restore VM instances, disks and files from cloud-native snapshots or image-level backups to the original location or to a new location.
- [Restore of Cloud SQL instances](#) — restore Cloud SQL instances (from cloud-native snapshots or image-level backups) and Cloud SQL databases (from image-level backups) to the original location or to a new location.
- [Restore of Cloud Spanner instances](#) — restore Cloud Spanner instances (from cloud-native snapshots or image-level backups) and Cloud Spanner databases (from cloud-native snapshots or image-level backups) to the original location or to a new location.
- [Instant Recovery](#) — immediately restore VM instances from image-level backups to VMware vSphere and Hyper-V environments, and to Nutanix AHV clusters.
- [VM instance disk export](#) — restore virtual disks and convert them to disks of the VMDK, VHD or VHDX format.
- [Disk publishing](#) — publish point-in-time disks and copy the necessary files and folders to the target server.
- [Restore to AWS](#) — restore VM instances from image-level backups to AWS as EC2 instances.
- [Restore to Microsoft Azure](#) — restore VM instances from image-level backups to Microsoft Azure as Azure VMs.
- [Restore to Nutanix AHV](#) — restore VM instances from image-level backups to Nutanix AHV as Nutanix AHV VMs.

VM Restore

The actions that you can perform with restore points of VM instances depend on whether you access the restore points using the Veeam Backup & Replication console or the Veeam Backup for Google Cloud Web UI.

VM Restore Using Console

Veeam Backup & Replication offers the following restore operations:

- [Instance restore](#) – start an entire VM instance from a restore point.
- [Guest OS file recovery](#) – restore individual files and folders of a VM instance.
- [Application restore](#) – restore applications such as Microsoft Active Directory, Microsoft Exchange, Microsoft SharePoint and Microsoft SQL Server.

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

NOTE

You can use restore points stored in standard repositories to perform all the listed recovery operations, while restore points stored in archive repositories can only be used to perform restore of VM instances to the original or to a new location.

Performing VM Instance Restore

In case a disaster strikes, you can restore an entire VM instance from a cloud-native snapshot or an image-level backup. Veeam Backup & Replication allows you to restore one or more VM instances at a time, to the original location or to a new location.

IMPORTANT

When restoring a VM instance, Veeam Backup for Google Cloud recovers data from all zonal and regional persistent disks (standard, balanced, extreme and SSD) attached to the instance. However, when it comes to local SSDs (SCSI and NVMe), Veeam Backup for Google Cloud is able to recover only the configuration of these disks due to [technical reasons](#).

How Instance Restore Works

To restore VM instances from cloud-native snapshots, Veeam Backup & Replication uses [native Google Cloud capabilities](#). To restore VM 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 [Performing Instance 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 Google Compute Engine Works](#).

How to Perform Instance Restore

To restore an entire VM instance, do the following:

1. [Launch the Restore to Google Compute Engine wizard](#).
2. [Select a restore point](#).
3. [Choose a restore mode](#).

4. [Select a project, region and an availability zone.](#)
5. [Specify instance type and encryption settings.](#)
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 Google Compute Engine Wizard


To launch the **Restore to Google Compute Engine** 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 a VM instance that you want to restore, select the necessary instance and click **Google CE** on the ribbon.

Alternatively, you can right-click the instance and select **Restore to Google CE**.

TIP

You can also launch the **Restore to Google Compute Engine** wizard from the **Home** tab. To do that, click **Restore** and select **GCP**. Then, in the **Restore** window, select **Google Compute Engine > Entire machine restore > Restore to public cloud > Restore to Google Compute Engine** and, depending on whether you want to restore from a backup or a snapshot, click either **Restore from GCE VM snapshot** or **Restore from Veeam backup**.



Step 2. Select Restore Point

At the **Machine** step of the wizard, choose a restore point that will be used to restore the selected VM 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:

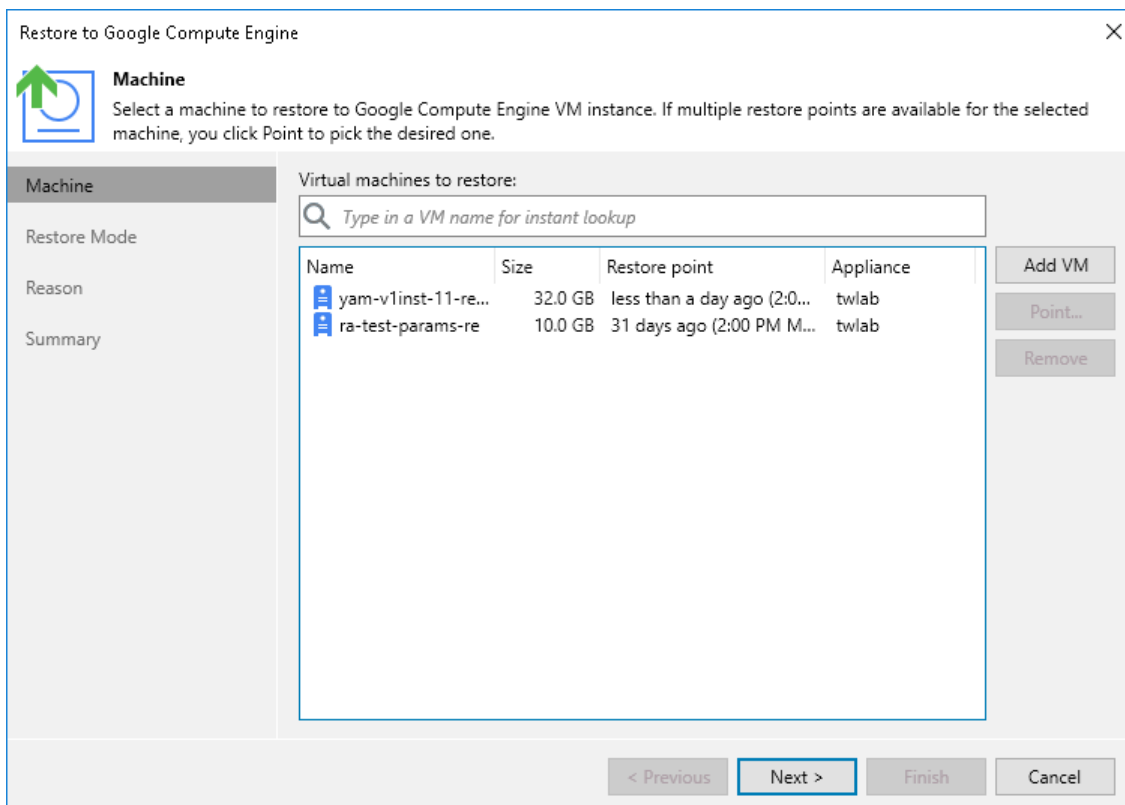
1. In the **Virtual machines to restore** list, select the VM instance and click **Point**.
2. In the **Restore Points** window, expand the backup policy that protects the VM instance, select the necessary restore point and click **OK**.

To help you choose a restore point, Veeam Backup & Replication provides the following information on each available restore point:

- o **Job** – the name of the backup policy that created the restore point and the date when the restore point was created.
- o **Type** – the type of the restore point.
- o **Location** – the 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 VM**, select more VM 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 VM instance to the original or to a new location.

IMPORTANT

When restoring a VM instance to the original location while the source VM instance still exists in Google Cloud, Veeam Backup for Google Cloud restores the instance with a different name, powers off the source VM instance, removes the source instance from the backup infrastructure, and then renames the restored VM instance. To allow the backup appliance to perform these operations, make sure that the [deletion protection](#) setting is disabled for the source instance, and the service account whose permissions Veeam Backup for Google Cloud uses to perform the restore operation is assigned the `compute.instances.setName` permission. Otherwise, you will not be able to perform restore to the original location.

Consider that the ability to rename VM instances is currently in pre-GA state. For more information, see [Google Cloud documentation](#).


2. Click **Pick account to use** to select a service account whose permissions will be used to perform the restore operation. For more information on the required permissions, see [Service Account Permissions](#).

For a service account to be displayed in the list of available accounts, it must be added to Veeam Backup for Google Cloud as described in section [Adding Service Accounts](#), and must be assigned the *VM Instances Restore* operational role as described in section [Adding Projects and Folders](#).

NOTE

By default, to perform restore operations, Veeam Backup & Replication uses permissions of service accounts that have been used to protect the source VM instances.

Restore to Google Compute Engine ✕

 **Restore Mode**
Specify whether selected VMs should be restored back to the original location, or to a new location or with different settings.

Machine	<input checked="" type="radio"/> Restore to the original location Quickly initiate the restore of selected VM to its original location, with the original name and settings. This option minimizes the chance of user input error.
Restore Mode	<input type="radio"/> Restore to a new location, or with different settings Customize the restored VM location, and change its settings. The wizard will automatically populate all controls with the original VM settings as the defaults. Pick account to use
Reason	
Summary	

< Previous **Next >** Finish Cancel

Step 4. Select Project, Region and Availability Zone

[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 a project that will be used to manage the restored VM instance, and specify a region and an availability zone where the restored VM instance will operate.

For a project to be displayed in the list of available projects, it must be created in Google Cloud as described in [Google Cloud documentation](#).

The screenshot shows a wizard window titled "Restore to Google Compute Engine" with a close button (X) in the top right corner. The "Data Center" step is active, indicated by a green arrow icon and a blue highlight in the left sidebar. The sidebar lists the following steps: Machine, Restore Mode, Data Center (selected), Machine Type, Name, Network, Reason, and Summary. The main content area contains the following fields and instructions:

- Project:** A dropdown menu with "RnD Backup" selected. Below it, the instruction reads: "Specify a project for the restored instance."
- Data center:** A dropdown menu with "europe-north1 (Finland)" selected. Below it, the instruction reads: "Select a data center based on the geographical proximity or pricing."
- Availability zone:** A dropdown menu with "europe-north1-c" selected. Below it, the instruction reads: "Specify an availability zone within data center region for the restored instance."

At the bottom of the window, there are four buttons: "< Previous" (disabled), "Next >" (active/highlighted), "Finish" (disabled), and "Cancel" (disabled).

Step 5. Specify Instance Type and 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 **Machine Type** step of the wizard, you can configure the restored VM instance settings. To do that, select the instance and do the following:

- If you want to specify a new machine type for the restored VM instance, click **Type** and select the necessary type in the **Machine Type** window.

To learn how to choose a machine type when creating a VM instance in Google Cloud, see [Google Cloud documentation](#).

- If you want to change the encryption settings of the restored VM instance, click **Encryption** and do the following in the **Disk Encryption** window:
 - If you do not want to encrypt persistent disks of the restored VM instance or want to apply the existing encryption scheme of the source VM instance, select the **Preserve the original encryption settings** option.
 - If you want to encrypt persistent disks of the restored VM instance with a Google Cloud KMS CMEK, select the **Use the following encryption key** option. Then, choose the necessary CMEK from the list.

For a CMEK to be displayed in the list of available encryption keys, it must be stored in the region selected at [step 4](#) of the wizard.

NOTE

The **Preserve the original encryption settings** option is disabled if the CMEK that was used to encrypt persistent disk of the source instance is not available in the region to which the VM instance will be restored.

Restore to Google Compute Engine

Machine Type
Specify the machine type and disk encryption settings for the restored virtual machine.

Machine

Restore Mode

Data Center

Machine Type

Name

Network

Reason

Summary

Virtual machine:

Name	Instance Type	Disk Encryption
yam-v1inst-11-redha...	e2-medium (2 cores, 4.0...	Preserve original settings
ra-test-params-re	f1-micro (1 core, 614 M...	Preserve original settings

Machine Type

Machine type:
e2-medium (2 cores, 4.00 GB memory)

vCPUs: 2

Memory: 4.00 GB

OK Cancel

Select multiple instances to apply settings change in bulk.

Type... Encryption...

< Previous Next > Finish Cancel

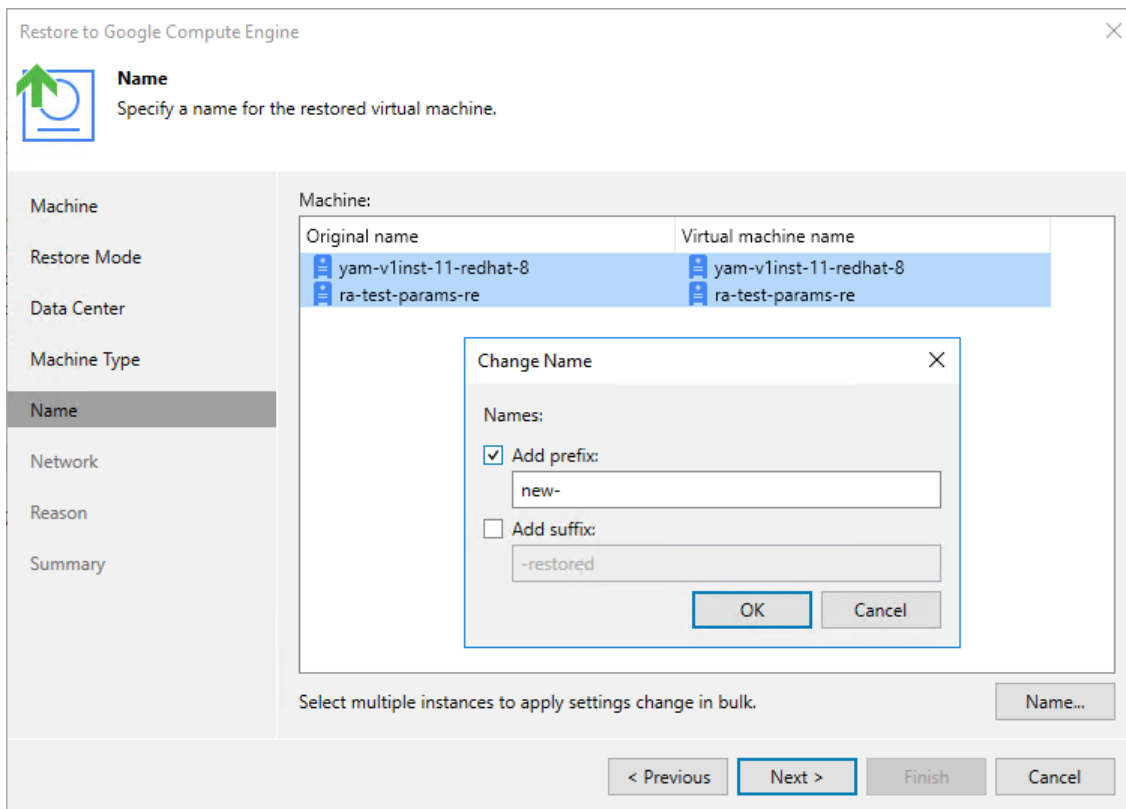
Step 6. Specify Instance Name

[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, specify a new name for the restored VM instance.

TIP

You can specify a single prefix or suffix and add it to the names of multiple VM 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

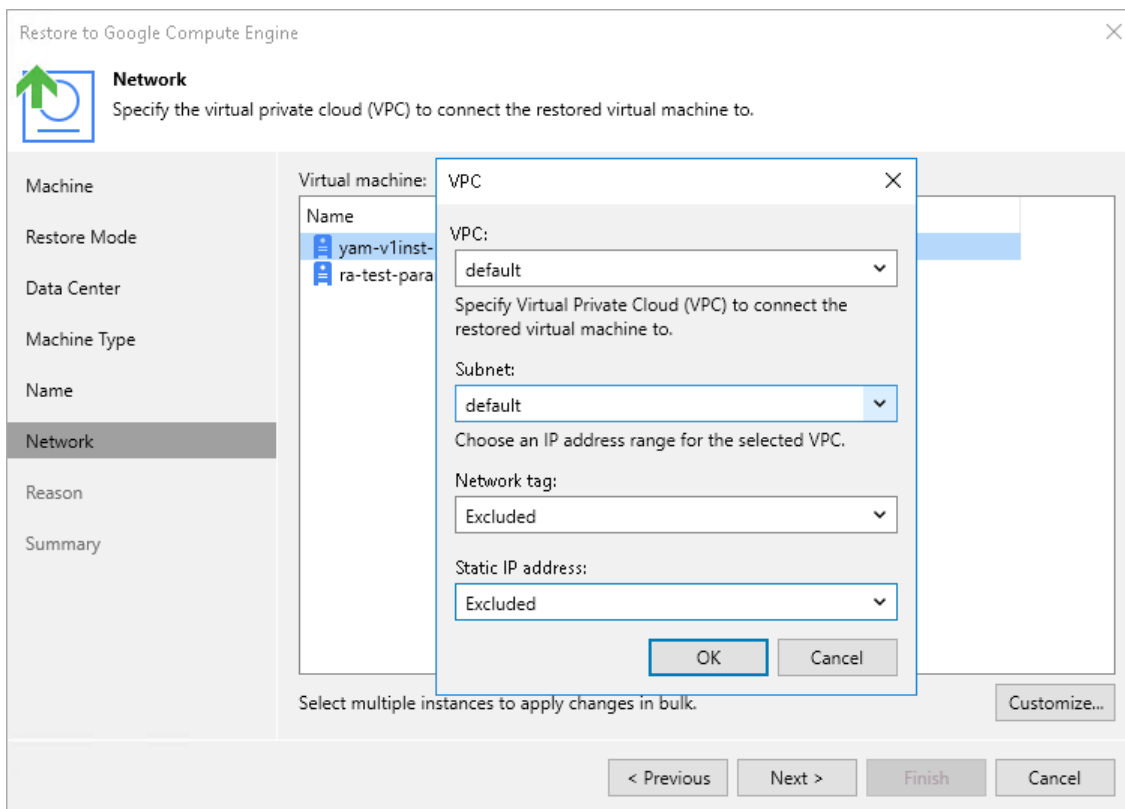
[This step applies only if you have selected the **Restore to a new location, or with different settings** option at the **Restore Mode** step of the wizard]

At the **Network** step of the wizard, you can select a VPC network and subnet to which the restored VM instance will be connected. To do that, select the VM instance and click **Customize**. You can also choose whether you want the restored VM instance to have the same network tags and the same reserved static external IP address as the source VM instance – to enable an option, select *Included* from the drop-down list.

NOTE

A static external IP address can be assigned to a restored VM instance only if this IP address has been reserved for the source VM instance. To learn how to reserve static external IP addresses for VM instances, see [Google Cloud documentation](#).

For a VPC network and subnet to be displayed in the lists of available networks, they must be created for the region specified at [step 4](#) in Google Cloud, as described in [Google Cloud documentation](#).



Step 8. Specify Restore Reason

At the **Reason** step of the wizard, specify a reason for restoring the VM instance. The information you provide will be saved in the session history and you can reference it later.

Restore to Google Compute Engine

Reason
Type in the reason for performing this restore operation. This information will be logged in the restore sessions history for later reference.

Machine
Restore Mode
Data Center
Machine Type
Name
Network
Reason
Summary

Restore reason:
Restore failed VMs

Do not show me this page again

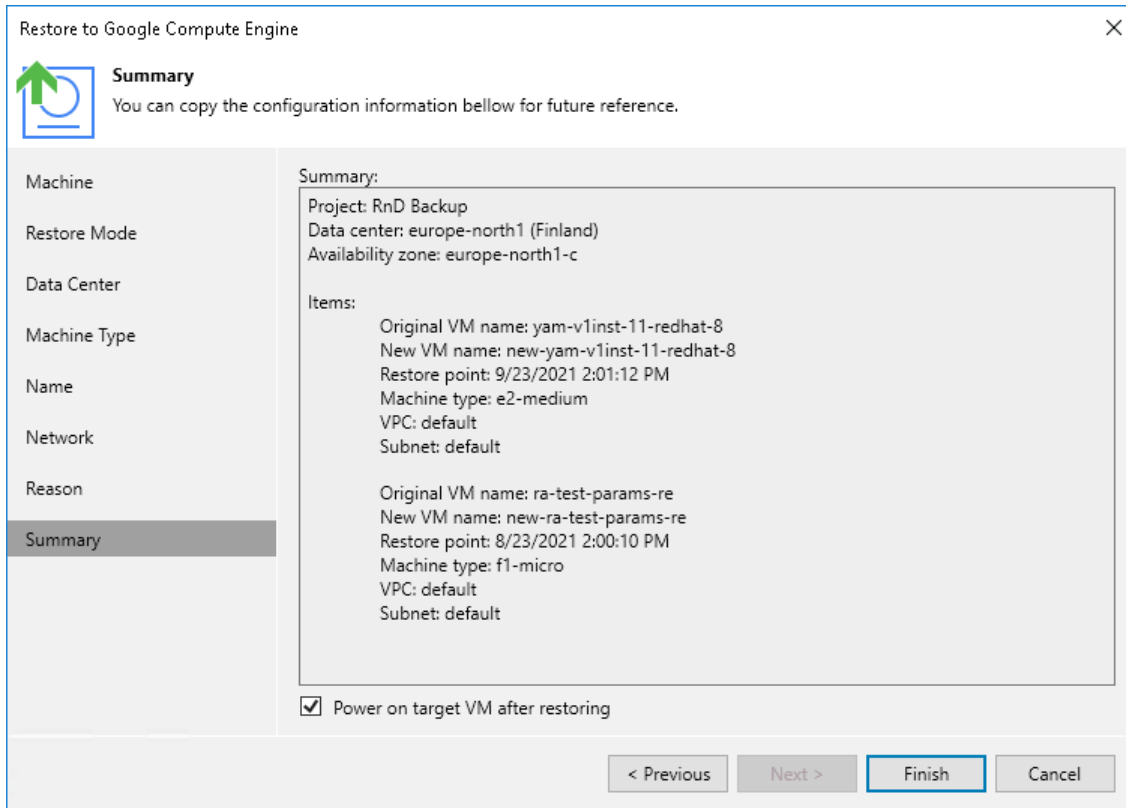
< Previous Next > Finish Cancel

Step 9. Finish Working with Wizard

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

TIP

If you want to start the VM instance immediately after restore, select the **Power on target VM after restoring** check box.



The screenshot shows a window titled "Restore to Google Compute Engine" 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: Machine, Restore Mode, Data Center, Machine Type, Name, Network, Reason, and Summary. The "Summary" step is currently selected and highlighted. The main content area has a sub-header "Summary" with a green arrow icon and the text "You can copy the configuration information bellow for future reference." Below this, there is a large text box containing the following information:

```
Summary:
Project: RnD Backup
Data center: europe-north1 (Finland)
Availability zone: europe-north1-c

Items:
Original VM name: yam-v1inst-11-redhat-8
New VM name: new-yam-v1inst-11-redhat-8
Restore point: 9/23/2021 2:01:12 PM
Machine type: e2-medium
VPC: default
Subnet: default

Original VM name: ra-test-params-re
New VM name: new-ra-test-params-re
Restore point: 8/23/2021 2:00:10 PM
Machine type: f1-micro
VPC: default
Subnet: default
```

At the bottom of the main content area, there is a checkbox labeled "Power on target VM after restoring" which is checked. At the bottom of the window, there are four buttons: "< Previous", "Next >", "Finish" (which is highlighted with a blue border), and "Cancel".

Performing Guest OS File Recovery

Veeam Backup & Replication allows you to use image-level backups to restore files and folders of various VM 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 backup repositories for which you have specified HMAC keys associated with the service accounts that are used to access the repositories. 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 Google Cloud Web UI. For more information, see [Performing File-Level Recovery](#).

Restoring Files of 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 [Considerations 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 a VM 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 Files of 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 [Considerations and Limitations](#).

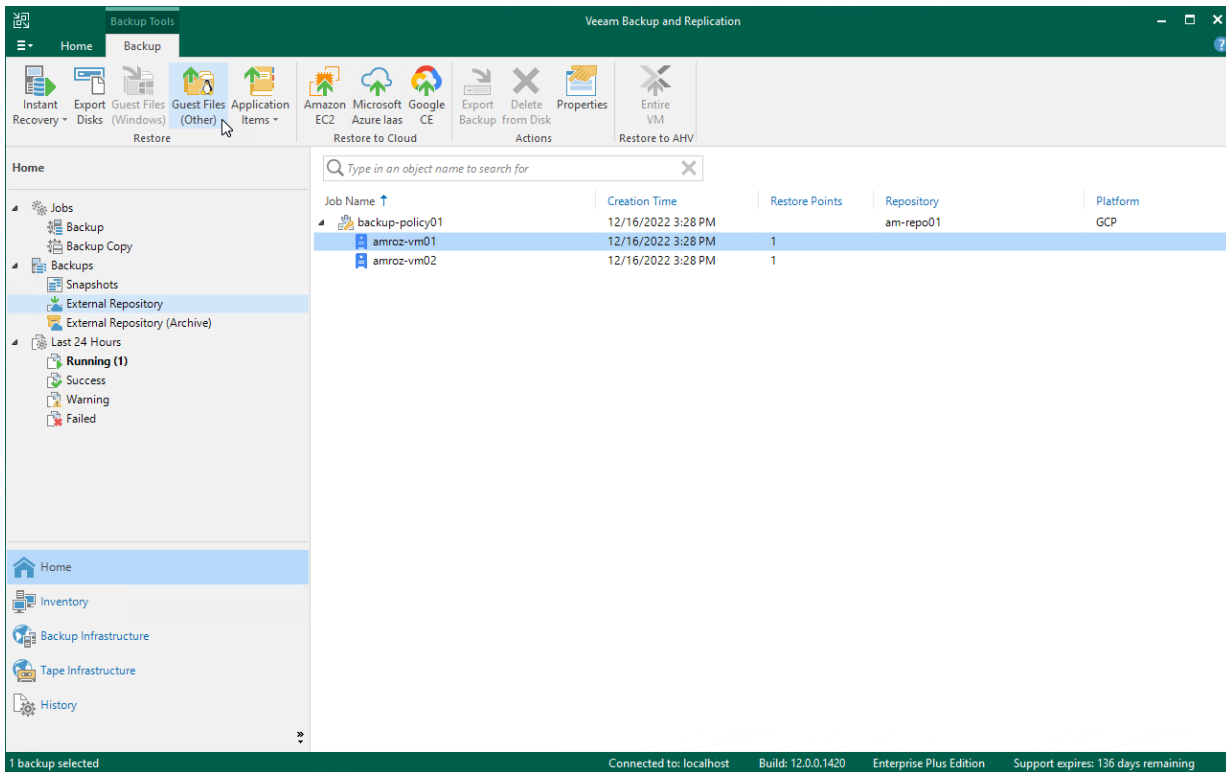
To restore guest OS files and folders, do the following:

1. In the Veeam Backup & Replication console, open the **Home** view.
2. Navigate to **Backups > External Repository**.
3. Expand the backup policy that protects a VM instance whose files and folders you want to restore, select the necessary instance and click **Guest Files (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](#).



VM Restore Using Web UI

Veeam Backup for Google Cloud offers the following restore operations:

- [Instance restore](#) – start an entire VM instance from a restore point.
- [Disk restore](#) – restore persistent disks attached to a VM instance.
- [File-level recovery](#) – recover individual files and folders of a VM instance.

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

Performing Application Restore

Veeam Backup & Replication provides auxiliary tools – Veeam Explorers – that allow you to restore application items directly from image-level backups of VM instances. You can restore items of the following applications: Microsoft Active Directory, Microsoft Exchange, Microsoft SharePoint, Microsoft SQL Server and Oracle Database. For more information, see the [Veeam Explorers User Guide](#).

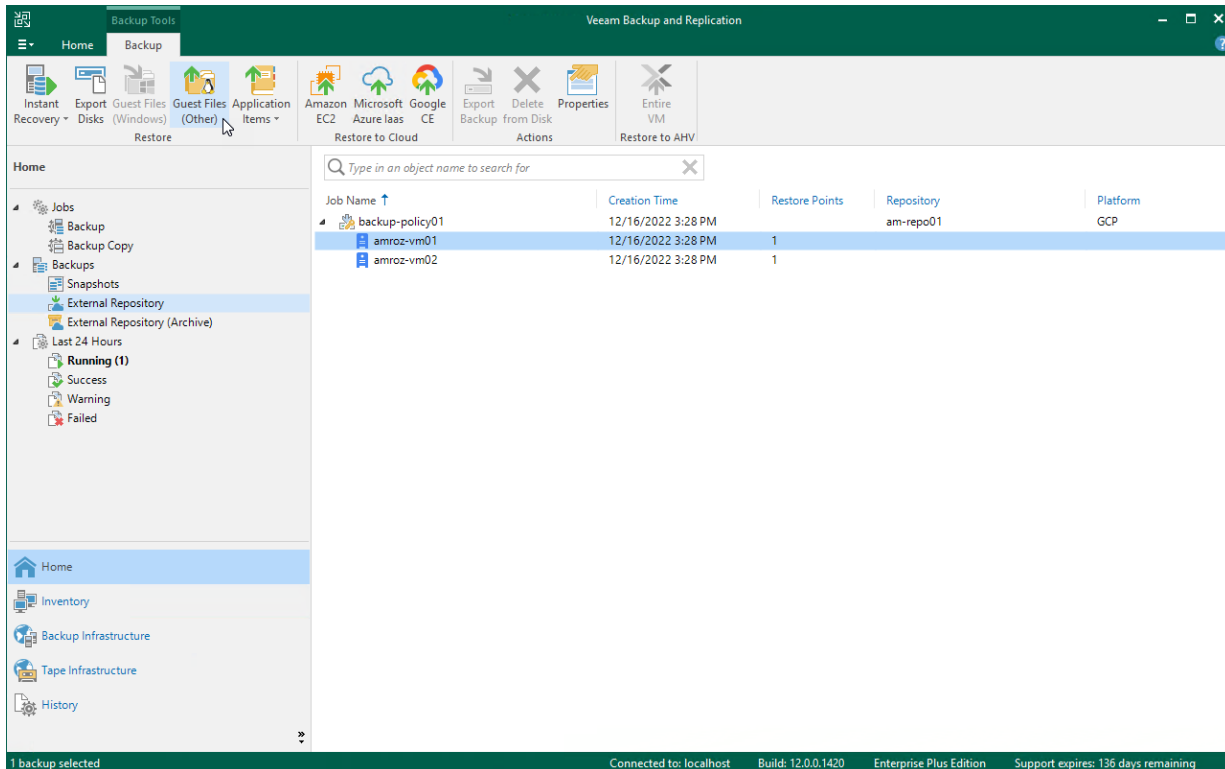
IMPORTANT

Application restore can be performed only using backup files stored in backup repositories for which you have specified HMAC keys associated with the service accounts that are used to access the repositories. 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**.
3. Expand the backup policy that protects a VM instance whose application item you want to restore, select the necessary instance and click **Application Items** on the ribbon. Then, select the necessary application.
4. In the restore wizard, select a restore point that will be used to restore the application, specify a restore reason and click **Browse**.

5. In the Veeam Explorer application, perform the steps described in the [Veeam Explorers User Guide](#).



Performing VM Instance Restore

In case a disaster strikes, you can restore an entire VM instance from a cloud-native snapshot or image-level backup. Veeam Backup for Google Cloud allows you to restore one or more VM instances at a time, to the original location or to a new location.

IMPORTANT

When restoring a VM instance, Veeam Backup for Google Cloud recovers data from all zonal and regional persistent disks (standard, balanced, extreme and SSD) attached to the instance. However, when it comes to local SSDs (SCSI and NVMe), Veeam Backup for Google Cloud is able to recover only the configuration of these disks due to [technical reasons](#).

To restore a protected VM instance, do the following:

1. [Launch the VM Instance Restore wizard](#).
2. [Select a restore point](#).
3. [Choose a restore mode](#).
4. [Select a service account](#).
5. [Select a project](#).
6. [Select a region and an availability zone](#).
7. [Enable encryption](#).
8. [Specify a new name and machine type for the instance](#).
9. [Configure network settings](#).

10. [Run configuration and permission checks.](#)
11. [Specify a restore reason.](#)
12. [Finish working with the wizard.](#)

IMPORTANT

Before you start VM instance restore, make sure that network settings are configured for each region where worker instances will be deployed during the restore process. For information on how to configure network settings, see [Adding Worker Configurations](#).

Step 1. Launch VM Instance Restore Wizard

To launch the **VM Instance Restore** wizard, do the following:

1. Navigate to **Protected Data > VM**.
2. Select the VM instance that you want to restore, and click **Restore > Instance Restore**.

The screenshot displays the Veeam Backup for Google Cloud interface. The top navigation bar includes the Veeam logo, the text "Veeam Backup for Google Cloud", the server time "Nov 20, 2023 9:14 PM", the user "administrator", and a "Configuration" button. The left sidebar shows a navigation menu with "Protected Data" selected. The main content area is titled "VM" and contains a search bar, a "Restore" dropdown menu (with "Instance Restore" highlighted), and a "Remove" button. Below this is a table of VM instances with columns for "Instance", "Project", "Restore Points", "Region", and "File-Level Recovery ...". The instance "dr-ga-backup-vm" is selected, and its "Instance Restore" option is highlighted in the dropdown menu. The table shows various VM instances with their respective projects, restore points, and regions.

Instance	Project	Restore Points	Region	File-Level Recovery ...
<input type="checkbox"/> atsvetkov-belgium	veeam-rnd-backup-4	—	europa-west1-b	—
<input type="checkbox"/> atsvetkov-belgium-v5-2	veeam-rnd-backup-4	—	europa-west1-b	—
<input type="checkbox"/> dr-asia-ne1	veeam-rnd-backup-4	—	asia-northeast1-b	—
<input type="checkbox"/> dr-asia-ne2	veeam-rnd-backup-4	asd	asia-northeast1-a	—
<input type="checkbox"/> dr-current-vm	veeam-rnd-backup-2	—	us-central1-a	—
<input type="checkbox"/> dr-fstabinstancer	veeam-rnd-backup-2	—	us-west3-a	—
<input checked="" type="checkbox"/> dr-ga-backup-vm	veeam-rnd-backup-4	—	us-east1-b	—
<input type="checkbox"/> dr-ga-backup-vm	veeam-rnd-backup-4	v4	us-east1-b	—
<input type="checkbox"/> dr-import2	veeam-rnd-backup-2	vms	us-central1-a	—
<input type="checkbox"/> dr-import2-1	veeam-rnd-backup-2	—	us-central1-a	—
<input type="checkbox"/> dr-import2-1asdasddd...	Scale Projects test 10	abigone	us-central1-a	—
<input type="checkbox"/> dr-import2asdasddd...	Scale Projects test 10	abigone	us-central1-a	—
<input type="checkbox"/> dr-scale-10	Scale Projects test 10	abigone	us-central1-a	—
<input type="checkbox"/> dr-scale-100	Scale Projects test 100	abigone	us-central1-a	—
<input type="checkbox"/> dr-scale-2-36	Scale Projects test 36	abigone	us-central1-a	—

Step 2. Select Restore Point

At the **Instances** step of the wizard, select a restore point that will be used to restore the selected VM instance. By default, Veeam Backup for Google Cloud uses the most recent valid restore point. However, you can restore the VM instance data to an earlier state.

To select a restore point, do the following:

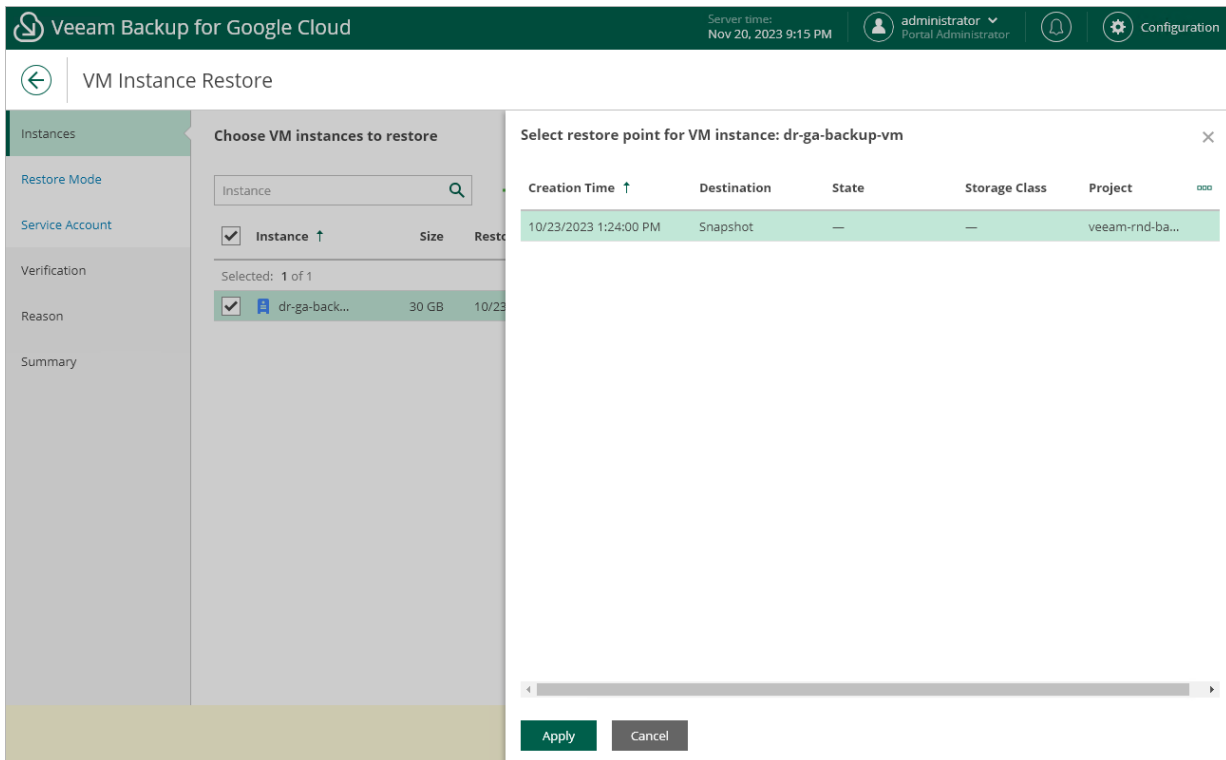
1. Select the VM instance.
2. Click **Restore Point**.
3. In the **Select restore point** window, select the necessary restore point and click **Apply**.

To help you choose a restore point, Veeam Backup for Google Cloud provides the following information on each available restore point:

- **Creation Time** – the date when the restore point was created.
- **Destination** – the type of the restore point:
 - § *Snapshot* – a cloud-native snapshot created by a backup policy.
 - § *Manual Snapshot* – a cloud-native snapshot created manually.
 - § *Backup* – an image-level backup created by a backup policy.
- **State** – the result of the latest health check performed for the restore point.
- **Storage Class** – the storage class of a backup repository where the restore point is stored (applies only to image-level backups).
- **Project** – a project that manages the protected VM instance.
- **Region** – a region in which the protected VM instance resides.
- **Retention** – a retention configured for the backup policy that created the restore point.

NOTE

You cannot restore entire VM instances using restore points in the *Incomplete* state. You can try running [disk restore](#) instead; however, the operation may fail to complete successfully.



The screenshot shows the Veeam Backup for Google Cloud interface. The main window is titled "VM Instance Restore". On the left, there is a sidebar with navigation options: "Instances", "Restore Mode", "Service Account", "Verification", "Reason", and "Summary". The "Instances" section is active, showing a search bar and a list of instances. One instance, "dr-ga-back...", is selected, with a size of 30 GB and a restore point of 10/23/2023 1:24:00 PM. A modal window titled "Select restore point for VM instance: dr-ga-backup-vm" is open, displaying a table with the following data:

Creation Time ↑	Destination	State	Storage Class	Project
10/23/2023 1:24:00 PM	Snapshot	—	—	veeam-rnd-ba...

At the bottom of the modal window, there are "Apply" and "Cancel" buttons.

Step 3. Choose Restore Mode

At the **Restore Mode** step of the wizard, choose whether you want to restore the selected VM instance to the original or to a custom location.

TIP

If restore to the original location is not available, the wizard will display a message notifying that some of the selected VM instances have issues with the original settings. To learn what these issues are, hover the mouse cursor over the message.

The screenshot shows the 'VM Instance Restore' wizard in Veeam Backup for Google Cloud. The interface includes a top navigation bar with the product name, server time (Nov 20, 2023 9:16 PM), user information (administrator, Portal Administrator), and a Configuration icon. The main content area is titled 'VM Instance Restore' and features a left-hand navigation menu with options: Instances, Restore Mode (highlighted), Project, Region, Encryption, Instance Settings, Network Settings, Verification, Reason, and Summary. The 'Choose restore mode' section contains two radio button options: 'Restore to original location, with original settings' (unselected) and 'Restore to new location, or with different settings' (selected). The 'Restore to original location' option includes a sub-note: 'Quickly restore the selected VM instances to their original location, with the same name and settings as the source instances.' The 'Restore to new location' option includes a sub-note: 'Restore the selected VM instances to a new location or use different configuration settings.' At the bottom of the wizard, there are three buttons: 'Previous', 'Next', and 'Cancel'.

Step 4. Select Service Account

[This step applies only if you have selected the **Restore to original location, with original settings** option at the **Restore Mode** step of the wizard]

At the **Account** step of the wizard, select a service account whose permissions will be used to perform the restore operation. For more information on the required permissions, see [Service Account Permissions](#).

For a service account to be displayed in the **Service account** drop-down list, it must be added to Veeam Backup for Google Cloud as described in section [Adding Service Accounts](#), and must be assigned the *VM Instances Restore* operational role as described in section [Adding Projects and Folders](#).

If you have not added the necessary service account to Veeam Backup for Google Cloud beforehand, you can do it without closing the **VM Instance Restore** wizard. To add an account, click **Add** and complete the **Add Service Account** wizard.

The screenshot shows the 'VM Instance Restore' wizard in the Veeam Backup for Google Cloud interface. The top navigation bar includes the Veeam logo, 'Veeam Backup for Google Cloud', the server time 'Nov 20, 2023 9:16 PM', and user information 'administrator Portal Administrator'. The wizard progress bar on the left highlights 'Service Account' as the current step, with other steps like 'Instances', 'Restore Mode', 'Verification', 'Reason', and 'Summary' visible. The main content area is titled 'Choose service account' and contains the instruction 'Specify a service account to be used to access resources for the restore operation.' Below this, there is a 'Service account:' label followed by a dropdown menu showing 'veeam-1649186685-sa@rmd-backup-2.iam.gservice...' and a '+ Add' button. A second, identical instruction and label are present below. At the bottom of the wizard, there are three buttons: 'Previous', 'Next', and 'Cancel'.

Step 5. Select Project

[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 **Project** step of the wizard, select a project that will be used to manage the restored VM instance and specify a service account whose permissions will be used to perform the restore operation. For more information on the required permissions, see [Service Account Permissions](#).

For a project to be displayed in the **Project** drop-down list, it must be added to Veeam Backup for Google Cloud as described in section [Adding Projects and Folders](#). If you have not added the necessary project to Veeam Backup for Google Cloud beforehand, you can do it without closing the **VM Instance Restore** wizard. To add a project, click **Add** and complete the **Add Projects and Folders** wizard.

For a service account to be displayed in the list of available accounts, it must be added to Veeam Backup for Google Cloud as described in section [Adding Service Accounts](#), and must be assigned permissions required to access the selected project as described in section [Adding Projects and Folders](#).

The screenshot shows the 'VM Instance Restore' wizard in the Veeam Backup for Google Cloud interface. The 'Project' step is selected in the left-hand navigation pane. The main content area is titled 'Choose project and service account' and contains two sections:

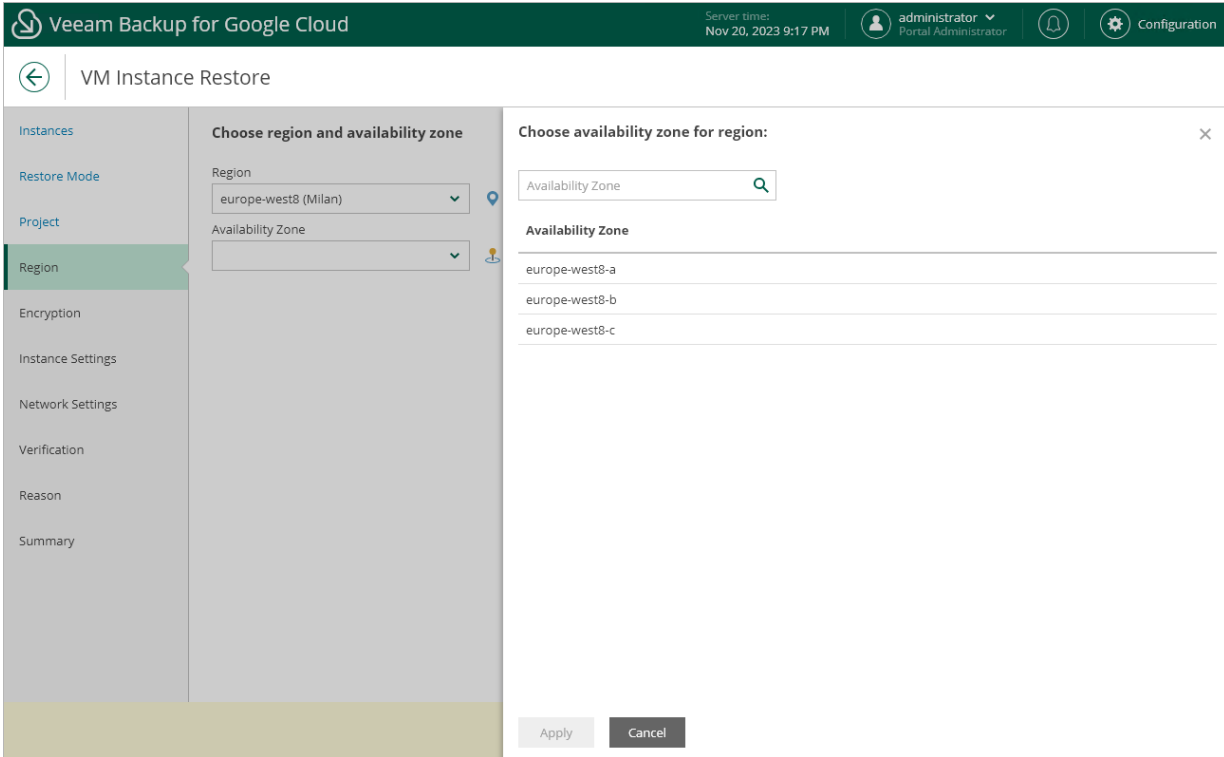
- Project:** A text box with the label 'Project' and a description: 'Specify a project where the restored VM instances will be created.' Below this is a dropdown menu showing 'veeam-rnd-backup-2 (rnd-backup-2)' and a '+ Add' button.
- Service account:** A text box with the label 'Service account' and a description: 'Specify a service account to be used to access resources for the restore operation.' Below this is a dropdown menu showing 'veeam-1649186685-sa@rnd-backup-2.lam.gservice...'

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

Step 6. Select Region and Availability Zone

[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 **Region** step of the wizard, select a region where the restored VM instance will operate and an availability zone for which you want to configure network settings.



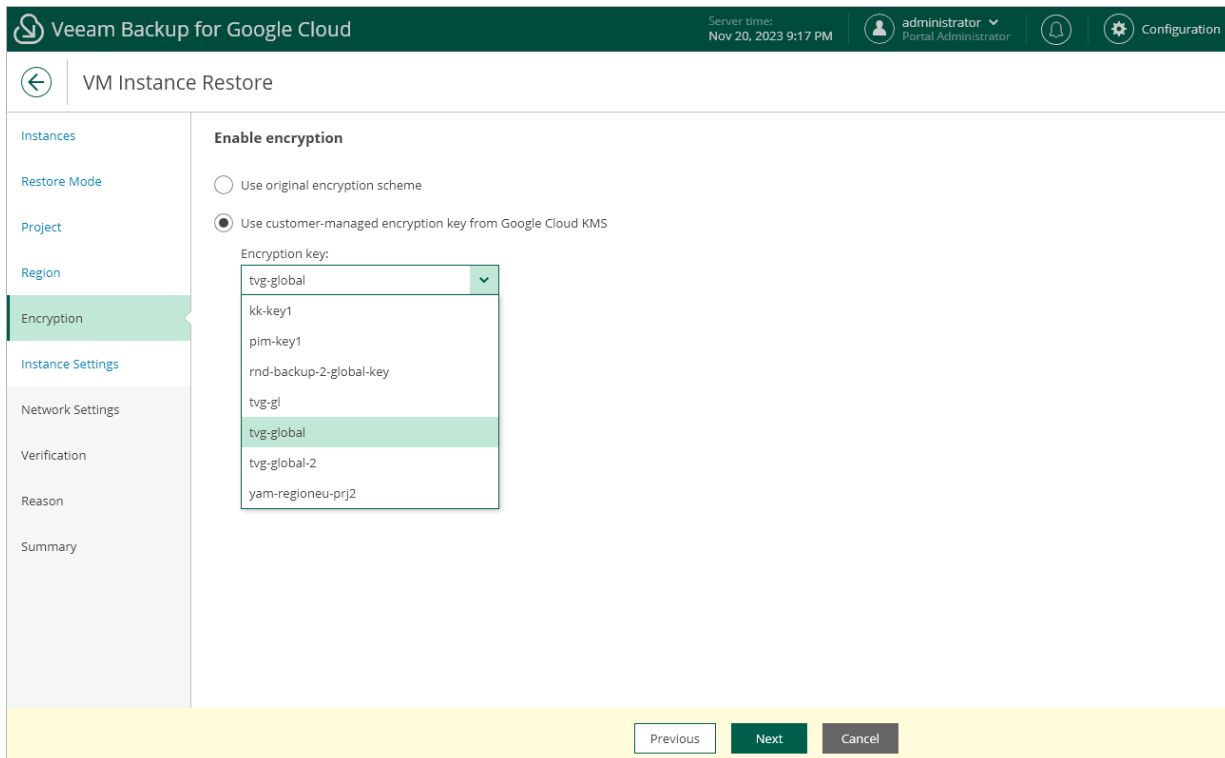
Step 7. Enable Encryption

[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, do the following:

- If you want to apply the existing encryption scheme of the source VM instance, select the **Use original encryption scheme** option.
- If you want to encrypt persistent disks of the restored VM instance with a Google Cloud KMS CMEK, select the **Use customer-managed encryption key from Google Cloud KMS** option and choose the necessary CMEK from the **Encryption key** drop-down list.

For a CMEK to be displayed in the list of available encryption keys, it must be stored in the region selected at [step 6](#) of the wizard.



Step 8. Specify Instance Name and Type

[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 **Instance Settings** step of the wizard, do the following:

1. Select the VM instance.
2. If you want to specify a new name and a new machine type for the restored VM instance, click **Edit**.

In the **Configure settings** window, specify the name and the machine type, and click **Apply**. To learn how to choose a machine type when creating a VM instance in Google Cloud, see [Google Cloud documentation](#).

TIP

If Veeam Backup for Google Cloud is unable to restore the VM instance using the specified name for some reason, the wizard will display an error icon in the **Instance** column. To learn what this reason is, hover your mouse over the icon.

The screenshot shows the Veeam Backup for Google Cloud interface. At the top, the title bar reads "Veeam Backup for Google Cloud" with a server time of "Nov 20, 2023 9:18 PM" and a user profile for "administrator". The main window is titled "VM Instance Restore". On the left, a navigation pane lists steps: "Instances", "Restore Mode", "Project", "Region", "Encryption", "Instance Settings" (highlighted), "Network Settings", "Verification", "Reason", and "Summary". The "Instance Settings" step is active, displaying a table for VM instances to be restored. The table has columns for "Instance" and "Machine Type". One instance is listed: "dr-ga-backup-vm" with a red error icon in the "Instance" column and "c2-standard-4" in the "Machine Type" column. A dialog box titled "Configure settings for VM instance: dr-ga-backup-vm" is open. It has a "Name" field containing "dr-ga-backup-vm-1" and a "Type" field with a search icon. Below these fields is a list of machine types with a scroll bar. The selected machine type is "e2-highcpu-32 (32 cores, 32GB)". Other visible machine types include "e2-highcpu-2 (2 cores, 2GB)", "e2-highcpu-4 (4 cores, 4GB)", "e2-highcpu-8 (8 cores, 8GB)", "e2-highcpu-16 (16 cores, 16GB)", "e2-highmem-2 (2 cores, 16GB)", "e2-highmem-4 (4 cores, 32GB)", "e2-highmem-8 (8 cores, 64GB)", "e2-highmem-16 (16 cores, 128GB)", and "e2-medium (2 cores, 4GB)". At the bottom of the dialog are "Apply" and "Cancel" buttons. A "Previous" button is visible at the bottom of the main window.

Step 9. 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 Settings** step of the wizard, do the following:

1. Select the VM instance.
2. Click **Edit**.
3. In the **Edit network settings** window, select a VPC network and a subnet to which the restored VM instance will be connected. You can also choose whether you want the restored VM instance to have the same reserved static external IP address and the same network tags as the source VM instance.

For a VPC network and a subnet to be displayed in the lists of available networks, they must be created in the Google Cloud console for the region specified at [step 6](#) of the wizard, as described in [Google Cloud documentation](#).

NOTE

Veeam Backup for Google Cloud cannot assign a static external IP address to a restored VM instance if the source instance does not have the address reserved. To learn how to reserve static external IP addresses for VM instances, see [Google Cloud documentation](#).

The screenshot displays the Veeam Backup for Google Cloud interface during the 'VM Instance Restore' process. The main window is titled 'VM Instance Restore' and shows a sidebar with navigation options: Instances, Restore Mode, Project, Region, Encryption, Instance Settings, Network Settings (highlighted), Verification, Reason, and Summary. The 'Configure network settings' section is active, showing a table with one instance selected. The 'Edit network settings for VM instance: dr-ga-backup-vm-1' dialog is open, showing options to select a VPC network and subnet, and checkboxes for 'Recover static external IP address' and 'Recover network tags'. A yellow highlight is under the 'Recover network tags' checkbox and its associated information message.

Instance	VPC Network	Subnet	Static IP
<input checked="" type="checkbox"/> dr-ga-backu...	—	—	Off

Selected: 1 of 1

Dialog: Edit network settings for VM instance: dr-ga-backup-vm-1

VPC network: tvg-net

Subnet: tvg-net

Recover static external IP address

Recover network tags

Network tags are used to apply firewall rules to the target VPC network

Buttons: Apply, Cancel, Previous

Step 10. Run Configuration Checks

At the **Verification** step of the wizard, Veeam Backup for Google Cloud will verify whether restore settings are configured properly and the specified service account has all the necessary permissions required to perform recovery tasks for the project that will manage the restored VM instance. For more information on the required permissions, see [Service Account Permissions](#).

To see the list of missing permissions that must be granted to the service account in order to perform an operation, click the link in the **Details** column. You can grant the missing permissions to the service account [using the Google Cloud console](#) or instruct Veeam Backup for Google Cloud to do it:

- To grant the missing permissions manually, click **Download Script**. Veeam Backup for Google Cloud will generate a gcloud script that you can run in the Google Cloud console to assign all the necessary permissions to the service account.

The account under which you run the script must have the permissions required both to get and set project IAM policies and to create custom IAM roles (for example, it can have the *iam.securityAdmin* and *iam.roleAdmin* roles assigned). To learn what permissions and roles are required to create custom roles in IAM, see [Google Cloud documentation](#).

- To let Veeam Backup for Google Cloud grant the missing permissions automatically, click **Grant** and then click **Sign in with Google** in the **Grant permissions** window. You will be redirected to the OAuth consent screen authorization page. Sign in using credentials of a Google account that will be used to grant the permissions.

The account under which you sign in to Google Cloud must have the permissions required both to get and set project IAM policies and to create custom IAM roles (for example, it can have the *iam.securityAdmin* and *iam.roleAdmin* roles assigned). To learn what permissions and roles are required to create service account, see [Google Cloud documentation](#).

NOTE

For Veeam Backup for Google Cloud to be able to authorize in Google Cloud, the OAuth consent screen must be configured as described in section [Registering Applications](#). Note that Veeam Backup for Google Cloud does not store in the configuration database the provided Google account credentials and access tokens received during authorization.

To make sure that the missing permissions have been successfully granted, click **Recheck**.

The screenshot shows the Veeam Backup for Google Cloud interface. The top navigation bar includes the Veeam logo, the product name, the server time (Nov 20, 2023 9:19 PM), the user role (administrator Portal Administrator), and a Configuration icon. The main content area is titled "VM Instance Restore" and features a left-hand navigation menu with options like Instances, Restore Mode, Project, Region, Encryption, Instance Settings, Network Settings, Verification (highlighted), Reason, and Summary. The main panel displays "Run verification checks" with a sub-header "Verify that permissions and configuration are correct." Below this are three action buttons: Recheck, Download Script, and Grant. A table lists the verification checks:

Check	Result	Details
VM Restore	Passed	All the required permissions are granted.
Worker	Passed	All the required permissions are granted.
Repository	Passed	All the required permissions are granted.

At the bottom of the interface, there are three buttons: "Previous", "Next", and "Cancel".

Step 11. Specify Restore Reason

At the **Reason** step of the wizard, specify a reason for restoring the VM instance. This information will be saved to the session history, and you will be able to reference it later.

The screenshot shows the 'VM Instance Restore' wizard in the Veeam Backup for Google Cloud interface. The top navigation bar includes the Veeam logo, the product name, the server time (Nov 20, 2023 9:19 PM), the user role (administrator Portal Administrator), and a Configuration icon. The main content area is titled 'VM Instance Restore' and features a left-hand navigation pane with the following steps: Instances, Restore Mode, Project, Region, Encryption, Instance Settings, Network Settings, Verification, Reason (highlighted in green), and Summary. The main pane is titled 'Enter reason for this restore operation' and contains a text input field labeled 'Restore reason:' with the text 'evaluating instance restore' entered. At the bottom of the wizard, there are three buttons: 'Previous' (disabled), 'Next' (active), and 'Cancel' (disabled).

Step 12. Finish Working with Wizard

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

TIP

If you want to keep the restored VM instance running as soon as the restore process completes, select the **Power on target VM instances after restore** check box. Otherwise, the instance will be powered off.

The screenshot shows the 'VM Instance Restore' wizard in the 'Summary' step. The interface includes a top navigation bar with the Veeam logo, 'Veeam Backup for Google Cloud', server time (Nov 20, 2023 9:19 PM), user (administrator), and configuration icons. A left sidebar lists steps: Instances, Restore Mode, Project, Region, Encryption, Instance Settings, Network Settings, Verification, Reason, and Summary (highlighted). The main content area is titled 'Review configured settings' and contains a 'Copy to Clipboard' button. It is divided into three sections: 'General' (Project name: veeam-rnd-backup-2, Project ID: rnd-backup-2, Service account: veeam-1649186685-sa@rnd-backup-2.iam.gserviceaccount.com, Restore mode: New location, Region: europe-west8 (Milan), Availability zone: europe-west8-b, Reason: evaluating instance restore), 'Restore settings' (VM instances: 1 instance, Encryption: tvq-global), and 'Validation' (Permission check: Passed, Instance settings: Passed). At the bottom, there is a checkbox for 'Power on target VM instances after restore' and three buttons: 'Previous', 'Finish', and 'Cancel'.

Performing Disk Restore

In case a disaster strikes, you can restore corrupted persistent disks of a VM instance from a cloud-native snapshot or image-level backup. Veeam Backup for Google Cloud allows you to restore persistent disks to the original location or to a new location.

IMPORTANT

You can restore zonal and regional persistent disks of all types: standard (pd-standard), balanced (pd-balanced), extreme (pd-extreme) and SSD (pd-ssd). Restore of local SSDs (SCSI and NVMe) is not supported due to [technical reasons](#).

To restore persistent disks attached to a protected VM instance, do the following:

1. [Launch the Disk Restore wizard](#).
2. [Select a restore point](#).
3. [Choose a restore mode](#).
4. [Select a service account](#).

5. [Select a project.](#)
6. [Select a region and an availability zone.](#)
7. [Enable encryption.](#)
8. [Specify new names for the disks.](#)
9. [Run configuration and permission checks.](#)
10. [Specify a restore reason.](#)
11. [Finish working with the wizard.](#)

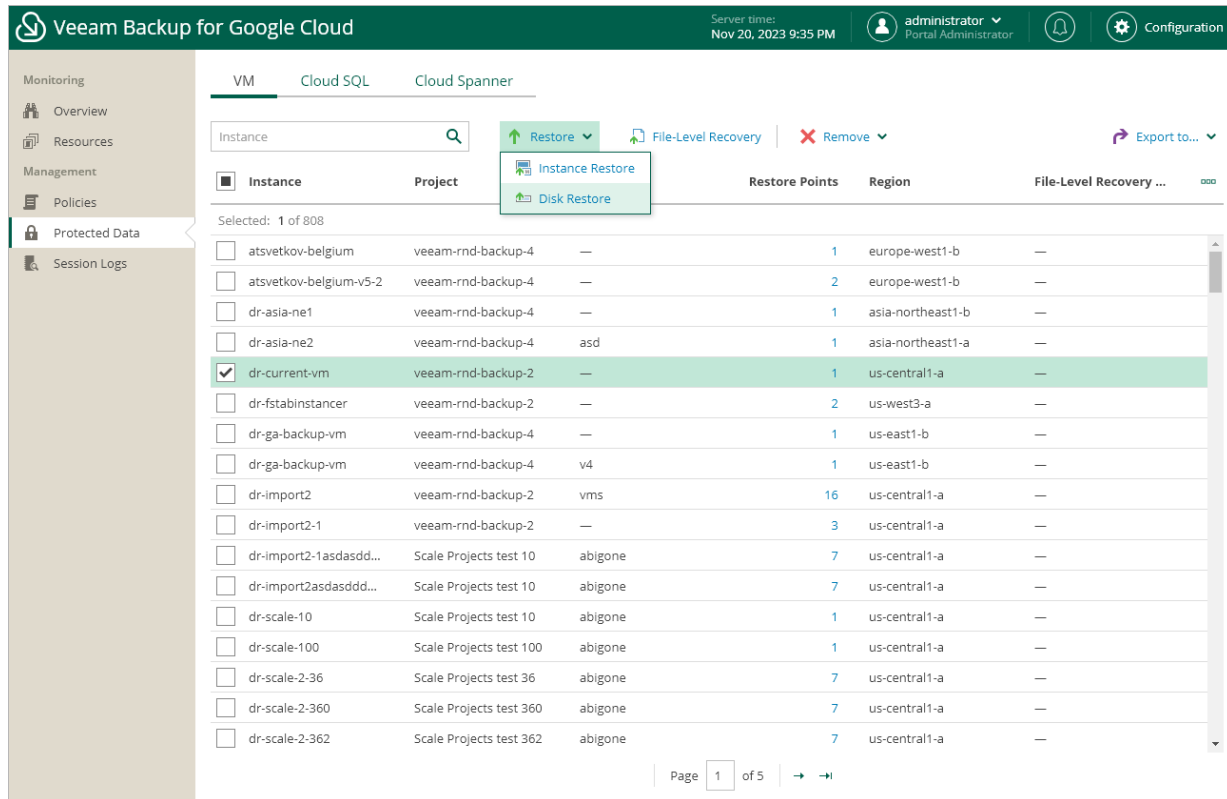
IMPORTANT

Before you start disk restore, make sure that network settings are configured for each region where worker instances will be deployed during the restore process. For information on how to configure network settings, see [Adding Worker Configurations](#).

Step 1. Launch Disk Restore Wizard

To launch the **Disk Restore** wizard, do the following:

1. Navigate to **Protected Data > VM**.
2. Select the VM instance whose persistent disks you want to restore, and click **Restore > Disk Restore**.



The screenshot shows the Veeam Backup for Google Cloud interface. The top navigation bar includes the Veeam logo, the text "Veeam Backup for Google Cloud", the server time "Nov 20, 2023 9:35 PM", and user information "administrator Portal Administrator". The left sidebar contains navigation options: Monitoring, Overview, Resources, Management, Policies, Protected Data (selected), and Session Logs. The main content area is titled "VM" and shows a list of VM instances. The "Instance" column is selected, and the "Disk Restore" option is highlighted in the "Restore" dropdown menu. The table below shows the list of VM instances with columns for Instance, Project, Restore Points, Region, and File-Level Recovery.

Instance	Project	Restore Points	Region	File-Level Recovery ...
<input type="checkbox"/> atsvetkov-belgium	veeam-rnd-backup-4	—	europa-west1-b	—
<input type="checkbox"/> atsvetkov-belgium-v5-2	veeam-rnd-backup-4	—	europa-west1-b	—
<input type="checkbox"/> dr-asia-ne1	veeam-rnd-backup-4	—	asia-northeast1-b	—
<input type="checkbox"/> dr-asia-ne2	veeam-rnd-backup-4	asd	asia-northeast1-a	—
<input checked="" type="checkbox"/> dr-current-vm	veeam-rnd-backup-2	—	us-central1-a	—
<input type="checkbox"/> dr-fstabinstancer	veeam-rnd-backup-2	—	us-west3-a	—
<input type="checkbox"/> dr-ga-backup-vm	veeam-rnd-backup-4	—	us-east1-b	—
<input type="checkbox"/> dr-ga-backup-vm	veeam-rnd-backup-4	v4	us-east1-b	—
<input type="checkbox"/> dr-import2	veeam-rnd-backup-2	vms	us-central1-a	—
<input type="checkbox"/> dr-import2-1	veeam-rnd-backup-2	—	us-central1-a	—
<input type="checkbox"/> dr-import2-1asdasddd...	Scale Projects test 10	abigone	us-central1-a	—
<input type="checkbox"/> dr-import2asdasddd...	Scale Projects test 10	abigone	us-central1-a	—
<input type="checkbox"/> dr-scale-10	Scale Projects test 10	abigone	us-central1-a	—
<input type="checkbox"/> dr-scale-100	Scale Projects test 100	abigone	us-central1-a	—
<input type="checkbox"/> dr-scale-2-36	Scale Projects test 36	abigone	us-central1-a	—
<input type="checkbox"/> dr-scale-2-360	Scale Projects test 360	abigone	us-central1-a	—
<input type="checkbox"/> dr-scale-2-362	Scale Projects test 362	abigone	us-central1-a	—

Page 1 of 5

Step 2. Select Restore Point

At the **Instances** step of the wizard, select a restore point that will be used to restore persistent disks of the selected VM instance. By default, Veeam Backup for Google Cloud uses the most recent valid restore point. However, you can restore the disks to an earlier state.

To select a restore point, do the following:

1. Select the VM instance.
2. Click **Choose Restore Point**.
3. In the **Select restore point** window, select the necessary restore point and click **Apply**.

To help you choose a restore point, Veeam Backup for Google Cloud provides the following information on each available restore point:

- **Creation Time** – the date when the restore point was created.
- **Destination** – the type of the restore point:
 - § *Snapshot* – a cloud-native snapshot created by a backup policy.
 - § *Manual Snapshot* – a cloud-native snapshot created manually.
 - § *Backup* – an image-level backup created by a backup policy.
- **State** – the result of the latest health check performed for the restore point.
- **Storage Class** – the storage class of a backup repository where the restore point is stored (applies only to image-level backups).
- **Project** – a project that manages the protected VM instance.
- **Region** – a region in which the protected VM instance resides.
- **Retention** – a retention configured for the backup policy that created the restore point.

TIP

If you want to restore only specific persistent disks of the selected VM, you can exclude the unnecessary disks from the restore process. To do that, click **Exclusions** to open the **Exclude disks from restore** window, select check boxes next to the disks that you do not want to restore, and click **Apply**.

The screenshot shows the Veeam Backup for Google Cloud interface. The top bar includes the Veeam logo, server time (Nov 20, 2023 9:37 PM), user profile (administrator), and configuration options. The main window is titled "Disk Restore" and is divided into several sections. On the left, there is a sidebar with navigation options: Instances, Restore Mode, Service Account, Verification, Reason, and Summary. The central area is titled "Select VM instance" and contains a "Choose Restore Point" button and an "Exclusions" button. Below these, a table lists VM instances with columns for Instance, Size, and Restore Point. One instance, "dr-current...", is highlighted. On the right, a modal window titled "Select restore point for VM instance: dr-current-vm" is open, displaying a table with columns for Creation Time, Destination, State, Storage Class, and Project. A single row is visible with a creation time of 11/10/2023 5:35:41 PM, a destination of Snapshot, and a project of veeam-rnd-ba... At the bottom of the modal, there are "Apply" and "Cancel" buttons.

Step 3. Choose Restore Mode

At the **Restore Mode** step of the wizard, choose whether you want to restore persistent disks of the selected VM instance to the original or to a custom location.

TIP

If restore to the original location is not available, the wizard will display a message notifying that some of the selected disks have issues with the original settings. To learn what these issues are, hover the mouse cursor over the message.

The screenshot shows the 'Disk Restore' wizard in Veeam Backup for Google Cloud. 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 product name 'Veeam Backup for Google Cloud', the server time 'Nov 20, 2023 9:37 PM', the user 'administrator Portal Administrator', and a 'Configuration' link. The left sidebar contains a list of steps: 'Instances', 'Restore Mode' (highlighted), 'Project', 'Region', 'Encryption', 'Disk Settings', 'Verification', 'Reason', and 'Summary'. The main content area is titled 'Choose restore mode' and contains two radio button options. The first option, 'Restore to original location, with original settings', is unselected and includes a sub-note: 'Quickly restore disks of the selected VM instances to their original location, with the same names and settings as the source disks.' The second option, 'Restore to new location, or with different settings', is selected and includes a sub-note: 'Restore disks of the selected VM instances to a new location or use different configuration settings.' At the bottom of the wizard, there are three buttons: 'Previous', 'Next' (highlighted in green), and 'Cancel'.

Step 4. Select Service Account

[This step applies only if you have selected the **Restore to original location, with original settings** option at the **Restore Mode** step of the wizard]

At the **Service Account** step of the wizard, select a service account whose permissions will be used to perform the restore operation. For more information on the required permissions, see [Service Account Permissions](#).

For a service account to be displayed in the **Service account** drop-down list, it must be added to Veeam Backup for Google Cloud for Google Cloud as described in section [Adding Service Accounts](#), and must be assigned the *VM Instances Restore* operational role as described in section [Adding Projects and Folders](#).

If you have not added the necessary service account to Veeam Backup for Google Cloud beforehand, you can do it without closing the **VM Instance Restore** wizard. To add an account, click **Add** and complete the **Add Service Account** wizard.

The screenshot shows the 'Disk Restore' wizard in Veeam Backup for Google Cloud. The interface is dark-themed with a green header. The header includes the Veeam logo, the product name 'Veeam Backup for Google Cloud', the server time 'Nov 20, 2023 9:37 PM', and the user 'administrator Portal Administrator'. A navigation bar on the left contains a back arrow and the title 'Disk Restore'. Below this is a vertical sidebar with menu items: 'Instances', 'Restore Mode', 'Service Account' (highlighted in green), 'Verification', 'Reason', and 'Summary'. The main content area is titled 'Specify service account' and contains the instruction 'Specify a service account that will be used to access the original project.' Below this, there is a 'Service account:' label followed by a dropdown menu showing 'veeam-1649186685-sa@rnd-backup-2.iam.gservice...' and a '+ Add' button. At the bottom of the wizard, there are three buttons: 'Previous', 'Next' (highlighted in green), and 'Cancel'.

Step 5. Select Project

[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 **Project** step of the wizard, select a project to which the restored persistent disks will belong and specify a service account whose permissions will be used to perform the restore operation. For more information on the required permissions, see [Service Account Permissions](#).

For a project to be displayed in the **Project** drop-down list, it must be added to Veeam Backup for Google Cloud as described in section [Adding Projects and Folders](#). If you have not added the necessary project to Veeam Backup for Google Cloud beforehand, you can do it without closing the **VM Instance Restore** wizard. To add a project, click **Add** and complete the **Add Projects and Folders** wizard.

For a service account to be displayed in the list of available accounts, it must be added to Veeam Backup for Google Cloud as described in section [Adding Service Accounts](#), and must be assigned permissions required to access the selected project as described in section [Adding Projects and Folders](#).

Veeam Backup for Google Cloud

Server time: Nov 20, 2023 9:37 PM administrator Portal Administrator Configuration

← Disk Restore

Instances

Restore Mode

Project

Region

Encryption

Disk Settings

Verification

Reason

Summary

Specify project

Project

Choose a target project.

Project: veeam-rnd-backup-2 (rnd-backup-2) + Add

Service account

Specify a service account that will be used to access the project.

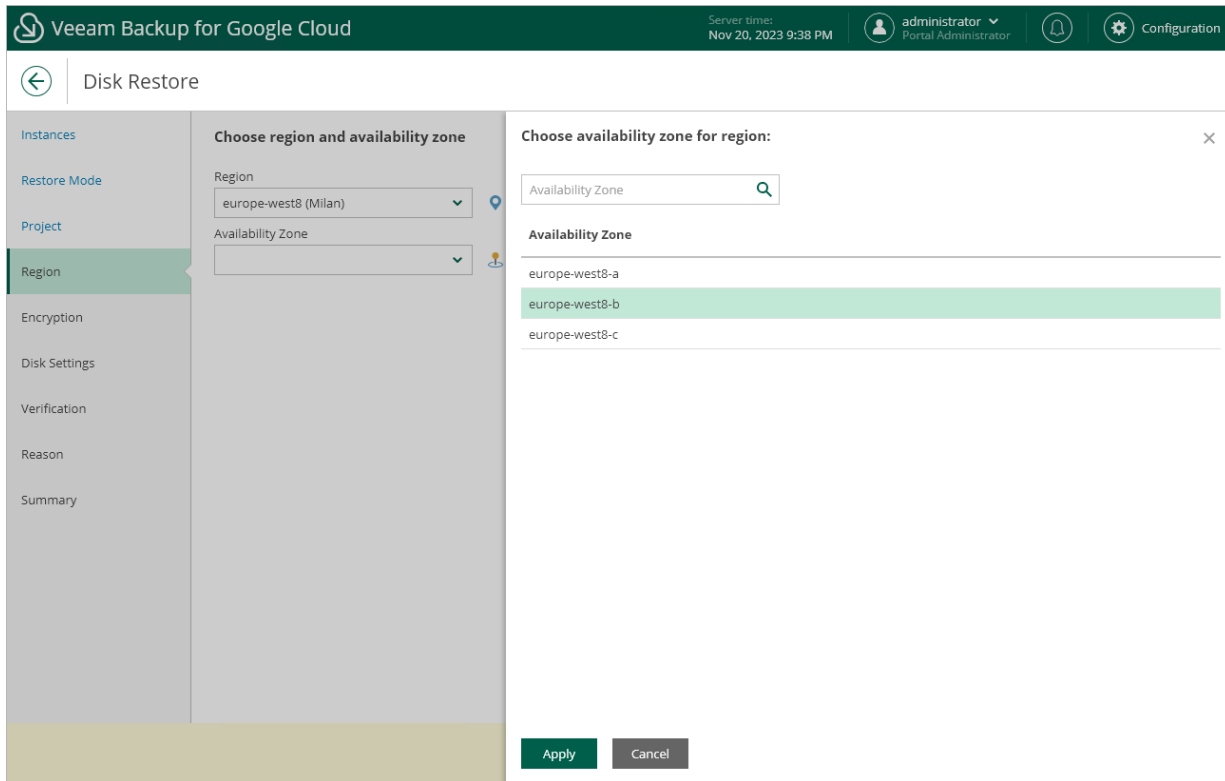
Service account: veeam-1649186685-sa@rnd-backup-2.iam.gservice...

Previous Next Cancel

Step 6. Select Region and Availability Zone

[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 **Region** step of the wizard, select a region and an availability zone to which the restored persistent disks will be placed.




Step 7. 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, do the following:

- If you do not want to change the existing encryption scheme of the restored persistent disks, select the **Use original encryption scheme** option.
- If you want to encrypt the restored persistent disks with a Google Cloud KMS CMEK, select the **Use customer-managed encryption key from Google Cloud KMS** option and choose the necessary CMEK from the **Encryption key** drop-down list.

For a CMEK to be displayed in the list of available encryption keys, it must be stored in the region selected at [step 6](#) of the wizard.



Step 8. Specify Disk Names

[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 **Disk Settings** step of the wizard, you can specify a new name for each restored persistent disk:

1. Select the necessary disk and click **Rename**.
2. In the **Rename disk** window, specify a name for the disk and click **Apply**.

TIP

If Veeam Backup for Google Cloud is unable to restore the disk using the specified name for some reason, the wizard will display a warning icon in the **Disk** column. To learn what this reason is, hover your mouse over the icon.

The screenshot shows the Veeam Backup for Google Cloud interface. The top navigation bar includes the Veeam logo, the text 'Veeam Backup for Google Cloud', the server time 'Nov 20, 2023 9:39 PM', and user information 'administrator Portal Administrator'. The main content area is titled 'Disk Restore' and has a sidebar with navigation options: Instances, Restore Mode, Project, Region, Encryption, Disk Settings (selected), Verification, Reason, and Summary. The 'Specify disk names' section is active, displaying a warning message: 'Due to technical limitations, the source and restored disks cannot have the same name. If a disk with the same name listed below still exists in the environment, Veeam Backup for Google Cloud will append a suffix to the name of the source disk when naming the restored disk (for more information, see the documentation). You can also specify any other name manually.' Below this message is a search box for 'Disk' and a 'Rename' button. A table lists disks with columns 'Disk' and 'Device'. The second row is highlighted, showing 'dr-current-vm-1671441087-boot-dis...' and 'dr-current-vm'. A modal window titled 'Rename disk: dr-new-testing-vm-1649186685-data-disk-restored-0' is open, with a 'Name:' field containing 'dr-new-testing-vm-1649186685-data-disk-restored-0'. At the bottom of the modal are 'Apply' and 'Cancel' buttons. The background interface also has 'Previous' and 'Apply' buttons at the bottom.

Step 9. Run Configuration Checks

At the **Verification** step of the wizard, Veeam Backup for Google Cloud will verify whether restore settings are configured properly and the specified service account has all the necessary permissions required to perform recovery tasks for the project to which the restored persistent disks will belong. For more information on the required permissions, see [Service Account Permissions](#).

To see the list of missing permissions that must be granted to the service account in order to perform an operation, click the link in the **Details** column. You can grant the missing permissions to the service account [using the Google Cloud console](#) or instruct Veeam Backup for Google Cloud to do it:

- To grant the missing permissions manually, click **Download Script**. Veeam Backup for Google Cloud will generate a gcloud script that you can run in the Google Cloud console to assign all the necessary permissions to the service account.

The account under which you run the script must have the permissions required both to get and set project IAM policies and to create custom IAM roles (for example, it can have the *iam.securityAdmin* and *iam.roleAdmin* roles assigned). To learn what permissions and roles are required to create custom roles in IAM, see [Google Cloud documentation](#).

- To let Veeam Backup for Google Cloud grant the missing permissions automatically, click **Grant** and then click **Sign in with Google** in the **Grant permissions** window. You will be redirected to the OAuth consent screen authorization page. Sign in using credentials of a Google account that will be used to grant the permissions.

The account under which you sign in to Google Cloud must have the permissions required both to get and set project IAM policies and to create custom IAM roles (for example, it can have the *iam.securityAdmin* and *iam.roleAdmin* roles assigned). To learn what permissions and roles are required to create service account, see [Google Cloud documentation](#).

NOTE

For Veeam Backup for Google Cloud to be able to authorize in Google Cloud, the OAuth consent screen must be configured as described in section [Registering Applications](#). Note that Veeam Backup for Google Cloud does not store in the configuration database the provided Google account credentials and access tokens received during authorization.

To make sure that the missing permissions have been successfully granted, click **Recheck**.

The screenshot shows the Veeam Backup for Google Cloud interface. At the top, the header includes the Veeam logo, the product name 'Veeam Backup for Google Cloud', the server time 'Nov 20, 2023 9:39 PM', and the user 'administrator Portal Administrator'. The main content area is titled 'Disk Restore' and features a left-hand navigation menu with options like 'Instances', 'Restore Mode', 'Project', 'Region', 'Encryption', 'Disk Settings', 'Verification', 'Reason', and 'Summary'. The 'Verification' section is active, displaying 'Run verification checks' with the instruction 'Verify that permissions and configuration are correct.' Below this, there are three buttons: 'Recheck' (highlighted with a green border), 'Download Script', and 'Grant'. A table follows, showing the results of three checks: 'VM Restore', 'Worker', and 'Repository', all of which have passed. At the bottom of the page, there are three buttons: 'Previous', 'Next', and 'Cancel'.

Check	Result	Details
VM Restore	Passed	All the required permissions are granted.
Worker	Passed	All the required permissions are granted.
Repository	Passed	All the required permissions are granted.

Step 10. Specify Restore Reason

At the **Reason** step of the wizard, specify a reason for restoring the persistent disks. This information will be saved to the session history, and you will be able to reference it later.

The screenshot shows the 'Disk Restore' wizard in Veeam Backup for Google Cloud. The interface is divided into a top navigation bar, a left sidebar, a main content area, and a bottom footer.

- Top Bar:** Veeam Backup for Google Cloud logo on the left. Server time: Nov 20, 2023 9:39 PM. User: administrator (Portal Administrator). Configuration icon on the right.
- Left Sidebar:** A vertical list of steps: Instances, Restore Mode, Project, Region, Encryption, Disk Settings, Verification, Reason (highlighted in green), and Summary.
- Main Content Area:**
 - Title: Enter reason for this restore operation
 - Label: Restore reason:
 - Text input field containing: evaluating disk restore
- Bottom Footer:** Three buttons: Previous (disabled), Next (active), and Cancel (disabled).

Step 11. Finish Working with Wizard

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

The screenshot shows the 'Disk Restore' wizard in the 'Summary' step. The interface includes a top navigation bar with the Veeam logo, server time (Nov 20, 2023 9:39 PM), user (administrator), and configuration options. A left sidebar lists steps: Instances, Restore Mode, Project, Region, Encryption, Disk Settings, Verification, Reason, and Summary (highlighted). The main content area is titled 'Review configured settings' and contains several sections: 'Copy to Clipboard', 'Project' (Name: veeam-rnd-backup-2, ID: rnd-backup-2, Service account: veeam-1649186685-sa@rnd-backup-2.iam.gserviceaccount.com), 'General' (Restore mode: New location, Region: europe-west8 (Milan), Availability zone: europe-west8-b, Encryption: tv-g-global, Reason: evaluating disk restore), 'Restore list' (Instance: dr-current-vm, Disks: 2 disks, Exclusions: —), and 'Validation' (Permission checks: Passed, Recheck; Disk settings: Passed). At the bottom, there are 'Previous', 'Finish', and 'Cancel' buttons.

Performing File-Level Recovery

In case a disaster strikes, you can recover corrupted or missing files of a VM instance from a cloud-native snapshot or image-level backup. Veeam Backup for Google Cloud allows you to download the necessary files and folders to a local machine or to their original location using the [File-Level Restore browser](#).

IMPORTANT

Consider the following:

- File-level recovery is supported for FAT, FAT32, NTFS, ext2, ext3, ext4, XFS and Btrfs file systems only. However, attributes of files and folders stored in FAT and FAT32 file systems cannot be restored to the original location.
- Restore of NTFS links (hard links, junction points, symbolic links) to the original location is not supported.

To recover files and folders of a protected VM instance, do the following:

1. [Launch the File-Level Recovery wizard](#).
2. [Select a restore point](#).
3. [Choose a restore mode](#).
4. [Run configuration and permission checks](#).
5. [Specify a recovery reason](#).

6. [Finish working with the wizard – start a recovery session.](#)
7. [Choose files and folders to recover.](#)
8. [Stop the recovery session.](#)

IMPORTANT

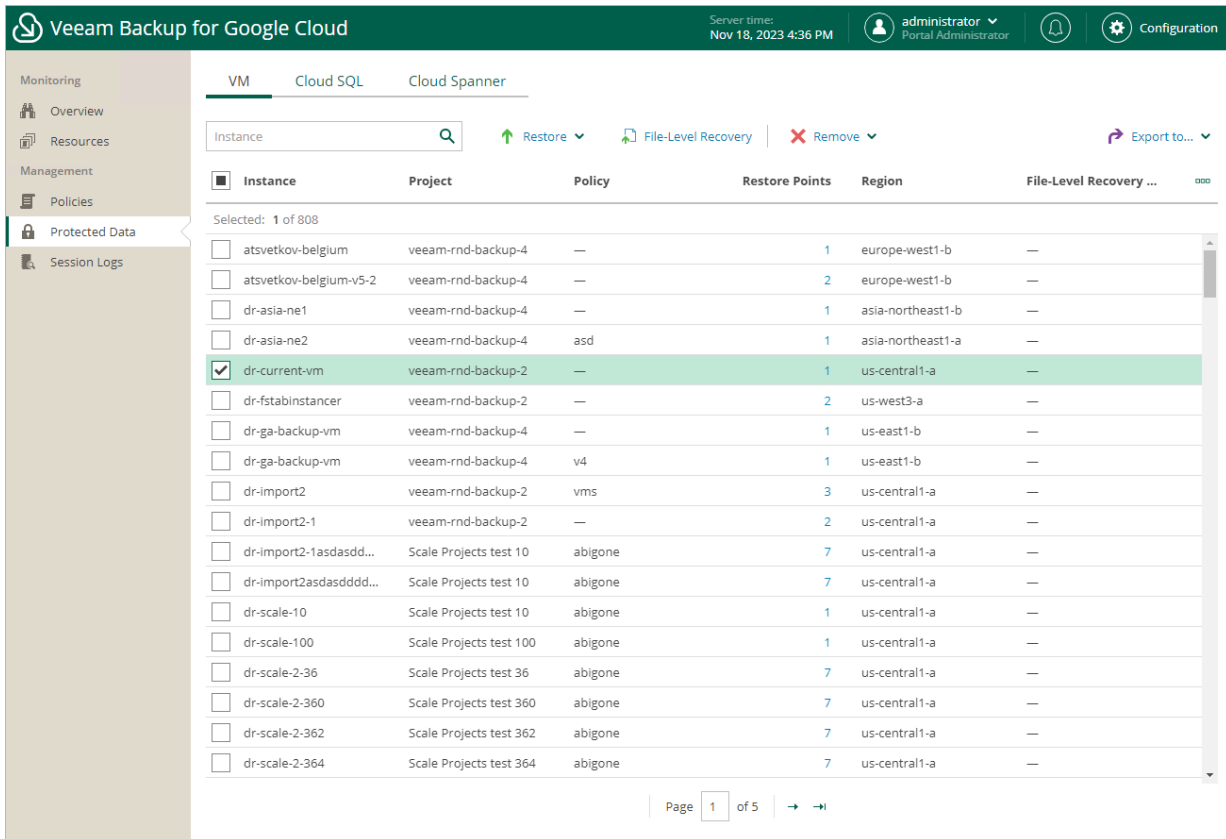
Before you start file-level recovery, check the following prerequisites:

- Make sure that network settings are configured for each region where worker instances will be deployed during the recovery process. For information on how to configure network settings, see [Adding Worker Configurations](#).
- Make sure that the machine where you plan to open the File-Level Restore browser is allowed to access the worker instances over the internet. To enable internet access for a worker instance, update the firewall rule specified in the instance network settings to add an inbound rule for HTTPS traffic on the port **443**. For information on how to update firewall rules, see [Google Cloud documentation](#).

Step 1. Launch File-Level Recovery Wizard

To launch the **File-Level Recovery** wizard, do the following:

1. Navigate to **Protected Data > VM**.
2. Select the VM instance whose files and folders you want to recover, and click **File-Level Recovery**.



Step 2. Select Restore Point

At the **Instances** step of the wizard, select a restore point that will be used to recover files and folders of the selected VM instance. By default, Veeam Backup for Google Cloud uses the most recent restore point. However, you can recover the items to an earlier state.

To select a restore point, do the following:

1. Select the VM instance.
2. Click **Choose Restore Point**.
3. In the **Select restore point** window, select the necessary restore point and click **Apply**.

To help you choose a restore point, Veeam Backup for Google Cloud provides the following information on each available restore point:

- **Creation Time** – the date when the restore point was created.
- **Destination** – the type of the restore point:
 - § *Snapshot* – a cloud-native snapshot created by a backup policy.
 - § *Manual Snapshot* – a cloud-native snapshot created manually.
 - § *Backup* – an image-level backup created by a backup policy.
- **State** – the result of the latest health check performed for the restore point.
- **Storage Class** – the storage class of a backup repository where the restore point is stored (applies only to image-level backups).
- **Project** – a project that manages the protected VM instance.
- **Region** – a region in which the protected VM instance resides.

- o **Retention** – a retention configured for the backup policy that created the restore point.

The screenshot shows the Veeam Backup for Google Cloud interface. The main window is titled 'File-Level Recovery'. On the left, there are navigation tabs: 'Instances', 'Restore mode', 'Reason', and 'Summary'. The 'Instances' tab is active, showing a 'Select VM instance' dialog with a search bar and a table of instances. One instance, 'dr-current...', is selected. A modal dialog titled 'Select restore point for VM instance: dr-current-vm' is open, displaying a table of restore points. The table has columns for 'Creation Time', 'Destination', 'State', 'Storage Class', and 'Project'. The restore points are listed in descending order of creation time, from 11/25/2023 3:00:29 PM to 12/06/2023 4:00:28 AM. At the bottom of the modal, there are 'Apply' and 'Cancel' buttons.

Creation Time ↑	Destination	State	Storage Class	Project
11/25/2023 3:00:29 PM	Snapshot	—	—	veeam-rnd-ba...
11/26/2023 3:00:24 PM	Snapshot	—	—	veeam-rnd-ba...
11/27/2023 3:00:27 PM	Snapshot	—	—	veeam-rnd-ba...
11/28/2023 3:00:24 PM	Snapshot	—	—	veeam-rnd-ba...
11/29/2023 3:00:29 PM	Snapshot	—	—	veeam-rnd-ba...
11/30/2023 3:00:23 PM	Snapshot	—	—	veeam-rnd-ba...
12/01/2023 3:00:26 PM	Snapshot	—	—	veeam-rnd-ba...
12/02/2023 3:00:24 PM	Snapshot	—	—	veeam-rnd-ba...
12/03/2023 3:00:33 PM	Snapshot	—	—	veeam-rnd-ba...
12/04/2023 3:00:25 PM	Snapshot	—	—	veeam-rnd-ba...
12/05/2023 3:00:28 PM	Snapshot	—	—	veeam-rnd-ba...
12/05/2023 6:00:39 PM	Snapshot	—	—	veeam-rnd-ba...
12/06/2023 4:00:28 AM	Snapshot	—	—	veeam-rnd-ba...

Step 3. Choose Restore Mode

At the **Restore Mode** step of the wizard, choose whether you want to download files and folders to a local machine or restore them to the original location. If you set the **Restore to original location** toggle to *On*, you must also specify a service account that has all the permissions required to perform the restore operation. For more information on the required permissions, see [Service Account Permissions](#).

For a service account to be displayed in the **Service account** drop-down list, it must be added to Veeam Backup for Google Cloud as described in section [Adding Service Accounts](#), and must be assigned the *VM Instances Restore* and *File-level Recovery to Original Location* operational roles as described in section [Adding Projects and Folders](#). If you have not added the necessary service account to Veeam Backup for Google Cloud beforehand, you can do it without closing the **VM Instance Restore** wizard. To add an account, click **Add** and complete the **Add Service Account** wizard.

IMPORTANT

Real-time protection (for example, Microsoft Defender antivirus) enabled on the target instance may significantly decrease the speed of the recovery process.

If you choose to perform restore to the original location, the target instance must meet the following requirements:

- The instance must be powered on.
- If the instance is a Linux-based VM, it must allow SSH access, and Veeam Backup for Google Cloud must have root access over SSH. To learn how to allow SSH access, see [Google Cloud documentation](#).
If the instance is a Windows-based VM, it must have Windows Remote Management (WinRM) configured. To learn how to configure WinRM, see [Microsoft documentation](#).
- The instance must be configured to allow the Cloud Pub/Sub API access. To learn how to allow Pub/Sub API access, see [Google Cloud documentation](#).
- The instance network must have the following firewall rule to allow access by IAP tunnel: IP address range 35.235.240.0/20, ports 22 and 5986. To learn how to configure firewall rules, see [Google Cloud documentation](#).

TIP

When Veeam Backup for Google Cloud performs restore to original location, it launches specific utilities on the target instance. If you plan to perform restore operations to the same instance in the future, you can select **Keep restore utilities on target instance after restore** check box to retain the utilities on the instance. This will allow Veeam Backup for Google Cloud to perform future restore operations faster.

File-Level Recovery

- Instances
- Restore mode**
- Verification
- Reason
- Summary

Choose restore mode

By default, files of the selected VM instance are downloaded to the local machine. Choose whether you want to restore them to the original location.

Restore to original location On

Service account: [veeam-1649186685-sa@rnd-backup-2.iam.gserviceaccount.com](#)

Keep restore utilities on target instance after restore ⓘ

Previous Next Cancel

Step 4. Run Configuration Checks

[This step applies only if you have set the **Restore to original location** toggle to *On* at the **Restore Mode** step of the wizard]

At the **Verification** step of the wizard, Veeam Backup for Google Cloud will verify whether restore settings are configured properly and the specified service account has all the necessary permissions required to perform recovery tasks for the target instance. For more information on the required permissions, see [Service Account Permissions](#).

To see the list of missing permissions that must be granted to the service account in order to perform an operation, click the link in the **Details** column. You can grant the missing permissions to the service account [using the Google Cloud console](#) or instruct Veeam Backup for Google Cloud to do it:

- To grant the missing permissions manually, click **Download Script**. Veeam Backup for Google Cloud will generate a gcloud script that you can run in the Google Cloud console to assign all the necessary permissions to the service account.

The account under which you run the script must have the permissions required both to get and set project IAM policies and to create custom IAM roles (for example, it can have the *iam.securityAdmin* and *iam.roleAdmin* roles assigned). To learn what permissions and roles are required to create custom roles in IAM, see [Google Cloud documentation](#).

- To let Veeam Backup for Google Cloud grant the missing permissions automatically, click **Grant** and then click **Sign in with Google** in the **Grant permissions** window. You will be redirected to the OAuth consent screen authorization page. Sign in using credentials of a Google account that will be used to grant the permissions.

The account under which you sign in to Google Cloud must have the permissions required both to get and set project IAM policies and to create custom IAM roles (for example, it can have the *iam.securityAdmin* and *iam.roleAdmin* roles assigned). To learn what permissions and roles are required to create service account, see [Google Cloud documentation](#).

NOTE

For Veeam Backup for Google Cloud to be able to authorize in Google Cloud, the OAuth consent screen must be configured as described in section [Registering Applications](#). Note that Veeam Backup for Google Cloud does not store in the configuration database the provided Google account credentials and access tokens received during authorization.

To make sure that the missing permissions have been successfully granted, click **Recheck**.

The screenshot shows the Veeam Backup for Google Cloud interface. At the top, the header includes the Veeam logo, the product name 'Veeam Backup for Google Cloud', the server time 'Nov 18, 2023 4:37 PM', and the user 'administrator Portal Administrator'. The main content area is titled 'File-Level Recovery' and features a left-hand navigation menu with options: 'Instances', 'Restore mode', 'Verification' (highlighted), 'Reason', and 'Summary'. The 'Verification' section is active, displaying 'Run verification checks' with the instruction 'Verify that permissions and configuration are correct.' Below this, there are three buttons: 'Recheck' (with a refresh icon), 'Download Script' (with a download icon), and 'Grant' (with a plus icon). A table below shows the results of the verification checks:

Check	Result	Details
Restore instance dr-current-vm	Passed	—
Worker	Passed	All the required permissions are granted.
Worker for File-Level Recovery to Origin...	Passed	All the required permissions are granted.
VM File-Level Recovery to Original Locat...	Passed	All the required permissions are granted.

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

Step 5. Specify Recovery Reason

At the **Reason** step of the wizard, 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 displays the Veeam Backup for Google Cloud interface. At the top, the header includes the Veeam logo, the product name 'Veeam Backup for Google Cloud', the server time 'Nov 18, 2023 4:38 PM', the user 'administrator Portal Administrator', and a 'Configuration' link. Below the header, the main area is titled 'File-Level Recovery'. On the left, a navigation pane lists 'Instances', 'Restore mode', 'Verification', 'Reason' (which is highlighted in green), and 'Summary'. The main content area is titled 'Enter reason for this restore operation' and contains a text input field with the text 'file-level recovery evaluation'. At the bottom of the interface, there are three buttons: 'Previous', 'Next', and 'Cancel'.

Step 6. 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 Google Cloud 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 Google Cloud will deploy a worker instance and attach persistent disks of the processed VM instance to it.

TIP

If you accidentally close the **FLR Running Sessions** window, navigate to **Protected Data** 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 state of the recovery session. In the **URL** column of the window, Veeam Backup for Google Cloud will display a link to the File-Level Restore browser. You can use the link in either of the following ways:

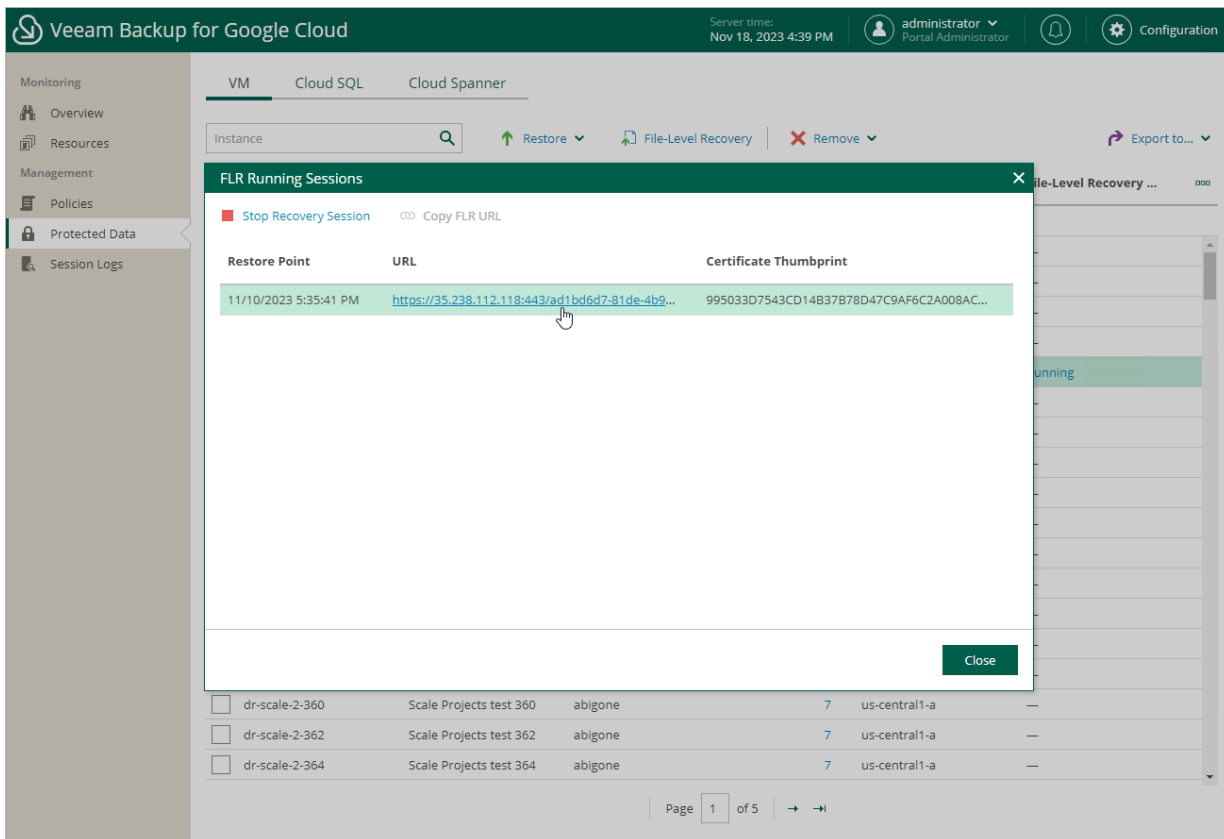
- Click the link to open the File-Level Restore browser on your local machine while the recovery session is running.
- Copy the link, close the **FLR Running Sessions** window and open the File-Level Restore browser on another machine.

IMPORTANT

When you click **Copy FLR URL**, Veeam Backup for Google Cloud copies the following information to the clipboard:

- A link to the File-Level Restore browser that includes an IP address of the worker instance hosting the browser and authentication information used to access the browser.
- A thumbprint of a TLS certificate that is installed on the worker instance hosting the File-Level Restore 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 Restore browser matches the provided certificate thumbprint.



The screenshot shows the Veeam Backup for Google Cloud interface. A dialog box titled "FLR Running Sessions" is open, displaying a table of running sessions. The table has three columns: "Restore Point", "URL", and "Certificate Thumbprint". The first row is highlighted, showing a restore point of "11/10/2023 5:35:41 PM", a URL starting with "https://35.238.112.118:443/ad1bd6d7-81de-4b9...", and a certificate thumbprint of "995033D7543CD14B37B78D47C9AF6C2A008AC...". The dialog box also includes a "Close" button at the bottom right and a "Copy FLR URL" button at the top right. The background interface shows a navigation menu on the left and a main area with a table of instances.

Restore Point	URL	Certificate Thumbprint
11/10/2023 5:35:41 PM	https://35.238.112.118:443/ad1bd6d7-81de-4b9...	995033D7543CD14B37B78D47C9AF6C2A008AC...

Step 7. Choose Items to Recover

In the File-Level Restore browser, you can find and recover items (files and folders) of the selected VM instance:

1. In the File-Level Restore browser, navigate to a folder that contains the necessary files.
2. In the working area, select check boxes next to the files and click **Add to Restore List**.
3. Repeat steps 1–2 for all other folders whose files you want to recover.
4. Switch to the **Restore List** tab.
5. On the **Restore List** tab, review the list of items to recover, select check boxes next to the items and do the following:
 - o To save all the recovered items as a single .ZIP archive to the default download directory on a machine from which you access the browser, click **Download**.
 - o To recover the items to the original location, click **Restore**.

NOTE

When recovering items to the original location, Veeam Backup for Google Cloud will be able to display the directory structure only in case the disks of the source VM were mounted either using drive letters (for Windows-based VMs) or using UUIDs/labels with mount records stored in the `/etc/fstab` file (for Linux-based VMs). If Veeam Backup for Google Cloud fails to display the structure correctly, you will be prompted to manually provide a path to the items you want to recover.

Restore List: tvg-win-regress-dc-2019

Restore Status: All

Restore | Download | Stop | Remove

	Location	Type	Size	Last Modified	Restore Point	Restore Status	
<input checked="" type="checkbox"/>	Recovery	C:		11/27/2023 4:48:47 PM	12/28/2023 11:05:12 AM	—	
<input checked="" type="checkbox"/>	NTUSER.DAT(edd21381-...	C:\Users\veeam_restore_user	.blf	64.0 kB	11/27/2023 5:55:14 PM	12/28/2023 11:05:12 AM	—
<input checked="" type="checkbox"/>	NTUSER.DAT(edd21381-...	C:\Users\veeam_restore_user	.regtrans-ms	512.0 kB	11/27/2023 5:55:08 PM	12/28/2023 11:05:12 AM	—
<input checked="" type="checkbox"/>	NTUSER.DAT(edd21381-...	C:\Users\veeam_restore_user	.regtrans-ms	512.0 kB	11/27/2023 5:55:08 PM	12/28/2023 11:05:12 AM	—
<input checked="" type="checkbox"/>	New Text Document	D:\New folder	.txt	1.3 kB	11/30/2023 1:20:15 PM	12/28/2023 11:05:12 AM	—

Session Log

Status: All

Action	Status	Start Time	End Time	Duration
Select a single item to view sessions details				

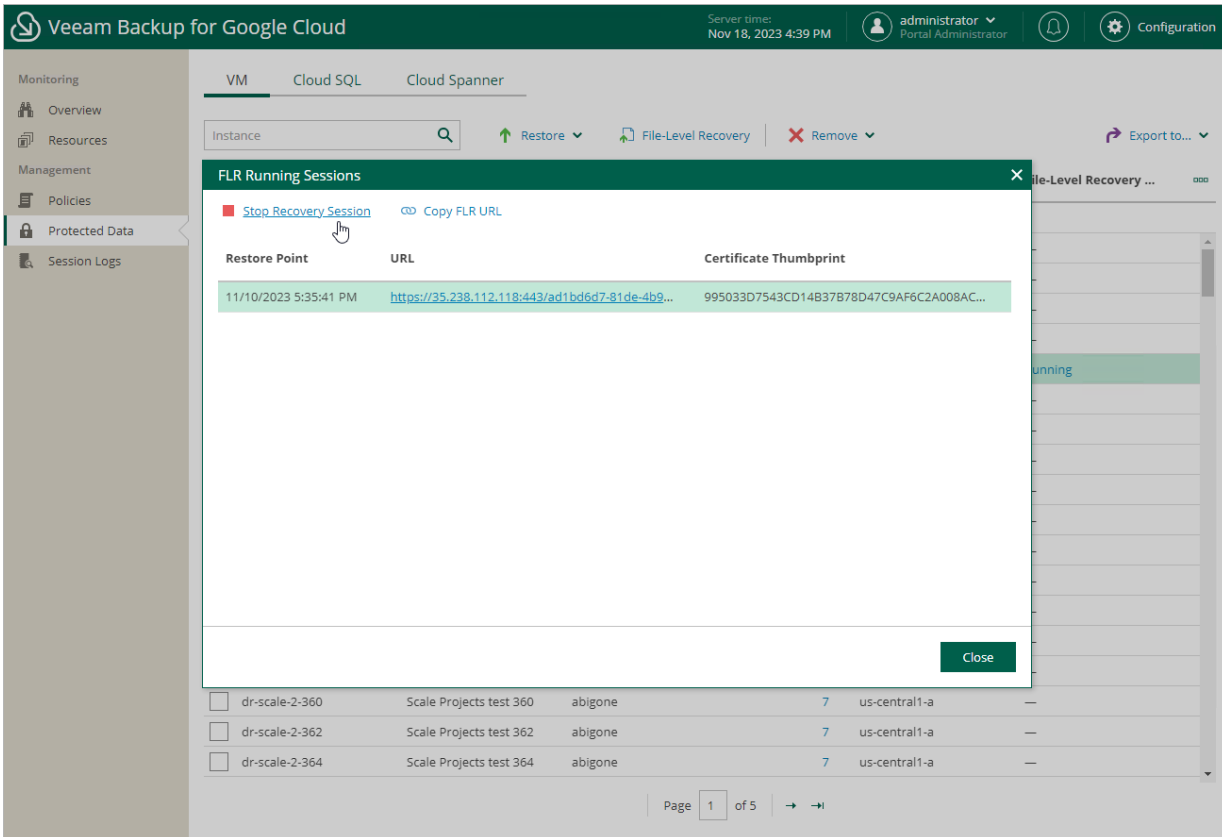
Step 8. Stop Recovery Session

After you finish working with the File-Level Restore browser, it is recommended that you stop the running recovery session so that Veeam Backup for Google Cloud can detach persistent disks of the processed VM instance from the deployed worker instance and remove the worker instance from Google Cloud.

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 Restore browser for 30 minutes, Veeam Backup for Google Cloud will stop the recovery session automatically.

TIP

If you accidentally close the **FLR Running Sessions** window, navigate to **Protected Data** and click the link in the **File-Level Recovery URL** column to open the window again.



SQL Restore

The actions that you can perform with restore points of Cloud SQL instances depend on whether you access the restore points using the Veeam Backup & Replication console or the Veeam Backup for Google Cloud Web UI.

SQL Restore Using Console

Veeam Backup & Replication offers the following restore operations:

- [Instance restore](#) – start an entire Cloud SQL instance from a restore point.
- [Database restore](#) – restore specific databases of a Cloud SQL instance from an image-level backup.

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

NOTE

You can use restore points stored in standard repositories to perform all the listed recovery operations, while restore points stored in archive repositories can only be used to perform restore of Cloud SQL instances to the original or to a new location.

Performing SQL Instance Restore

In case a disaster strikes, you can restore an entire Cloud SQL instance from a cloud-native snapshot or an image-level backup. Veeam Backup & Replication allows you to restore one or more Cloud SQL instances at a time, to the original location or to a new location. To learn how Cloud SQL restore works, see [Performing Instance Restore](#).

To restore a Cloud SQL instance, do the following:

1. [Launch the Restore to Google Cloud SQL wizard](#).
2. [Select a restore point](#).
3. [Choose a restore mode](#).
4. [Select a project, region and an availability zone](#).
5. [Specify instance type and name](#).
6. [Configure network settings](#).
7. [Configure security settings](#).
8. [Enable flag assignment](#).
9. [Specify a restore reason](#).
10. [Finish working with the wizard](#).

Step 1. Launch Restore to Google Cloud SQL Wizard

To launch the **Restore to Google Cloud SQL** 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.

NOTE

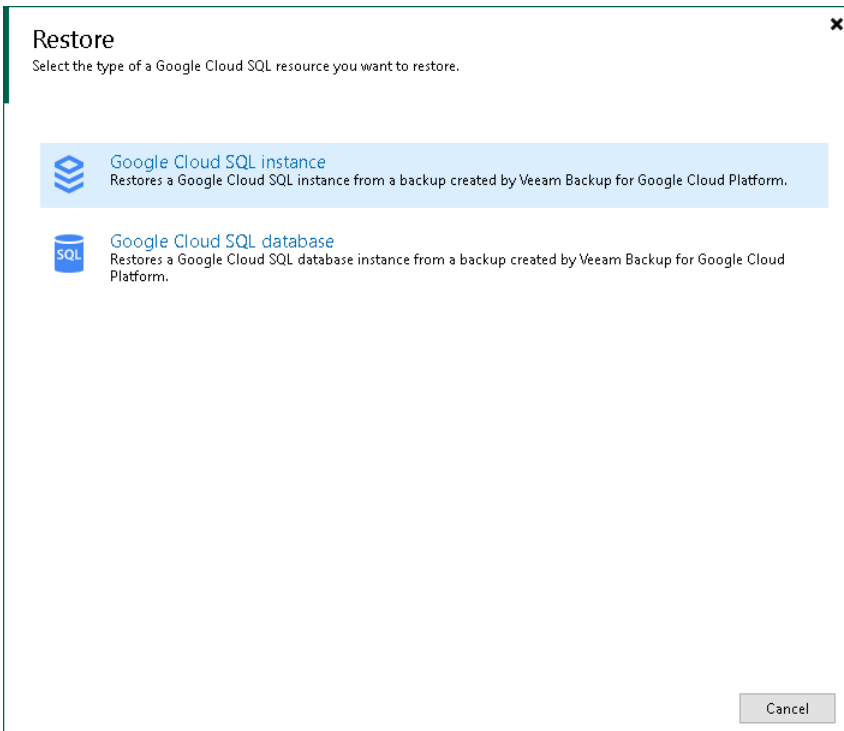
Note that restore of Cloud SQL instances to the original location is supported only from image-level backups.

3. In the working area, expand the backup policy that protects a Cloud SQL instance you want to restore and select the necessary instance. Then, click **Google Cloud SQL** on the ribbon and select **Google Cloud SQL instance** in the **Restore** window.

Alternatively, you can right-click the instance and select **Restore to Google Cloud SQL**. In the **Restore** window, select **Google Cloud SQL instance**.

TIP

You can also launch the **Restore to Google Cloud SQL** wizard from the **Home** tab. To do that, click **Restore** and select **GCP**. Then, in the **Restore** window, select **Google Cloud SQL** and, depending on whether you want to restore from a backup or a snapshot, select either **Restore from Google Cloud SQL snapshot** or **Restore from Veeam backup**.



Step 2. Select Restore Point

At the **SQL instance** step of the wizard, choose a restore point that will be used to restore the selected Cloud SQL instance. By default, Veeam Backup & Replication uses the most recent valid restore points. However, you can restore the instance data to an earlier state.

To select a restore point, do the following:

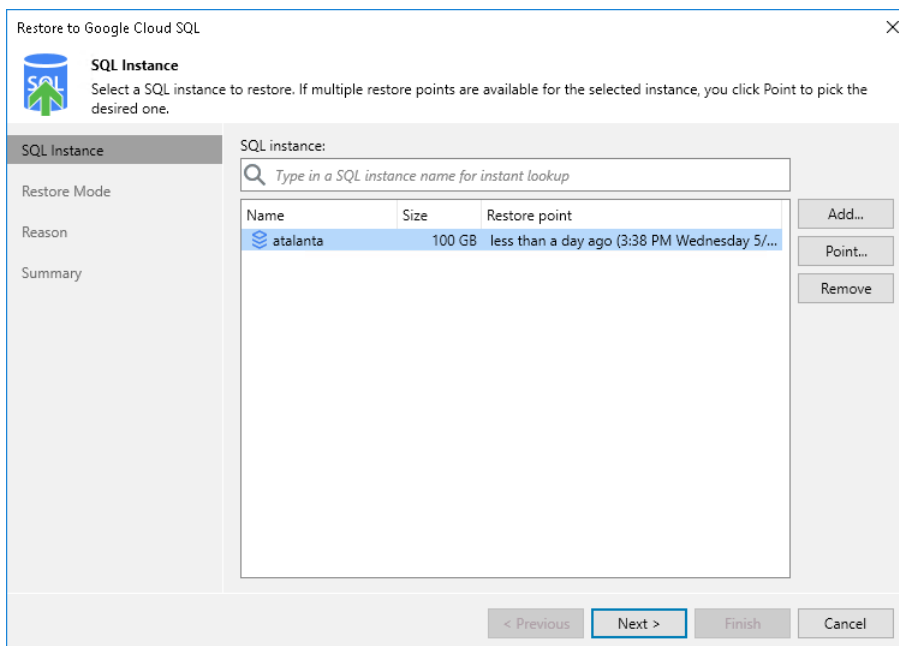
1. In the **Select SQL instance** list, select the Cloud SQL instance and click **Point**.
2. In the **Restore Points** window, expand the backup policy that protects the Cloud SQL instance, select the necessary restore point and click **OK**.

To help you choose a restore point, Veeam Backup & Replication provides the following information on each available restore point:

- o **Job** – the name of the backup policy that created the restore point and the date when the restore point was created.
- o **Type** – the type of the restore point.
- o **Location** – the 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 Cloud SQL instances to restore and select 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 Cloud SQL instance to the original or to a new location.

NOTE

Due to [technical limitations in Google Cloud](#), Veeam Backup & Replication does not support restore to the original location if the source Cloud SQL instance is still present in the location, if it has been recently deleted (less than a week ago), or if its name is reserved.

2. Click **Pick account to use** to select a service account whose permissions will be used to perform the restore operation. For more information on the required permissions for service accounts, see [Service Account Permissions](#).

For a service account to be displayed in the list of available accounts, it must be added to the backup appliance as described in section [Adding Service Accounts](#), and must be assigned the *Cloud SQL Instances Restore* operational role as described in section [Adding Projects and Folders](#).

NOTE

By default, to perform the restore operation, Veeam Backup & Replication uses permissions of the service account that has been used to protect the source Cloud SQL instance.

The screenshot shows a wizard window titled "Restore to Google Cloud SQL" with a close button (X) in the top right corner. The window contains a sidebar on the left with a list of steps: "SQL Instance", "Restore Mode" (highlighted), "Data Center", "SQL Instance", "Network", "Security", "Flags", "Reason", and "Summary". The main area of the wizard is titled "Restore Mode" and includes a sub-header "Specify whether selected SQL instances should be restored back to the original location, or to a new location or with different settings." Below this, there are two radio button options: "Restore to the original location" (unselected) and "Restore to a new location, or with different settings" (selected). The selected option has a description: "Customize the restored SQL instance location, and change its settings. The wizard will automatically populate all controls with the original SQL instance settings as the defaults." Below the description is a blue link that says "Pick account to use". At the bottom of the wizard, there are four buttons: "< Previous", "Next >" (highlighted with a blue border), "Finish", and "Cancel".

Step 4. Select Project, Region and Availability Zone

[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 a project that will be used to manage the restored Cloud SQL instance, and specify a region and an availability zone where the restored instance will operate.

For a project to be displayed in the list of available projects, it must be created in Google Cloud as described in [Google Cloud documentation](#).

TIP

To configure the restored Cloud SQL instance for high availability, select the **Multiple zones (survives datacenter outage)** option, and choose a primary and secondary zone where the restored instance will be located within the selected region. The high availability configuration allows you to reduce downtime when a zone or the instance becomes unavailable. For more information on high availability in Google Cloud, see [Google Cloud documentation](#).

Note that this option is available only for restore points created for Cloud SQL instances with high availability enabled.

The screenshot shows the 'Data Center' step of the 'Restore to Google Cloud SQL' wizard. The window title is 'Restore to Google Cloud SQL'. The left sidebar contains a list of steps: SQL Instance, Restore Mode, Data Center (highlighted), SQL Instance, Network, Security, Flags, Reason, and Summary. The main content area is titled 'Data Center' and includes the instruction 'Specify a data center and availability settings for the restored SQL instance.' The form contains the following fields and options:

- Project:** A dropdown menu with 'RnD Backup 4' selected.
- Data center:** A dropdown menu with 'europe-north1 (Finland)' selected.
- Availability settings:** Two radio button options:
 - Single zone (survives host outage only)
 - Multiple zones (survives datacenter outage)
- Availability zone:** A dropdown menu with 'europe-north1-c' selected (this field is only visible when the 'Single zone' option is selected).
- Primary zone:** A dropdown menu with 'europe-north1-c' selected (this field is only visible when the 'Multiple zones' option is selected).
- Secondary zone:** An empty dropdown menu (this field is only visible when the 'Multiple zones' option is selected).

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

Step 5. Specify Instance Type and 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 **SQL Instance** step of the wizard, specify a new name for the restored Cloud SQL instance.

TIP

You can specify a single prefix or suffix and add it to the names of multiple Cloud SQL 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**.

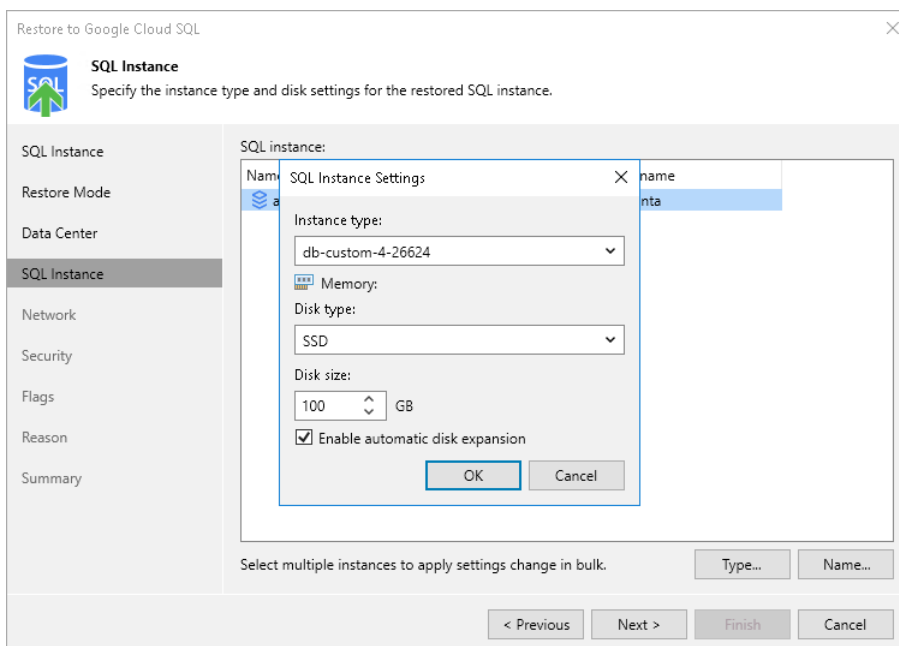
You can also configure the following settings:

- You can specify a new machine type for the restored Cloud SQL instance. To do that, select the instance and click **Type**. Then, select the necessary type in the **SQL Instance Settings** window.

To learn how to choose the machine type when creating a Cloud SQL instance in Google Cloud, see [Google Cloud documentation](#).

- You can choose a new disk storage type or increase (either manually or automatically) storage capacity for the restored Cloud SQL instance. To do that, select the instance and click **Type**. Then, use the options in the **Memory** section of the **SQL Instance Settings** window. Note, however, that the amount of storage capacity allocated to an instance affects its cost.

To learn how to configure storage settings when creating a Cloud SQL instance in Google Cloud, see [Google Cloud documentation](#).



Step 6. 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 settings for the restored Cloud SQL instance. To do that, select the instance and do the following:

- If you want to connect the restored Cloud SQL instance to a VPC network with a private IP address, click **Access**. In the **Access Settings** window, select the **Assign a private IP address from following Virtual Private Network** check box, choose a VPC network to which the instance will be connected, and click **OK**.

For a VPC network to be displayed in the lists of available networks, it must be created in the Google Cloud for the region specified at [step 4](#) of the wizard, as described in [Google Cloud documentation](#).

IMPORTANT

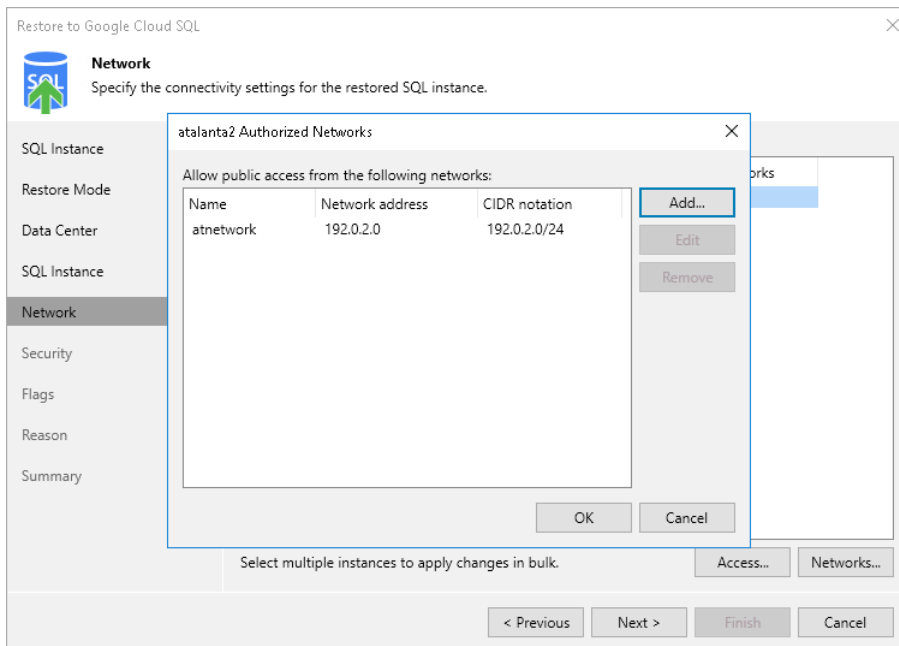
The specified VPC network must have Private Service Connect configured. For more information, see [Google Cloud documentation](#).

- If you want to assign a public IPv4 address to the restored Cloud SQL instance and to accept connections to it from specific IP address ranges, click **Access**. In the **Access Settings** window, select the **Assign a public IP address** check box and click **OK**.

Then, click **Network**. In the **Authorized Networks** window, add the allowed IP address ranges and click **OK**. The IP address ranges must be specified in the CIDR notation (for example, `12.23.34.0/24`).

TIP

To let all IP addresses access the restored Cloud SQL instance, you can enter `0.0.0.0/0`. However, note that allowing access from all IP addresses is unsafe and thus not recommended in production environments.



Step 7. Configure Security 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 **Security** step of the wizard, you can configure specific security settings for the restored Cloud SQL instance. To do that, select the instance and do the following:

- If you want to connect to the restored Cloud SQL instance using TLS only, click **Security** and select the **Allow only secure connections (TLS)** option in the **Security Settings** window.

NOTE

Since TLS connections use digital certificates to provide encrypted access, make sure that you have obtained a Certificate Authority (CA) certificate, a client public key certificate, and a client private key – before you connect to the restored instance using TLS. For more information, see [Google Cloud documentation](#).

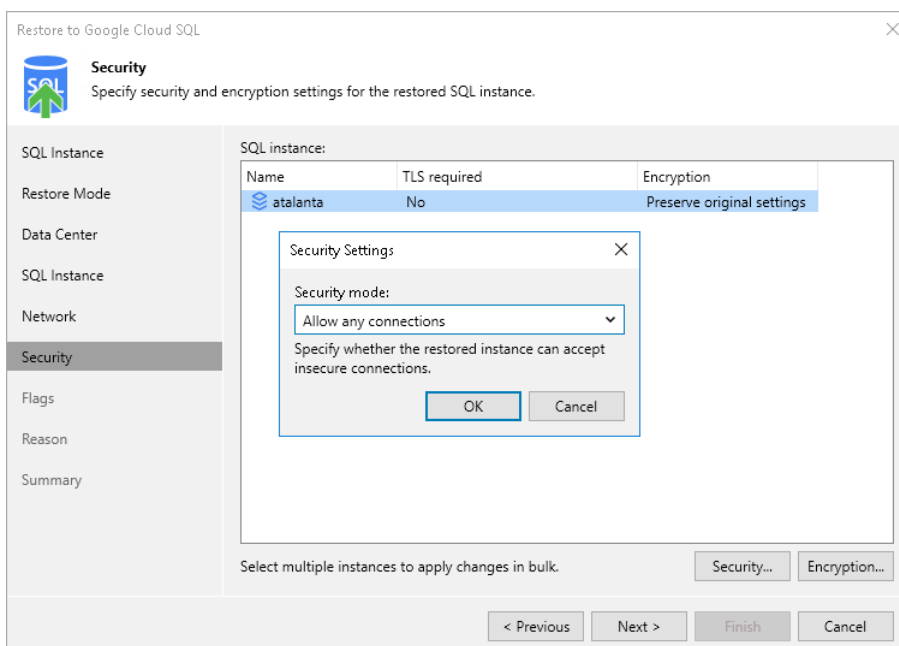
If you do not want to connect to the restored Cloud SQL instance using TLS, select the **Allow any connections** option.

- If you want to change the encryption settings of the restored Cloud SQL instance, click **Encryption** and do the following in the **Disk Encryption** window:
 - If you do not want to encrypt the restored data or want to apply the existing encryption scheme, select the **Preserve the original encryption settings** option.
 - If you want to encrypt the restored data with a Google Cloud KMS CMEK, select the **Use the following encryption key** option. Then, select the necessary CMEK from the drop-down list.

For a CMEK to be displayed in the list of available encryption keys, it must be stored in the region selected at [step 4](#) of the wizard.

NOTE

The **Preserve the original encryption settings** option is disabled if the CMEK that was used to encrypt data of the source instance is not available in the region to which the Cloud SQL instance will be restored.




Step 8. Enable Flag Assignment

[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 **Flags** step of the wizard, you can modify flags set on databases of the restored Cloud SQL instance. To do that, select the instance and do the following:

1. Click **Flags**.
2. In the **Flags** window, choose whether you want flags of the restored databases to have the same value as the source databases or new modified values.

If you want to set a new value for a database flag, select the flag and click **Edit**. To save changes made to the flag settings, click **OK**.



Step 9. Specify Restore Reason

At the **Reason** step of the wizard, specify a reason for restoring the Cloud SQL instance. The information you provide will be saved in the session history and you can reference it later.

Restore to Google Cloud SQL

Reason
Type in the reason for performing this restore operation. This information will be logged in the restore sessions history for later reference.

SQL Instance
Restore Mode
Data Center
SQL Instance
Network
Security
Flags
Reason
Summary

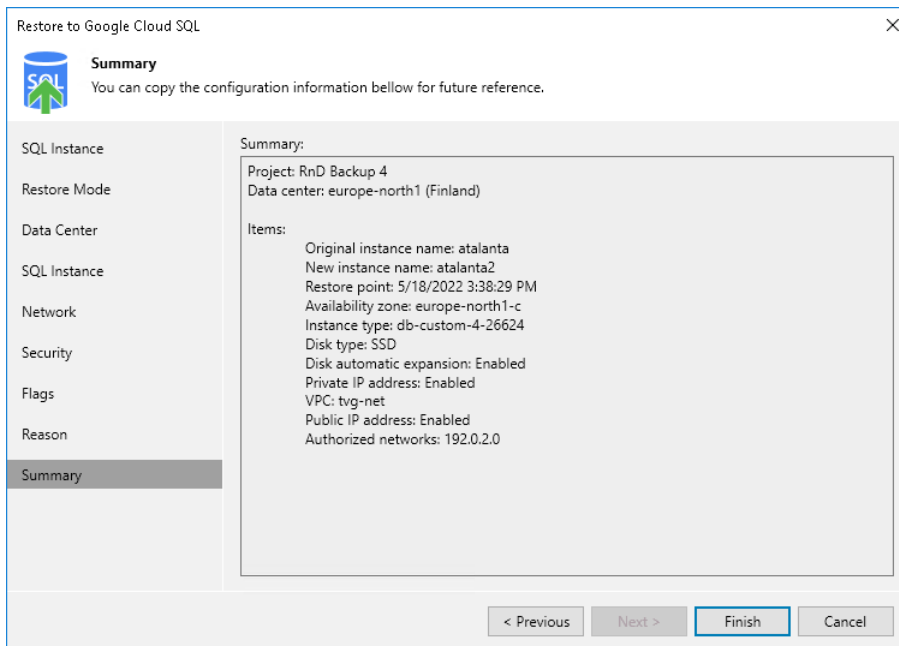
Restore reason:
Restore a production database.

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**.



Performing Database Restore

In case a disaster strikes, you can restore corrupted databases of Cloud SQL instance from an image-level backup. Veeam Backup & Replication allows you to restore databases to the original location or to a new location.

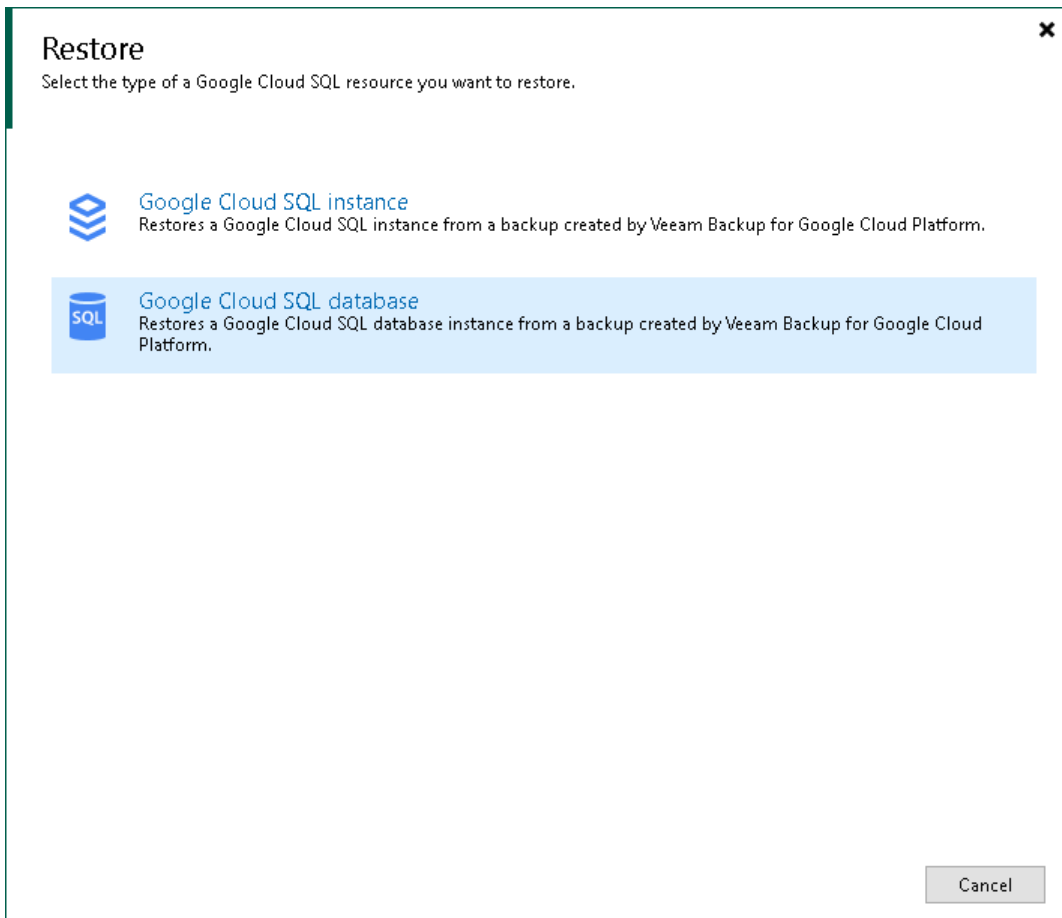
To restore a database, do the following:

1. In the Veeam Backup & Replication console, open the **Home** view.
2. Navigate to **Backups**.
3. Expand the backup policy that protects a Cloud SQL instance whose database you want to restore, select the necessary instance and click **Google Cloud SQL** on the ribbon.

Alternatively, you can right-click the Cloud SQL instance and select **Restore to Google SQL**.

4. In the **Restore** window, select **Google Cloud SQL database**.

Veeam Backup & Replication will open the **Data Restore** wizard in a web browser. Complete the wizard as described in section [Performing Database Restore](#).



SQL Restore Using Web UI

Veeam Backup for Google Cloud offers the following restore operations:

- [Instance restore](#) – start an entire Cloud SQL instance from a restore point.
- [Database restore](#) – restore specific databases of a Cloud SQL instance.

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

Performing SQL Instance Restore

In case a disaster strikes, you can restore an entire Cloud SQL instance from a cloud-native snapshot or image-level backup. Veeam Backup for Google Cloud allows you to restore one or more Cloud SQL instances at a time, to the original location or to a new location.

NOTE

Veeam Backup for Google Cloud does not support restore to the original location if the source Cloud SQL instance is still present in the location or if its name is reserved. However, note that if you delete an instance from Google Cloud, all its cloud-native snapshots will be deleted as well due to [technical limitations in Google Cloud](#).

To restore a protected Cloud SQL instance, do the following:

1. [Launch the Cloud SQL Instance Restore wizard](#).
2. [Select a restore point](#).
3. [Choose a restore mode](#).
4. [Select a service account](#).
5. [Select a project](#).
6. [Select a region and an availability zone](#).
7. [Specify a new name and machine type for the instance](#).
8. [Configure network settings](#).
9. [Configure security settings](#).
10. [Enable flag assignment](#).
11. [Run configuration and permission checks](#).
12. [Specify a restore reason](#).
13. [Finish working with the wizard](#).

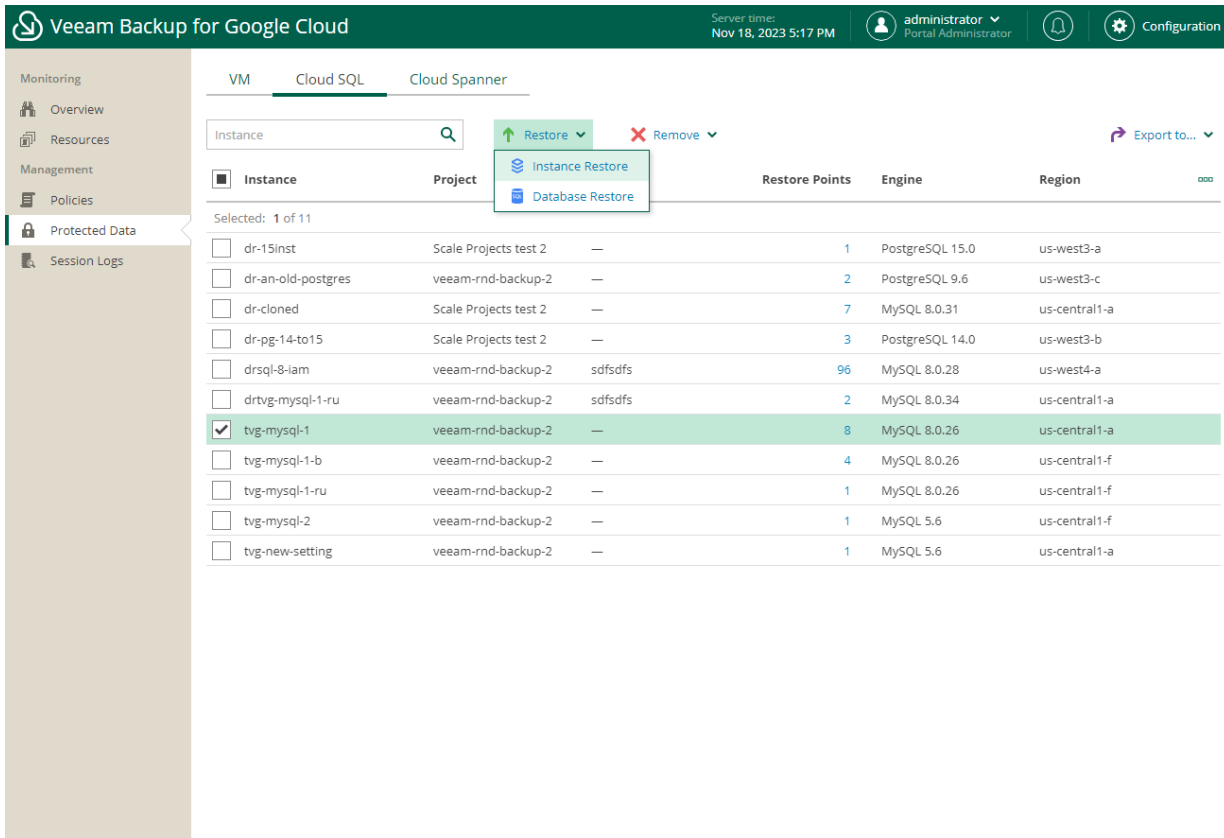
IMPORTANT

Before you start Cloud SQL instance restore, make sure that network settings are configured for each region where worker instances will be deployed during the restore process. For information on how to configure network settings, see [Adding Worker Configurations](#).

Step 1. Launch Cloud SQL Instance Restore Wizard

To launch the **Cloud SQL Instance Restore** wizard, do the following:

1. Navigate to **Protected Data > Cloud SQL**.
2. Select the Cloud SQL instance that you want to restore, and click **Restore > Instance Restore**.



The screenshot shows the Veeam Backup for Google Cloud interface. The top navigation bar includes the Veeam logo, the product name, the server time (Nov 18, 2023 5:17 PM), the user (administrator), and a Configuration icon. The left sidebar shows the navigation menu with 'Protected Data' selected. The main content area is titled 'Cloud SQL' and displays a table of instances. The 'Instance Restore' option is highlighted in the 'Restore' dropdown menu.

Instance	Project	Restore Points	Engine	Region
<input type="checkbox"/> dr-15inst	Scale Projects test 2	1	PostgreSQL 15.0	us-west3-a
<input type="checkbox"/> dr-an-old-postgres	veeam-rnd-backup-2	2	PostgreSQL 9.6	us-west3-c
<input type="checkbox"/> dr-cloned	Scale Projects test 2	7	MySQL 8.0.31	us-central1-a
<input type="checkbox"/> dr-pg-14-to15	Scale Projects test 2	3	PostgreSQL 14.0	us-west3-b
<input type="checkbox"/> drsql-8-iam	veeam-rnd-backup-2	96	MySQL 8.0.28	us-west4-a
<input type="checkbox"/> drtv-mysql-1-ru	veeam-rnd-backup-2	2	MySQL 8.0.34	us-central1-a
<input checked="" type="checkbox"/> tvg-mysql-1	veeam-rnd-backup-2	8	MySQL 8.0.26	us-central1-a
<input type="checkbox"/> tvg-mysql-1-b	veeam-rnd-backup-2	4	MySQL 8.0.26	us-central1-f
<input type="checkbox"/> tvg-mysql-1-ru	veeam-rnd-backup-2	1	MySQL 8.0.26	us-central1-f
<input type="checkbox"/> tvg-mysql-2	veeam-rnd-backup-2	1	MySQL 5.6	us-central1-f
<input type="checkbox"/> tvg-new-setting	veeam-rnd-backup-2	1	MySQL 5.6	us-central1-a

Step 2. Select Restore Point

At the **Instances** step of the wizard, select a restore point that will be used to restore the selected Cloud SQL instance. By default, Veeam Backup for Google Cloud 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. Select the Cloud SQL 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 Google Cloud provides the following information on each available restore point:

- **Creation Time** – the date when the restore point was created.
- **Destination** – the type of the restore point:
 - § *Snapshot* – a cloud-native snapshot 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 result of the latest health check performed for the restore point.
- **Storage Class** – the storage class of a backup repository where the restore point is stored (applies only to image-level backups).
- **Policy** – a backup policy that created the restore point.
- **Region** – a region in which the protected Cloud SQL instance resides.
- **Project** – a project that manages the protected Cloud SQL instance.

- o **Retention** – a retention configured for the backup policy that created the restore point.

The screenshot shows the Veeam Backup for Google Cloud interface. The top navigation bar includes the Veeam logo, the product name 'Veeam Backup for Google Cloud', the server time 'Nov 18, 2023 5:19 PM', and user information 'administrator Portal Administrator'. The main content area is titled 'Cloud SQL Instance Restore'. On the left, there is a sidebar with navigation options: 'Instances', 'Restore Mode', 'Service Account', 'Verification', 'Reason', and 'Summary'. The main panel is divided into two sections: 'Choose Cloud SQL instances to restore' and 'Choose restore point'. The 'Choose Cloud SQL instances to restore' section has a search bar and a table with one entry: 'tvq-mysql-1' with a restore point of '10/31/2023 5:4...'. The 'Choose restore point' section is a modal dialog with a table of restore points. The table has columns: 'Creation...', 'Destinat...', 'State', 'Storage ...', 'Policy', 'Region', and 'Project'. The row for '09/12/20...' is highlighted in green. At the bottom of the dialog are 'Apply' and 'Cancel' buttons.

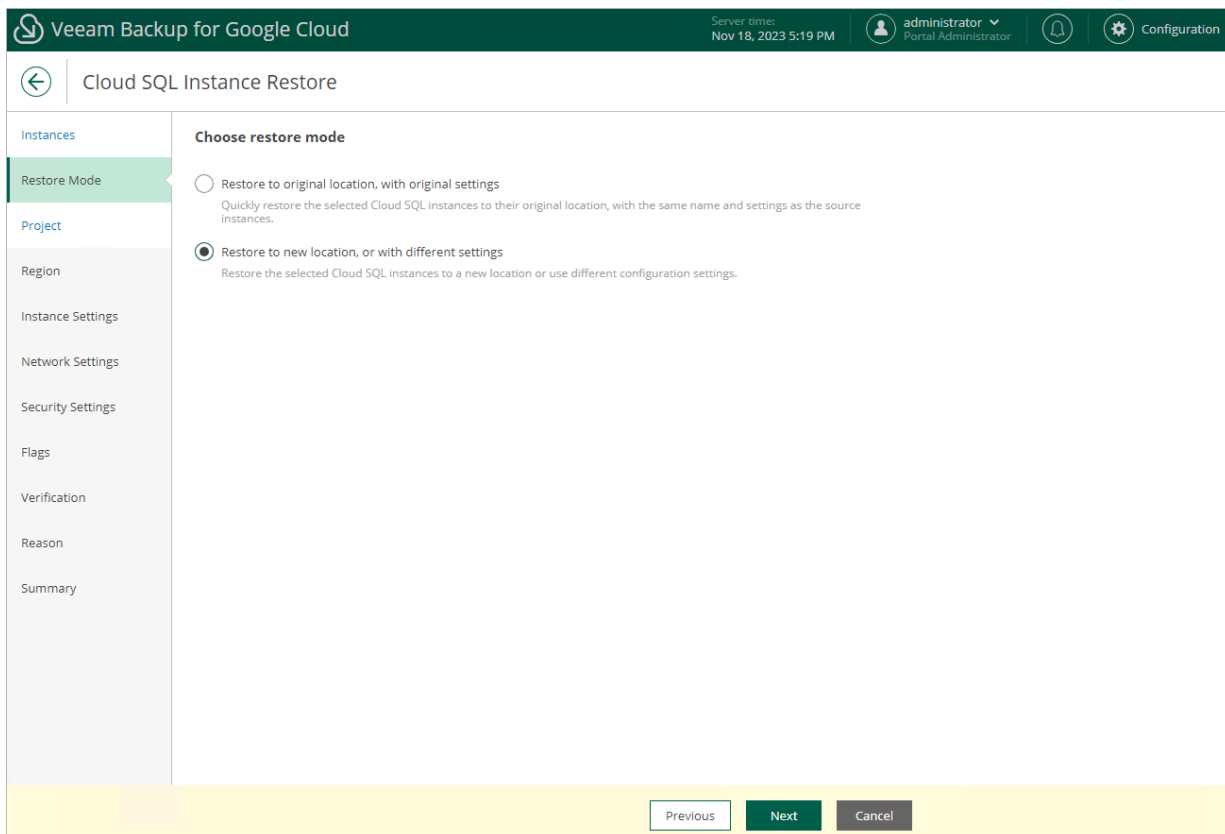
Creation...	Destinat...	State	Storage ...	Policy	Region	Project
10/31/20...	Snapshot	—	—	—	us-centra...	veeam-rn...
10/02/20...	Snapshot	—	—	—	us-centra...	veeam-rn...
09/12/20...	Snapshot	—	—	—	us-centra...	veeam-rn...
09/12/20...	Snapshot	—	—	—	us-centra...	veeam-rn...
09/11/20...	Snapshot	—	—	—	us-centra...	veeam-rn...
08/15/20...	Snapshot	—	—	—	us-centra...	veeam-rn...
08/14/20...	Snapshot	—	—	—	us-centra...	veeam-rn...

Step 3. Choose Restore Mode

At the **Restore Mode** step of the wizard, choose whether you want to restore the selected Cloud SQL instance to the original or to a custom location.

IMPORTANT

Restore to the original location is supported only using restore points of the *Backup* and *Archive* types. If you select a restore point of the *Snapshot* or *Manual Snapshot* type at [step 2](#) of the wizard, you will be able to select the **Restore to original** option and proceed with the wizard but only up to the **Verification** step — at this step, the verification check will notify you that the restore settings have not been configured properly. As a result, Veeam Backup for Google Cloud will not be able to perform the operation



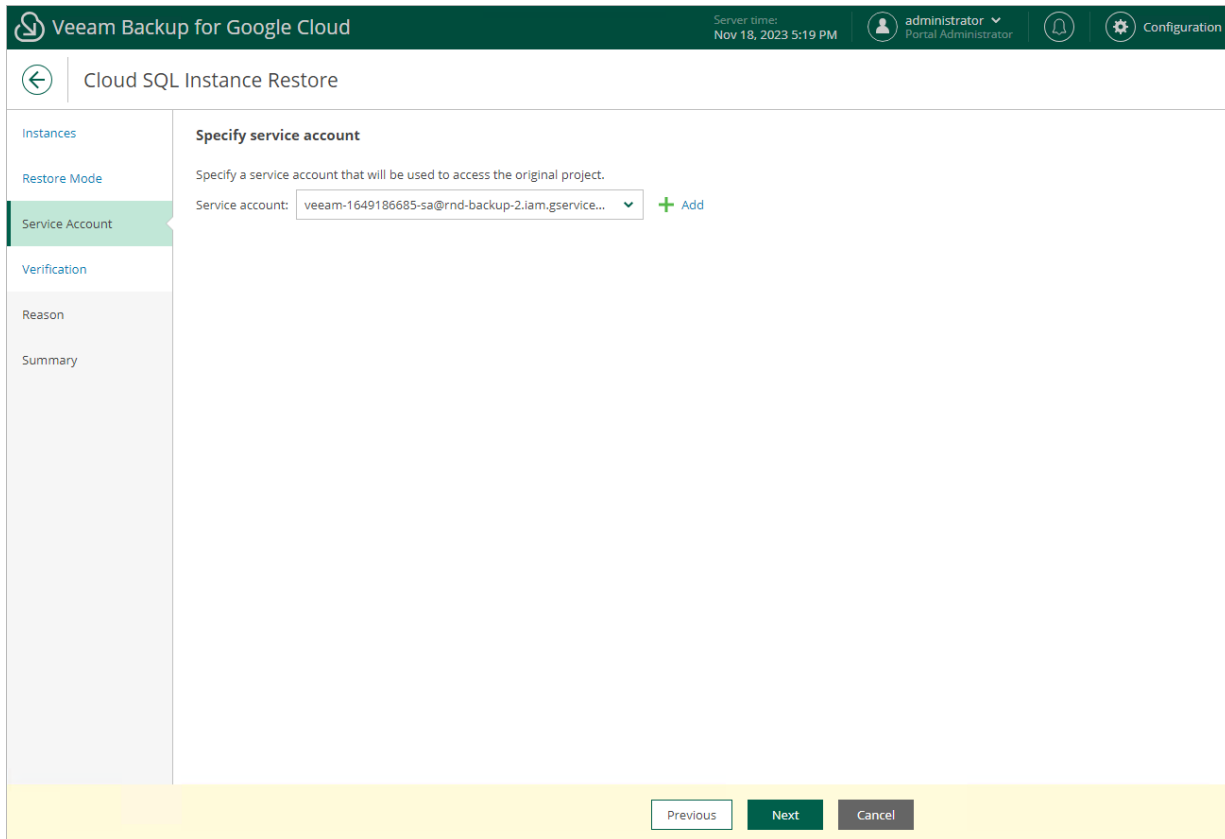
Step 4. Select Service Account

[This step applies only if you have selected the **Restore to original location, with original settings** option at the **Restore Mode** step of the wizard]

At the **Service Account** step of the wizard, select a service account whose permissions will be used to perform the restore operation. For more information on the required permissions, see [Service Account Permissions](#).

For a service account to be displayed in the **Service account** drop-down list, it must be added to Veeam Backup for Google Cloud for Google Cloud as described in section [Adding Service Accounts](#), and must be assigned the *Cloud SQL Instances Restore* operational role as described in section [Adding Projects and Folders](#).

If you have not added the necessary service account to Veeam Backup for Google Cloud beforehand, you can do it without closing the **Cloud SQL Instance Restore** wizard. To add an account, click **Add** and complete the **Add Service Account** wizard.



Step 5. Select Project

[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 **Project** step of the wizard, select a project that will be used to manage the restored Cloud SQL instance and specify a service account whose permissions will be used to perform the restore operation. For more information on the required permissions, see [Service Account Permissions](#).

For a project to be displayed in the **Project** drop-down list, it must be added to Veeam Backup for Google Cloud as described in section [Adding Projects and Folders](#). If you have not added the necessary project to Veeam Backup for Google Cloud beforehand, you can do it without closing the **Cloud SQL Instance Restore** wizard. To add a project, click **Add** and complete the **Add Projects and Folders** wizard.

For a service account to be displayed in the list of available accounts, it must be added to Veeam Backup for Google Cloud as described in section [Adding Service Accounts](#), and must be assigned permissions required to access the selected project as described in section [Adding Projects and Folders](#).

The screenshot shows the 'Cloud SQL Instance Restore' wizard in the Veeam Backup for Google Cloud interface. The top navigation bar includes the Veeam logo, the title 'Veeam Backup for Google Cloud', the server time 'Nov 18, 2023 5:19 PM', and user information 'administrator Portal Administrator'. The left sidebar contains a list of steps: Instances, Restore Mode, Project (highlighted), Region, Instance Settings, Network Settings, Security Settings, Flags, Verification, Reason, and Summary. The main content area is titled 'Specify project' and contains the following instructions and fields:

- Specify project**
Choose a project where the restored Cloud SQL instances will be created, and specify a service account that will be used to access the project. By default, the settings saved in the selected restore point will be used.
- Project**
Choose a target project.
Project: + Add
- Service account**
Specify a service account.
Service account:

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

Step 6. Select Region and Availability Zone

[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 **Region** step of the wizard, select a region where the restored Cloud SQL instance will operate and an availability zone for which you want to configure network settings.

To configure the restored Cloud SQL instance for high availability, select the **Multiple zones** option, and choose a primary and a secondary zone where the restored Cloud SQL instance will be located within the selected region. The high availability configuration allows you to reduce downtime when a zone or the instance becomes unavailable. For more information on high availability in Google Cloud, see [Google Cloud documentation](#).

TIP

If some of the restored Cloud SQL instances cannot be configured for high availability, the wizard will display a message notifying that the instances have issues with the original zone settings. To learn what these issues are, click the **instance** link in the message.

The screenshot shows the 'Cloud SQL Instance Restore' wizard in Veeam Backup for Google Cloud. The interface includes a top navigation bar with the Veeam logo, server time (Nov 18, 2023 5:20 PM), user information (administrator, Portal Administrator), and a Configuration icon. The main content area is titled 'Cloud SQL Instance Restore' and features a left-hand navigation menu with options: Instances, Restore Mode, Project, Region (highlighted), Instance Settings, Network Settings, Security Settings, Flags, Verification, Reason, and Summary. The 'Choose region and availability zone' section contains the following elements: a 'Region' field with 'us-east5 (Columbus)' selected; an 'Availability zone' section with a radio button for 'Single zone' selected and a dropdown menu showing 'Any', 'us-east5-a', 'us-east5-b', and 'us-east5-c'; and a 'Secondary zone:' dropdown menu with 'Any (different from primary)' selected. At the bottom of the wizard, there are three buttons: 'Previous', 'Next', and 'Cancel'.

Step 7. Specify Instance Name and Type

[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 **Instance Settings** step of the wizard, do the following:

1. Select the Cloud SQL instance.
2. If you want to specify a new name and a new machine type for the restored Cloud SQL instance, or to configure storage settings for the instance, click **Edit**.

In the **Configure general settings** window, specify the name and the machine type, and click **Apply**. To learn how to choose a machine type when creating a Cloud SQL instance in Google Cloud, see [Google Cloud documentation](#).

NOTE

Restore of PostgreSQL instances to Cloud SQL instances of the *db-f1-micro* and *db-g1-small* machine types is not supported. If you want to restore a PostgreSQL instance to one of the specified machine types, you must first manually create a Cloud SQL instance of the necessary type in the Google Cloud console as described in [Google Cloud documentation](#), and then restore the backed-up databases to the created instance as described in section [Performing Database Restore](#).

You can also choose a new storage type and manually increase storage capacity for the restored Cloud SQL instance. If you want Veeam Backup for Google Cloud to increase the storage capacity to fit the instance size automatically, select the **Automatic increase** check box. Note, however, that the amount of storage capacity allocated to an instance affects its cost. To learn how to configure storage settings when creating a Cloud SQL instance in Google Cloud, see [Google Cloud documentation](#).

TIP

If Veeam Backup for Google Cloud is unable to restore the Cloud SQL instance using the specified name for some reason, the wizard will display an error icon in the **Instance** column. To learn what this reason is, hover your mouse over the icon.

The screenshot shows the Veeam Backup for Google Cloud interface during a Cloud SQL Instance Restore process. The main window is titled "Cloud SQL Instance Restore". On the left, there is a navigation menu with options like "Instances", "Restore Mode", "Project", "Region", "Instance Settings", "Network Settings", "Security Settings", "Flags", "Verification", "Reason", and "Summary". The "Instance Settings" section is currently selected.

The "Configure instance settings" panel is active, showing a table with the following data:

Instance	Machine Type	Storage Type	Capacity	Auto-Incr
<input checked="" type="checkbox"/> tvg-mysql-1	—	HDD	100 GB	<input checked="" type="checkbox"/> Enable

The "Instance" column for the first row contains a red error icon. Below the table, it says "Selected: 1 of 1".

The "Configure general settings" panel is also visible, showing the following fields:

- Name: tvg-mysql-12
- Machine type: db-n1-standard-4 (...)
- Type: HDD, SSD
- Capacity: 100
- Automatic Increase

Buttons for "Apply" and "Cancel" are located at the bottom of the "Configure general settings" panel. At the bottom of the main window, there are "Previous" and "Next" buttons.

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 Settings** step of the wizard, do the following:

1. Select the Cloud SQL instance.
2. Click **Edit**.
3. In the **Edit network settings** window, choose whether you want to configure public IP and private IP connectivity for the restored Cloud SQL instance:
 - o To connect the restored Cloud SQL instance to a VPC network with a private IP address, select the **Private IP** check box and specify a VPC network to which the instance will be connected. For a VPC network to be displayed in the lists of available networks, it must contain a subnet that exists in the region specified at [step 6](#) of the wizard.

IMPORTANT

The specified VPC network must have private services access configured. To learn how to configure private services access for a VPC network, see [Google Cloud documentation](#).

- o To assign a public IPv4 address to the restored Cloud SQL instance and to accept connections to it from specific IP address ranges, set the **Public IP** toggle to *On*, click **Add** and then enter the allowed IP address ranges in the **Add Network Connection** window.

TIP

The IP address ranges must be specified in the CIDR notation (for example, 12.23.34.0/24). To let all IP addresses access the restored Cloud SQL instance, you can enter 0.0.0.0/0. However, note that allowing access from all IP addresses is unsafe and thus not recommended in production environments.

The screenshot shows the Veeam Backup for Google Cloud interface during a Cloud SQL Instance Restore. The main window is titled 'Cloud SQL Instance Restore' and features a sidebar with navigation options: Instances, Restore Mode, Project, Region, Instance Settings, Network Settings (selected), Security Settings, Flags, Verification, Reason, and Summary.

The 'Configure network settings' dialog is open, displaying a table with the following data:

Instance	Private IP	VPC
Selected: 1 of 1		
<input checked="" type="checkbox"/> tvg-mysql-12	Disabled	—

Below the table, there are 'Apply' and 'Cancel' buttons. The 'Previous' button is also visible at the bottom left of the dialog.

The 'Configure private IP settings' section includes a checkbox for 'Enable private IP' (checked), a VPC dropdown menu set to 'tvg-net', and a note: 'Only networks with Private Service Connection enabled are shown'.

The 'Configure public IP settings' section includes a 'Public IP' toggle set to 'Enabled', and buttons for '+ Add', 'Edit', and 'X Remove'.

Below these settings is a table for 'Allowed Network' and 'IP Range':

Allowed Network	IP Range
priv-net	89.177.50.123

Step 9. Configure Security 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 **Security Settings** step of the wizard, do the following:

1. Select the Cloud SQL instance.
2. Click **Edit**.
3. In the opened window, choose whether you want to connect to the restored Cloud SQL instance using SSL only, and whether you want the instance data to be encrypted with a Google Cloud KMS CMEK:

- If you want to secure connections to the restored Cloud SQL instance, set the **Allow only SSL connections** toggle to *On*.

Since SSL connections use digital certificates to provide encrypted access, make sure that you have obtained a Certificate Authority (CA) certificate, a client public key certificate, and a client private key – before you connect to the restored instance using SSL. For more information, see [Google Cloud documentation](#).

- If you want to apply the existing encryption scheme of the source Cloud SQL instance, select the **Use original encryption scheme** option.
- If you want to encrypt the restored data with a CMEK, select the **Use customer-managed encryption key from Google Cloud KMS** option and choose the necessary CMEK from the **Encryption key** drop-down list.

For a CMEK to be displayed in the list of available encryption keys, it must be stored in the region selected at [step 6](#) of the wizard.

NOTE

Due to technical limitations in Google Cloud, Veeam Backup for Google Cloud does not support data encryption with multi-regional keys. For more information, see [Cloud SQL for MySQL documentation](#) and [Cloud SQL for PostgreSQL documentation](#).

The screenshot displays the Veeam Backup for Google Cloud interface. At the top, the header shows the product name, server time (Nov 18, 2023 5:22 PM), and user information (administrator, Portal Administrator). The main content area is titled "Cloud SQL Instance Restore" and is divided into several sections:

- Instances:** A list of instances with columns for Instance, SSL Only, and Encryption. One instance, "tv-g-mysql-12", is selected and marked as "Disabled".
- Configure security settings:** A section with the heading "Enable encryption and allow connections through SSL." It includes a search bar and an "Edit" button.
- Configure SSL settings:** A modal dialog box with the following options:
 - Allow only SSL connections
 - Enable encryption:** Choose encryption settings for the restored data.
 - Use original encryption scheme
 - Use customer-managed encryption key from Google Cloud KMS
 - Encryption key: A dropdown menu.
 - An information icon and a link to the User Guide.
 - Buttons for "Apply" and "Cancel".

Step 10. Enable Flag Assignment

[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 **Flags** step of the wizard, you can instruct Veeam Backup for Google Cloud to modify flags set on databases of the restored Cloud SQL instance:

1. Select the Cloud SQL instance.
2. Click **Edit**.
3. In the **Configure flags** window, choose whether you want flags of the restored databases to have the same values as the source databases or new modified values.

If you want to set a new value for a database flag, select the flag and click **Edit**. If you want to clear all flags to their original values, click **Reset to Original**. To save changes made to the flag settings, click **Apply**.

The screenshot shows the Veeam Backup for Google Cloud interface. The top navigation bar includes the Veeam logo, the product name, the server time (Nov 18, 2023 5:23 PM), the user (administrator), and the current step (Configuration). The main content area is titled 'Cloud SQL Instance Restore' and is divided into two main sections: 'Modify database flags' and 'Configure flags'. The 'Modify database flags' section shows a search bar for the instance name and a list of instances, with 'tvq-mysql-12' selected. The 'Configure flags' section shows a table of database flags with their current values and states. An 'Edit Flag' dialog box is open over the 'autovacuum' flag, allowing the user to specify a new value. The dialog box contains a text input field with 'on' entered, a 'Reset' button, and 'Save' and 'Cancel' buttons. The 'Apply' and 'Cancel' buttons are also visible at the bottom of the 'Configure flags' section.

Flag	Value	State
shared_buffers	589824	Original
maintenance_work_mem	2147483647	Original
max_wal_size	2147483647	Original
pglogical.synchronous_commit		
autovacuum		
wal_buffers		

Step 11. Run Configuration Checks

At the **Verification** step of the wizard, Veeam Backup for Google Cloud will verify whether restore settings are configured properly and the specified service account has all the necessary permissions required to perform recovery tasks for the project that will manage the restored Cloud SQL instance. For more information on the required permissions, see [Service Account Permissions](#).

To see the list of missing permissions that must be granted to the service account in order to perform an operation, click the link in the **Details** column. You can grant the missing permissions to the service account [using the Google Cloud console](#) or instruct Veeam Backup for Google Cloud to do it:

- To grant the missing permissions manually, click **Download Script**. Veeam Backup for Google Cloud will generate a gcloud script that you can run in the Google Cloud console to assign all the necessary permissions to the service account.

The account under which you run the script must have the permissions required both to get and set project IAM policies and to create custom IAM roles (for example, it can have the *iam.securityAdmin* and *iam.roleAdmin* roles assigned). To learn what permissions and roles are required to create custom roles in IAM, see [Google Cloud documentation](#).

- To let Veeam Backup for Google Cloud grant the missing permissions automatically, click **Grant** and then click **Sign in with Google** in the **Grant permissions** window. You will be redirected to the OAuth consent screen authorization page. Sign in using credentials of a Google account that will be used to grant the permissions.

The account under which you sign in to Google Cloud must have the permissions required both to get and set project IAM policies and to create custom IAM roles (for example, it can have the *iam.securityAdmin* and *iam.roleAdmin* roles assigned). To learn what permissions and roles are required to create service account, see [Google Cloud documentation](#).

NOTE

For Veeam Backup for Google Cloud to be able to authorize in Google Cloud, the OAuth consent screen must be configured as described in section [Registering Applications](#). Note that Veeam Backup for Google Cloud does not store in the configuration database the provided Google account credentials and access tokens received during authorization.

To make sure that the missing permissions have been successfully granted, click **Recheck**.

The screenshot shows the Veeam Backup for Google Cloud interface. At the top, the header includes the Veeam logo, the product name 'Veeam Backup for Google Cloud', the server time 'Nov 18, 2023 5:23 PM', the user 'administrator Portal Administrator', and a 'Configuration' button. The main content area is titled 'Cloud SQL Instance Restore' and features a left-hand navigation menu with options like 'Instances', 'Restore Mode', 'Project', 'Region', 'Instance Settings', 'Network Settings', 'Security Settings', 'Flags', 'Verification' (which is highlighted), 'Reason', and 'Summary'. The main panel displays 'Run verification checks' with a sub-header 'Verify that permissions and configuration are correct.' Below this, there are three buttons: 'Recheck', 'Download Script', and 'Grant'. A table shows a single verification check for 'Cloud SQL Restore' with a 'Passed' result and a detail message: 'All the required permissions are gr...'. At the bottom of the interface, there are three buttons: 'Previous', 'Next', and 'Cancel'.

Check	Result	Details
Cloud SQL Restore	Passed	All the required permissions are gr...

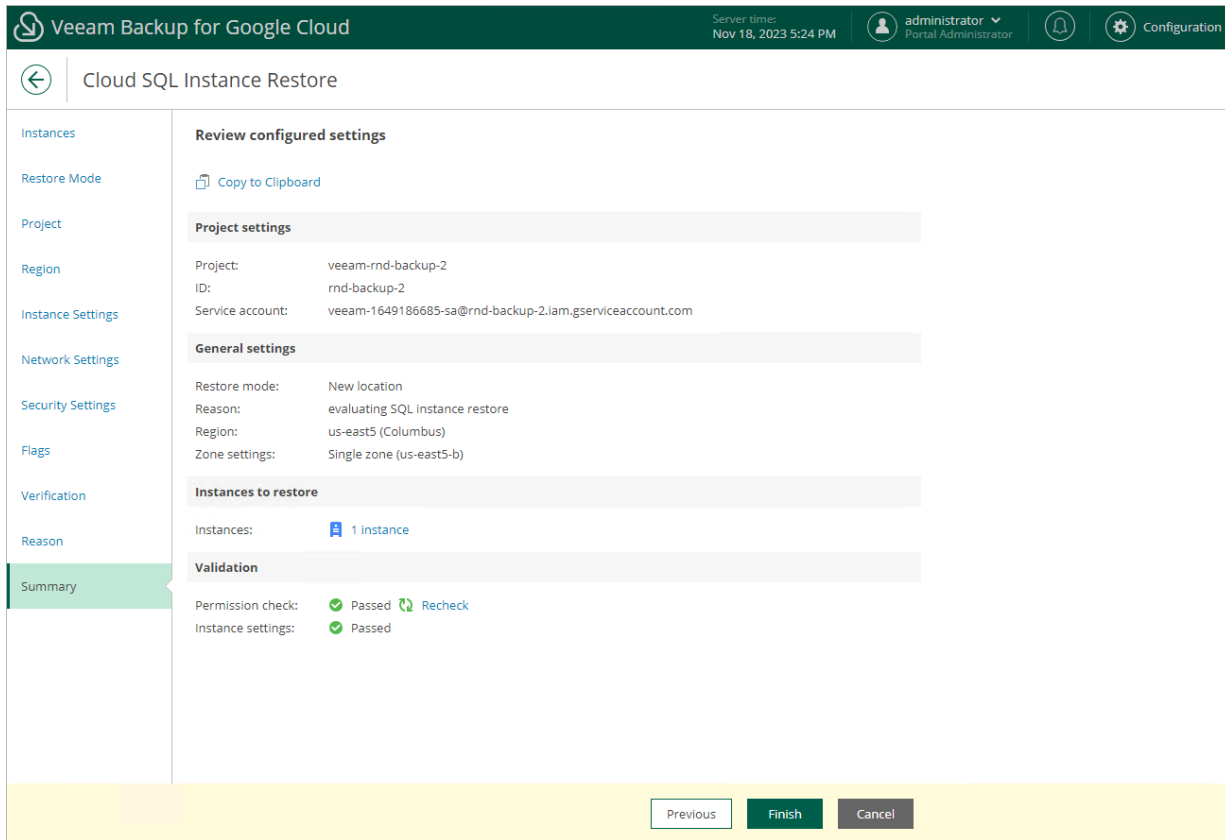
Step 12. Specify Restore Reason

At the **Reason** step of the wizard, specify a reason for restoring the Cloud SQL instance. This information will be saved to the session history, and you will be able to reference it later.

The screenshot shows the 'Cloud SQL Instance Restore' wizard in the Veeam Backup for Google Cloud interface. The top navigation bar includes the Veeam logo, the product name, the server time (Nov 18, 2023 5:23 PM), the user role (administrator / Portal Administrator), and a Configuration icon. The main content area is titled 'Cloud SQL Instance Restore' and features a left-hand navigation pane with the following steps: Instances, Restore Mode, Project, Region, Instance Settings, Network Settings, Security Settings, Flags, Verification, Reason (highlighted in green), and Summary. The main pane is titled 'Enter reason for this restore operation' and contains a text input field labeled 'Restore reason:' with the text 'evaluating SQL instance restore' entered. At the bottom of the wizard, there are three buttons: 'Previous' (disabled), 'Next' (active), and 'Cancel' (disabled).

Step 13. Finish Working with Wizard

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



Performing Database Restore

In case a disaster strikes, you can restore corrupted databases of a Cloud SQL instance from an image-level backup. Veeam Backup for Google Cloud allows you to restore databases to the original location or to a new location.

NOTE

Due to [technical limitations in Google Cloud](#), Veeam Backup for Google Cloud does not support restore to the original location if the source database is still present in the location.

To restore databases of a protected Cloud SQL instance, do the following:

1. [Launch the Database Restore wizard.](#)
2. [Select databases.](#)
3. [Select a project.](#)
4. [Configure target instance settings.](#)
5. [Check required permissions.](#)
6. [Specify a restore reason.](#)
7. [Finish working with the wizard.](#)

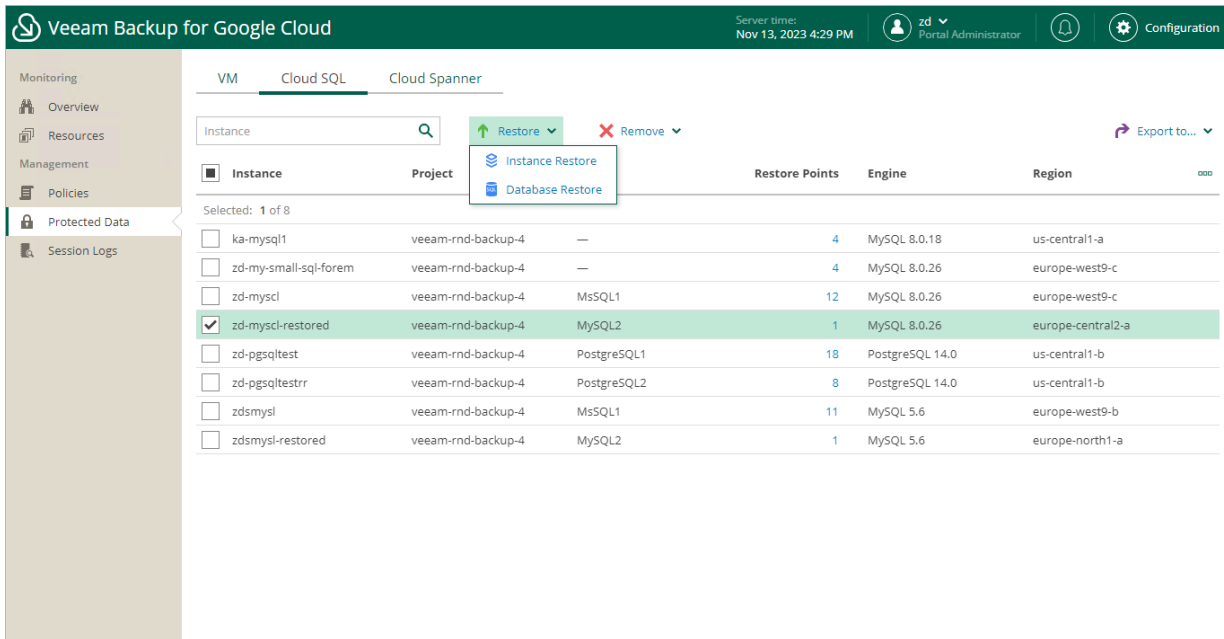
IMPORTANT

Before you start Cloud SQL database restore, make sure that network settings are configured for each region where worker instances will be deployed during the restore process. For information on how to configure network settings, see [Adding Worker Configurations](#).

Step 1. Launch Database Restore Wizard

To launch the **Database Restore** wizard, do the following:

1. Navigate to **Protected Data > Cloud SQL**.
2. Select the Cloud SQL instance whose databases you want to restore, and click **Restore > Database Restore**.



The screenshot shows the Veeam Backup for Google Cloud interface. The top navigation bar includes the Veeam logo, the title "Veeam Backup for Google Cloud", the server time "Nov 13, 2023 4:29 PM", and user information "zd Portal Administrator". The left sidebar shows the navigation menu with "Protected Data" selected. The main content area is titled "Cloud SQL" and displays a table of instances. The table has columns for Instance, Project, Restore Points, Engine, and Region. The instance "zd-myscl-restored" is selected, and a context menu is open over it, showing "Instance Restore" and "Database Restore" options.

Instance	Project	Restore Points	Engine	Region
ka-mysql1	veeam-rnd-backup-4	4	MySQL 8.0.18	us-central1-a
zd-my-small-sql-forem	veeam-rnd-backup-4	4	MySQL 8.0.26	europa-west9-c
zd-myscl	veeam-rnd-backup-4	12	MySQL 8.0.26	europa-west9-c
<input checked="" type="checkbox"/> zd-myscl-restored	veeam-rnd-backup-4	1	MySQL 8.0.26	europa-central2-a
zd-pgsqtest	veeam-rnd-backup-4	18	PostgreSQL 14.0	us-central1-b
zd-pgsqtestrr	veeam-rnd-backup-4	8	PostgreSQL 14.0	us-central1-b
zdsmysl	veeam-rnd-backup-4	11	MySQL 5.6	europa-west9-b
zdsmysl-restored	veeam-rnd-backup-4	1	MySQL 5.6	europa-north1-a

Step 2. Select Databases

At the **Databases** step of the wizard, click **Add** to select databases to restore, and then choose a restore point that will be used to restore the selected databases. By default, Veeam Backup for Google Cloud uses the most recent valid restore point. However, you can restore the backed-up data to an earlier state.

To select a restore point, do the following:

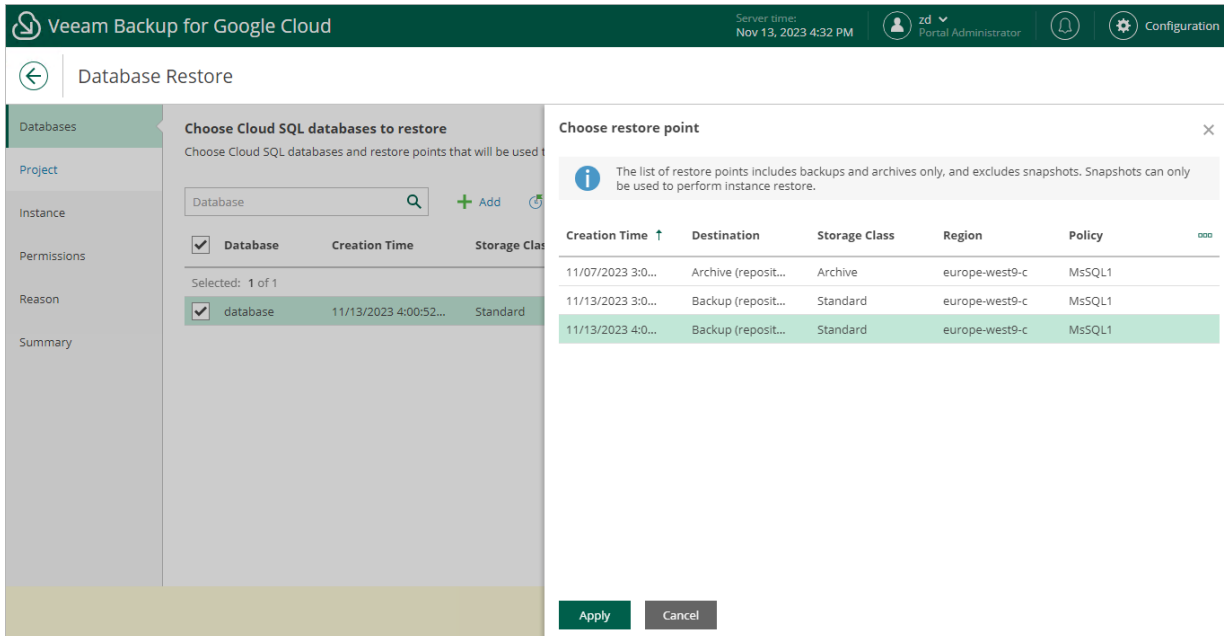
1. Select a database.
2. Click **Restore Point**.
3. In the **Choose restore point** window, select the necessary restore point and click **Apply**.

To help you choose a restore point, Veeam Backup for Google Cloud provides the following information on each available restore point:

- **Creation Time** – the date when the restore point was created.
- **Destination** – the type of the restore point:
 - § *Snapshot* – a cloud-native snapshot 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.
- **Storage Class** – the storage class of the backup repository where the restore point is stored (applies only to image-level backups).
- **Region** – a region in which the protected Cloud SQL instance resides.
- **Policy** – a backup policy that created the restore point.
- **Retention** – a retention configured for the backup policy that created the restore point.

NOTE

Veeam Backup for Google Cloud does not support restore of the **postgres** database, that is, the default database automatically added to PostgreSQL instances upon creation. Consider that it is not recommended to use this database to store any data. For more information, see [Google Cloud documentation](#).



The screenshot shows the Veeam Backup for Google Cloud interface. The top navigation bar includes the Veeam logo, the product name, the server time (Nov 13, 2023 4:32 PM), the user (zd Portal Administrator), and a Configuration icon. The main content area is titled 'Database Restore' and is divided into two panels.

The left panel, 'Choose Cloud SQL databases to restore', has a search bar and a table with the following data:

Database	Creation Time	Storage Class
<input checked="" type="checkbox"/> database	11/13/2023 4:00:52...	Standard

The right panel, 'Choose restore point', contains an information message and a table of restore points:

The list of restore points includes backups and archives only, and excludes snapshots. Snapshots can only be used to perform instance restore.

Creation Time ↑	Destination	Storage Class	Region	Policy
11/07/2023 3:0...	Archive (reposit...	Archive	europa-west9-c	M5SQL1
11/13/2023 3:0...	Backup (reposit...	Standard	europa-west9-c	M5SQL1
11/13/2023 4:0...	Backup (reposit...	Standard	europa-west9-c	M5SQL1

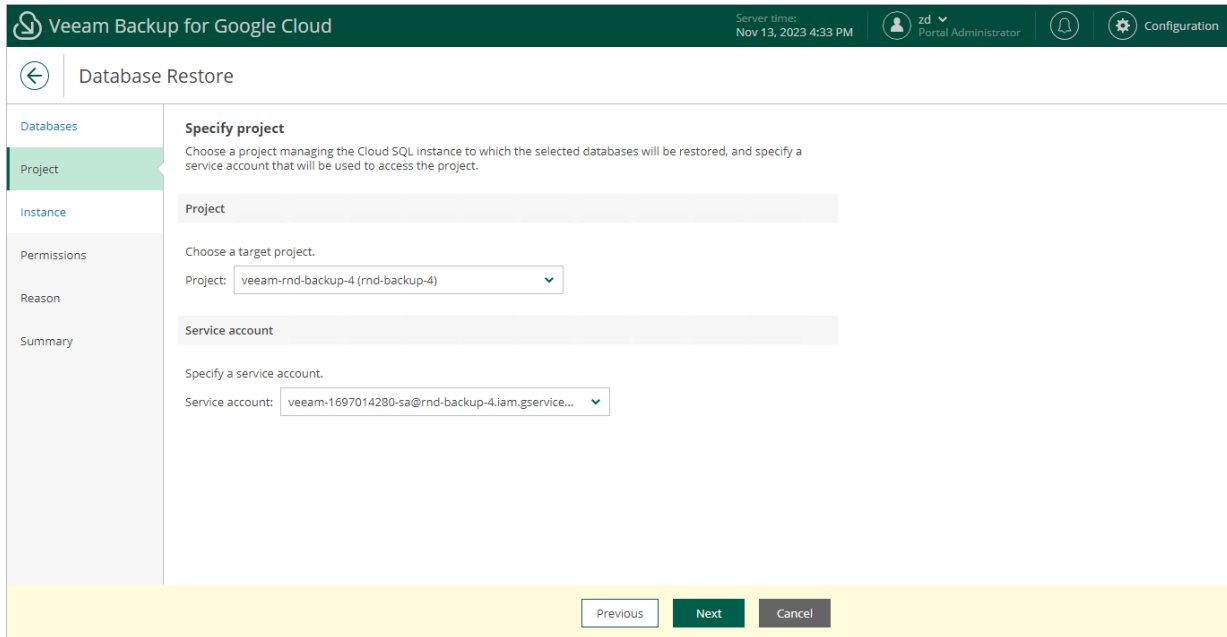
At the bottom of the right panel, there are 'Apply' and 'Cancel' buttons.

Step 3. Select Project

At the **Project** step of the wizard, select a project that manages a Cloud SQL instance to which you want to restore the selected databases and specify a service account whose permissions will be used to perform the restore operation. For more information on the required permissions, see [Service Account Permissions](#).

For a project to be displayed in the **Project** drop-down list, it must be added to Veeam Backup for Google Cloud as described in section [Adding Projects and Folders](#).

For a service account to be displayed in the list of available accounts, it must be added to Veeam Backup for Google Cloud as described in section [Adding Service Accounts](#), and must be assigned the *Cloud SQL Instances Restore* operational role as described in section [Adding Projects and Folders](#).



Step 4. Configure Target Instance Settings

At the **Instance** step of the wizard, choose a Cloud SQL instance that will host the restored databases. To do that, click the link in the **Instance** field, select the necessary Cloud SQL instance from the **Choose Cloud SQL instance** list, and click **Apply**. For a Cloud SQL instance to be displayed in the list of available instances, it must belong to the selected project and be running on a supported database engine.

NOTES

- Restore to Cloud SQL instances configured to accept SSL connections is not supported.
- PostgreSQL databases can be restored only to PostgreSQL instances running the same database engine version.

You must also specify a Cloud SQL account whose credentials will be used to perform the restore operation. To do that, click a link in the **Credentials** field and choose an account from the list of available Cloud SQL accounts. For an account to be displayed in the list, it must be added to Veeam Backup for Google Cloud as described in section [Adding Cloud SQL Accounts](#). If you have not added the necessary account to Veeam Backup for Google Cloud beforehand, you can do it without closing the **Database Restore** wizard. To add an account, click **Add** and complete the **Add Account** wizard.

TIP

Veeam Backup for Google Cloud will perform a number of configuration checks for the selected instance and databases:

- If any of the checks fail to complete successfully for an instance, the wizard will display an error in the **Configuration checks** field.
- If any of the checks fail to complete successfully for a database, the wizard will display an error in the **Checks column** of the **Databases to restore** table.

You can click the link to get more information on an error.

The screenshot shows the 'Database Restore' wizard in the 'Instance' step. The interface includes a top navigation bar with the Veeam logo, 'Veeam Backup for Google Cloud', server time (Nov 13, 2023 4:52 PM), and user information (zd Portal Administrator). A left sidebar contains navigation options: Databases, Project, Instance (selected), Permissions, Reason, and Summary. The main content area is titled 'Choose instance' and includes the instruction: 'Specify a Cloud SQL instance where the selected databases will be restored.' The instance details are: Instance: zd-myscl, Credentials: Default (veeam-1697014280-sa@rnd-backup-4.iam.gserviceaccount.com), Engine: MySQL, Version: 8.0.26, and Configuration checks: Passed. An information box states: 'Stored procedures and triggers will also be restored to the specified instance.' Below this is a table titled 'Databases to restore' with columns: Database, Checks, Restore Point, Storage Class, and Region. The table contains one row: database, Passed, 11/13/2023 4:00:52 PM, Standard, europe-west... At the bottom, there are 'Previous', 'Next', and 'Cancel' buttons.

Database	Checks	Restore Point	Storage Class	Region
database	Passed	11/13/2023 4:00:52 PM	Standard	europe-west...

Step 5. Check Required Permissions

At the **Permissions** step of the wizard, Veeam Backup for Google Cloud will verify whether the specified service account has all the necessary permissions required to perform data recovery tasks for the selected project. For more information on the required permissions, see [Service Account Permissions](#).

To see the list of missing permissions that must be granted to the service account in order to perform an operation, click the link in the **Details** column. You can grant the missing permissions to the service account [using the Google Cloud console](#) or instruct Veeam Backup for Google Cloud to do it:

- To grant the missing permissions manually, click **Download Script**. Veeam Backup for Google Cloud will generate a gcloud script that you can run in the Google Cloud console to assign all the necessary permissions to the service account.

The account under which you run the script must have the permissions required both to get and set project IAM policies and to create custom IAM roles (for example, it can have the *iam.securityAdmin* and *iam.roleAdmin* roles assigned). To learn what permissions and roles are required to create custom roles in IAM, see [Google Cloud documentation](#).

- To let Veeam Backup for Google Cloud grant the missing permissions automatically, click **Grant** and then click **Sign in with Google** in the **Grant permissions** window. You will be redirected to the OAuth consent screen authorization page. Sign in using credentials of a Google account that will be used to grant the permissions.

The account under which you sign in to Google Cloud must have the permissions required both to get and set project IAM policies and to create custom IAM roles (for example, it can have the *iam.securityAdmin* and *iam.roleAdmin* roles assigned). To learn what permissions and roles are required to create service account, see [Google Cloud documentation](#).

NOTE

For Veeam Backup for Google Cloud to be able to authorize in Google Cloud, the OAuth consent screen must be configured as described in section [Registering Applications](#). Note that Veeam Backup for Google Cloud does not store in the configuration database the provided Google account credentials and access tokens received during authorization.

To make sure that the missing permissions have been successfully granted, click **Recheck**.

The screenshot shows the 'Check permissions' step of the 'Database Restore' wizard. The interface includes a top navigation bar with the Veeam logo, server time (Nov 13, 2023 4:56 PM), user profile (zd Portal Administrator), and a Configuration icon. A left sidebar contains navigation links: Databases, Project, Instance, Permissions (highlighted), Reason, and Summary. The main content area is titled 'Check permissions' and contains the instruction 'Verify whether all the required permissions are granted.' Below this are three buttons: 'Recheck' (with a refresh icon), 'Download Script' (with a download icon), and 'Grant' (with a plus icon). A table displays the results of the permission checks:

Check	Result	Details
Cloud SQL Restore	Passed	All the required permissions are ...
Worker	Passed	All the required permissions are ...
Repository	Passed	All the required permissions are ...

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

Step 6. Specify Restore Reason

At the **Reason** step of the wizard, specify a reason for restoring the Cloud SQL databases. This information will be saved to the session history, and you will be able to reference it later.

The screenshot shows the 'Database Restore' wizard in Veeam Backup for Google Cloud. 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 product name 'Veeam Backup for Google Cloud', the server time 'Nov 13, 2023 4:58 PM', the user 'zd Portal Administrator', and a 'Configuration' link. The left sidebar contains a list of steps: 'Databases', 'Project', 'Instance', 'Permissions', 'Reason' (which is highlighted in green), and 'Summary'. The main content area is titled 'Enter reason for this restore operation' and features a text input field with the text 'evaluating database restore'. At the bottom of the wizard, 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 'Database Restore' wizard in the 'Summary' step. The interface includes a top navigation bar with the Veeam logo, server time (Nov 13, 2023 5:00 PM), and user information (zd Portal Administrator). A left sidebar lists navigation options: Databases, Project, Instance, Permissions, Reason, and Summary (highlighted). The main content area is titled 'Review configured settings' and includes a 'Copy to Clipboard' button. It is divided into four sections: 'Project settings' (Project: veeam-rnd-backup-4, ID: rnd-backup-4, Service account: veeam-1697014280-sa@rnd-backup-4.iam.gserviceaccount.com), 'General settings' (Reason: evaluating database restore, Region: europe-west9), 'Databases to restore' (Instance: zd-my-small-sql-forem, Databases: 1 database), and 'Validation' (Permission check: Passed, Instance settings: Passed). At the bottom, there are three buttons: 'Previous', 'Finish', and 'Cancel'.

Review configured settings

[Copy to Clipboard](#)

Project settings

Project: veeam-rnd-backup-4
ID: rnd-backup-4
Service account: veeam-1697014280-sa@rnd-backup-4.iam.gserviceaccount.com

General settings

Reason: evaluating database restore
Region: europe-west9

Databases to restore

Instance: [zd-my-small-sql-forem](#)
Databases: [1 database](#)

Validation

Permission check: ✓ Passed [Recheck](#)
Instance settings: ✓ Passed

[Previous](#) [Finish](#) [Cancel](#)

Spanner Restore

The actions that you can perform with restore points of Cloud Spanner instances depend on whether you access the restore points using the Veeam Backup & Replication console or the Veeam Backup for Google Cloud Web UI.


Spanner Restore Using Console

You can recover corrupted Cloud Spanner instances and databases in the Veeam Backup for Google Cloud Web UI only. However, you can launch the **Cloud Spanner 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 Cloud Spanner instances that you want to recover, select the necessary instance and click **Google Cloud Spanner** on the ribbon.

Alternatively, you can right-click the selected instance and click **Restore to Google Cloud Spanner**.

Veeam Backup & Replication will open the **Cloud Spanner Restore** wizard in a web browser. Complete the wizard as described in section [Performing Spanner Restore](#).



Spanner Restore Using Web UI

Veeam Backup for Google Cloud offers the following restore operations:

- [Instance restore](#) — start an entire Cloud Spanner instance from a restore point.
- [Database restore](#) — restore specific databases of a Cloud Spanner instance.

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

Performing Spanner Instance Restore

In case a disaster strikes, you can restore an entire Cloud Spanner instance from a cloud-native snapshot or image-level backup. Veeam Backup for Google Cloud allows you to restore one or more Cloud Spanner instances at a time, to the original location or to a new location.

NOTE

Veeam Backup for Google Cloud does not support restore to the original location if the source Cloud Spanner instance is still present in the location or if its name is reserved. You can delete the instance; however, keep in mind that you must delete all its cloud-native snapshots first — due to [technical limitations in Google Cloud](#).

To restore a protected Cloud Spanner instance, do the following:

1. [Launch Cloud Spanner instance restore wizard](#)
2. [Select a restore point](#).
3. [Choose a restore mode](#).
4. [Select a service account](#).
5. [Select a project](#).
6. [Select a region and an availability zone](#).
7. [Specify a new name and machine type for the instance](#).
8. [Configure encryption settings](#).
9. [Run configuration and permission checks](#).
10. [Specify a restore reason](#).
11. [Finish working with the wizard](#).

IMPORTANT

Before you start Cloud Spanner instance restore, make sure that network settings are configured for each region where worker instances will be deployed during the restore process. For information on how to configure network settings, see [Adding Worker Configurations](#).

Step 1. Launch Cloud Spanner Instance Restore Wizard

To launch the **Cloud Spanner Instance Restore** wizard, do the following:

1. Navigate to **Protected Data > Cloud Spanner**.
2. Select the Cloud Spanner instance that you want to restore, and click **Restore > Instance Restore**.

The screenshot displays the Veeam Backup for Google Cloud interface. The top navigation bar includes the Veeam logo, the title "Veeam Backup for Google Cloud", the server time "Nov 8, 2023 2:12 PM", and user information "administrator Portal Administrator". The left sidebar shows a navigation menu with "Protected Data" selected. The main content area is titled "Cloud Spanner" and features a search bar, a "Restore" dropdown menu (with "Instance Restore" and "Database Restore" options), and a "Remove" button. Below this is a table of instances with columns for "Instance", "Project", "Policy", "Restore Points", and "Configuration". The instance "pim-multireg1" is selected, and its "Restore" dropdown menu is open, showing "Instance Restore" and "Database Restore" options.

Instance	Project	Policy	Restore Points	Configuration
<input type="checkbox"/> pim-4tb	RnD Backup	—	7	—
<input type="checkbox"/> pim-huge-dbs	RnD Backup	—	16	—
<input type="checkbox"/> pim-huge-dbs	RnD Backup	—	9	—
<input checked="" type="checkbox"/> pim-multireg1	RnD Backup	—	1	—

Step 2. Select Restore Point

At the **Instances** step of the wizard, select a restore point that will be used to restore the selected Cloud Spanner instance. By default, Veeam Backup for Google Cloud 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. Select the Cloud Spanner 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 Google Cloud provides the following information on each available restore point:

- **Creation Time** – the date when the restore point was created.
- **Destination** – the type of the restore point:
 - § *Snapshot* – a cloud-native snapshot 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 result of the latest health check performed for the restore point.
- **Storage Class** – the storage class of a backup repository where the restore point is stored (applies only to image-level backups).
- **Policy** – a backup policy that created the restore point.
- **Project** – a project that manages the protected Cloud Spanner instance.

- o **Retention** – a retention configured for the backup policy that created the restore point.

Choose Cloud Spanner instances to restore

Choose Cloud Spanner instances and the restore point

Instance

Instance	Restore Point
<input checked="" type="checkbox"/> pim-multireg1	09/15/2023 7:34:00

Selected: 1 of 1

Choose restore point

Creation Ti...	Destination	State	Storage Cla...	Policy	Project
09/15/2023 ...	Snapshot	—	—	—	RnD Backup

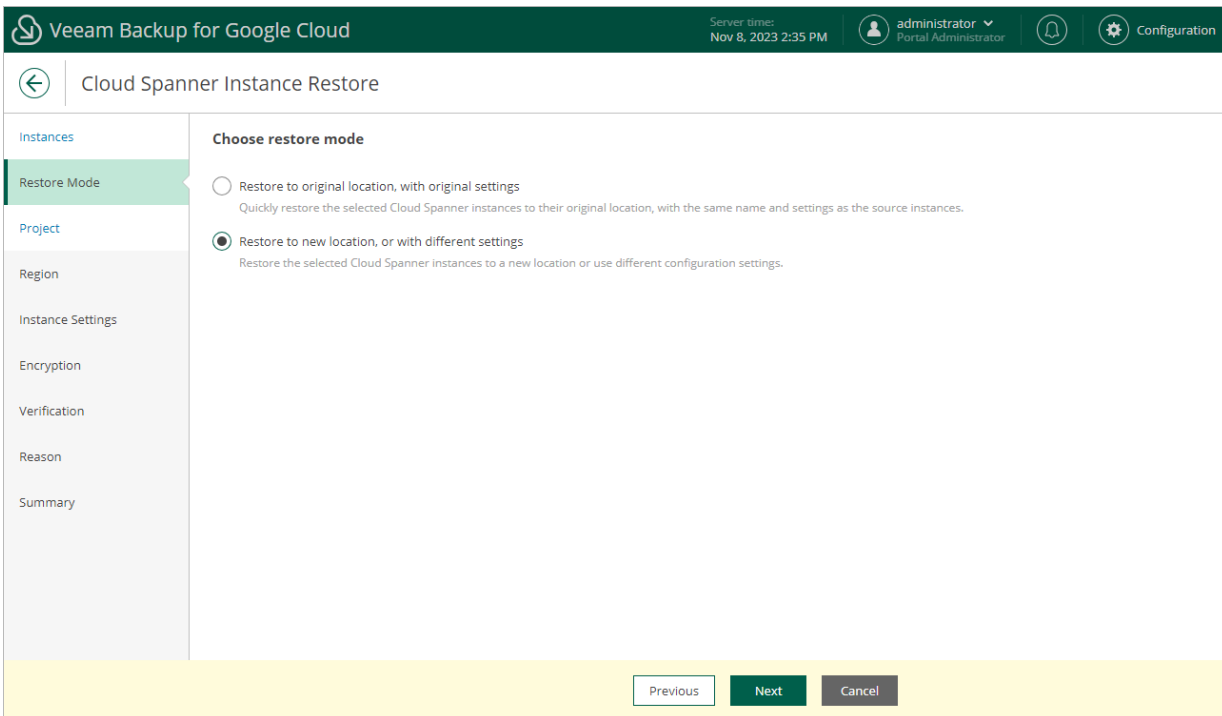
Apply Cancel

Step 3. Choose Restore Mode

At the **Restore Mode** step of the wizard, choose whether you want to restore the selected Cloud Spanner instance to the original or to a new location.

IMPORTANT

Restore to the original location is supported only using restore points of the *Backup* and *Archive* types. If you select a restore point of the *Snapshot* or *Manual Snapshot* type at [step 2](#) of the wizard, you will be able to select the **Restore to original** option and proceed with the wizard but only up to the **Verification** step — at this step, the verification check will notify you that the restore settings have not been configured properly. As a result, Veeam Backup for Google Cloud will not be able to perform the operation.



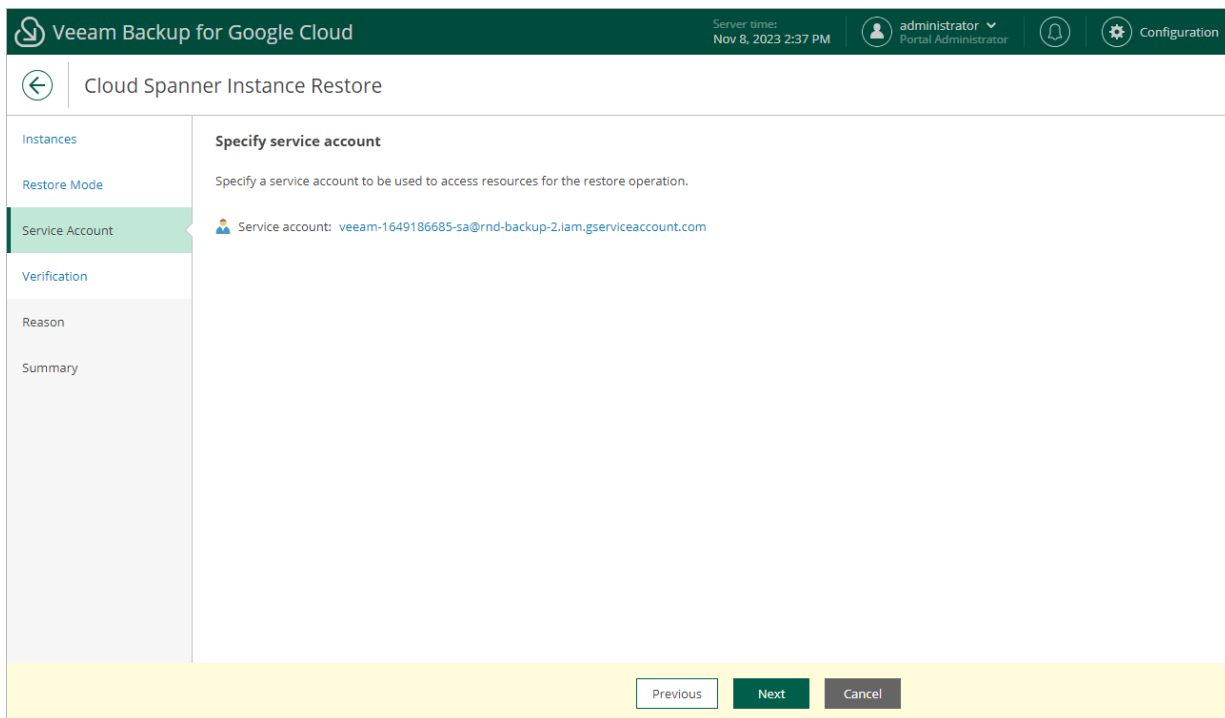
Step 4. Select Service Account

[This step applies only if you have selected the **Restore to original location, with original settings** option at the **Restore Mode** step of the wizard]

At the **Service Account** step of the wizard, select a service account whose permissions will be used to perform the restore operation. For more information on the required permissions, see [Service Account Permissions](#).

For a service account to be displayed in the **Service account** drop-down list, it must be added to Veeam Backup for Google Cloud as described in section [Adding Service Accounts](#), and must be assigned the *Cloud Spanner Instances Restore* operational role as described in section [Adding Projects and Folders](#).

If you have not added the necessary service account to Veeam Backup for Google Cloud beforehand, you can do it without closing the **Cloud Spanner Instance Restore** wizard. To add an account, click **Add** and complete the **Add Service Account** wizard.



Step 5. Select Project

[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 **Project** step of the wizard, select a project that will be used to manage the restored Cloud Spanner instance and specify a service account whose permissions will be used to perform the restore operation. For more information on the required permissions, see [Service Account Permissions](#).

For a project to be displayed in the **Project** drop-down list, it must be added to Veeam Backup for Google Cloud as described in section [Adding Projects and Folders](#). If you have not added the necessary project to Veeam Backup for Google Cloud beforehand, you can do it without closing the **Cloud Spanner Instance Restore** wizard. To add a project, click **Add** and complete the **Add Projects and Folders** wizard.

For a service account to be displayed in the list of available accounts, it must be added to Veeam Backup for Google Cloud as described in section [Adding Service Accounts](#), and must be assigned permissions required to access the selected project as described in section [Adding Projects and Folders](#).

The screenshot shows the 'Specify project' step of the 'Cloud Spanner Instance Restore' wizard. The interface includes a top navigation bar with the Veeam logo, 'Veeam Backup for Google Cloud', server time 'Nov 8, 2023 2:42 PM', and user 'administrator Portal Administrator'. A left sidebar lists steps: Instances, Restore Mode, Project (highlighted), Region, Instance Settings, Encryption, Verification, Reason, and Summary. The main content area is titled 'Specify project' and contains instructions: 'Choose a project where the restored Cloud Spanner instances will be created, and specify a service account that will be used to access the project. By default, the settings saved in the selected restore point will be used.' Below this, there are two sections: 'Project' with a dropdown menu showing 'rnd-qa-monitoring (rnd-qa-monitoring)' and an '+ Add' button; and 'Service account' with a text field showing 'veeam-1649186685-sa@rnd-backup-2.iam.gserviceaccount.com'. At the bottom, there are 'Previous', 'Next', and 'Cancel' buttons.

Step 6. Configure Regional 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 **Region** step of the wizard, select an instance configuration for the restored Cloud Spanner instance. The configuration defines the geographic location where the instance data will be stored.

To configure the restored Cloud Spanner instance for high availability, select the *Multi-region* configuration, and choose base configurations that contain regions where replicas of the restored Cloud Spanner instance will be placed. The high availability configuration allows you to reduce the chance of downtime in case a zone or an entire region becomes unavailable. For more information on high availability and instance configurations in Google Cloud, see [Google Cloud documentation](#).

TIP

If some of the restored Cloud Spanner instances cannot be configured for high availability, the wizard will display a message notifying that the instances have issues with the original zone settings. To learn what these issues are, click the **Instances** link in the message.

You can also add optional read-only replicas for both regional and multi-region instance configurations to increase your read capacity and data availability. If you set the **Additional read-only replicas** toggle to *On*, you must specify both regions where read-only replicas of the restored Cloud Spanner instance will be placed and the number of replicas in each region. If the **Read-only region** list does not include the location that you want to add, you can request a new optional read-only replica region as described in [Google Cloud documentation](#).

However, note that adding read-only replicas may increase read latency in case a read-only replica is added to a region belonging to a continent other than the one where replicas of the restored Cloud Spanner instance are located. To maintain low read latency in this scenario, it is recommended that you add 2 read-only replicas to a three-continent configuration as described in [Google Cloud documentation](#).

The screenshot shows the 'Configure regional settings' step of the 'Cloud Spanner Instance Restore' wizard. The interface includes a top navigation bar with the Veeam logo, server time (Nov 8, 2023 2:43 PM), and user information (administrator, Portal Administrator). A left sidebar lists navigation options: Instances, Restore Mode, Project, Region (highlighted), Instance Settings, Encryption, Verification, Reason, and Summary. The main content area is titled 'Configure regional settings' and contains the following elements:

- Configuration:** A dropdown menu set to 'Regional'.
- Project:** A location pin icon followed by 'us-central1 (Iowa)'.
- Additional read-only replicas:** A toggle switch that is currently turned on.
- Read-only region:** A dropdown menu set to 'europe-west9'.
- Number of replicas:** A numeric input field set to '1' with a '+ Add' button to its right.
- Table:** A table with two columns: 'Region' and 'Replicas'. It contains one row with 'europe-west9' in the 'Region' column and '1' in the 'Replicas' column. A red 'X' icon is positioned to the right of the row.

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

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]

At the **Instance Settings** step of the wizard, do the following:

1. Select the Cloud Spanner instance.
2. If you want to specify a new name and a new ID for the restored Cloud Spanner instance, or to configure compute capacity settings for the instance, click **Edit**.

In the **Configure general settings** window, specify the name and the ID, and click **Apply**.

You can also choose a new measurement unit and manually increase compute capacity for the restored Cloud Spanner instance. Note, however, that the amount of compute capacity allocated to an instance affects its cost. To learn how to configure compute capacity settings when creating a Cloud Spanner instance in Google Cloud, see [Google Cloud documentation](#).

TIP

If Veeam Backup for Google Cloud is unable to restore the Cloud Spanner instance using the specified ID for some reason, the wizard will display an error icon in the **Instance** column. To learn what this reason is, hover your mouse over the icon.

The screenshot shows the Veeam Backup for Google Cloud interface during the 'Cloud Spanner Instance Restore' process. The 'Instance Settings' tab is selected in the left sidebar. The main area displays 'Configure instance settings' for a table of instances. One instance, 'pim-multireg1', is selected. A 'Configure general settings' dialog box is open, allowing configuration of the instance's name, ID, measurement unit, and compute capacity. The 'Apply' button is highlighted.

Instance	ID	Measurement Unit
<input checked="" type="checkbox"/> pim-multireg1	pim-multireg1	Nodes

Selected: 1 of 1

Configure general settings

Name: pim-multireg1

ID: pim-multireg1

Configure compute capacity settings

Measurement unit: Nodes

Compute capacity: 1

The minimum number of nodes is 1. A node is equal to 1000 processing units. For values below 1000, use the Processing units option.

Apply Cancel

Step 8. Configure Encryption 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 **Encryption** step of the wizard, do the following:

1. Select the Cloud Spanner instance.
2. Click **Edit**.
3. In the opened window, choose whether you want the instance data to be encrypted with a Google Cloud KMS CMEK:
 - If you want to apply the existing encryption scheme of the source Cloud Spanner instance, select the **Use original encryption scheme** option.
 - If you want to apply Google-managed encryption scheme, select the **Use Google-managed encryption key** option.
 - If you want to encrypt the restored data with a CMEK, select the **Use customer-managed encryption key from Google Cloud KMS** option and choose the necessary CMEK from the **Encryption key** drop-down list.

For a CMEK to be displayed in the list of available encryption keys, it must be stored in the region selected at [step 6](#) of the wizard.

NOTES

- Due to [technical limitations in Google Cloud](#), Veeam Backup for Google Cloud does not support data encryption with multi-regional keys.
- Due to [technical limitations in Google Cloud](#), encrypting data with CMEKs is not supported for custom instance configurations with optional read-only replicas. If you want the instance data to be encrypted with a CMEK, the key must be stored in the same location as the restored Cloud Spanner instance (that is, for regional configuration – in the same region, and for multi-regional configuration – in the same multi-regional location).

The screenshot displays the Veeam Backup for Google Cloud interface. At the top, the header shows 'Veeam Backup for Google Cloud' and 'Server time: Nov 8, 2023 2:44 PM'. The user is logged in as 'administrator Portal Administrator'. The main content area is titled 'Cloud Spanner Instance Restore'. On the left, a navigation pane lists 'Instances', 'Restore Mode', 'Project', 'Region', 'Instance Settings', 'Encryption', 'Verification', 'Reason', and 'Summary'. The 'Encryption' section is active, showing 'Enable encryption' with a search bar and an 'Edit' button. Below this, a table lists instances with checkboxes for 'Instance' and 'Encryption'. The instance 'pim-multireg1' is selected, and its encryption type is 'Google-managed'. A modal dialog titled 'Configure encryption settings' is open, prompting the user to 'Choose encryption settings for the restored data.' The dialog offers three options: 'Use original encryption scheme', 'Use Google-managed encryption key' (which is selected), and 'Use customer-managed encryption key from Google Cloud KMS'. There is an 'Encryption key:' dropdown menu and an information icon with a link to the 'User Guide'. At the bottom of the dialog are 'Apply' and 'Cancel' buttons.

Step 9. Run Verification Checks

At the **Verification** step of the wizard, Veeam Backup for Google Cloud will verify whether restore settings are configured properly and the specified service account has all the necessary permissions required to perform recovery tasks for the project that will manage the restored Cloud Spanner instance. For more information on the required permissions, see [Service Account Permissions](#).

To see the list of missing permissions that must be granted to the service account in order to perform an operation, click the link in the **Details** column. You can grant the missing permissions to the service account [using the Google Cloud console](#) or instruct Veeam Backup for Google Cloud to do it:

- To grant the missing permissions manually, click **Download Script**. Veeam Backup for Google Cloud will generate a gcloud script that you can run in the Google Cloud console to assign all the necessary permissions to the service account.

The account under which you run the script must have the permissions required both to get and set project IAM policies and to create custom IAM roles (for example, it can have the *iam.securityAdmin* and *iam.roleAdmin* roles assigned). To learn what permissions and roles are required to create custom roles in IAM, see [Google Cloud documentation](#).

- To let Veeam Backup for Google Cloud grant the missing permissions automatically, click **Grant** and then click **Sign in with Google** in the **Grant permissions** window. You will be redirected to the OAuth consent screen authorization page. Sign in using credentials of a Google account that will be used to grant the permissions.

The account under which you sign in to Google Cloud must have the permissions required both to get and set project IAM policies and to create custom IAM roles (for example, it can have the *iam.securityAdmin* and *iam.roleAdmin* roles assigned). To learn what permissions and roles are required to create service account, see [Google Cloud documentation](#).

NOTE

For Veeam Backup for Google Cloud to be able to authorize in Google Cloud, the OAuth consent screen must be configured as described in section [Registering Applications](#). Note that Veeam Backup for Google Cloud does not store in the configuration database the provided Google account credentials and access tokens received during authorization.

To make sure that the missing permissions have been successfully granted, click **Recheck**.

The screenshot shows the Veeam Backup for Google Cloud interface. The top header includes the Veeam logo, the product name, the server time (Nov 8, 2023 2:44 PM), the user (administrator Portal Administrator), and a Configuration icon. The main content area is titled 'Cloud Spanner Instance Restore' and features a left-hand navigation menu with options: Instances, Restore Mode, Project, Region, Instance Settings, Encryption, Verification (highlighted), Reason, and Summary. The main panel displays 'Run verification checks' with the instruction 'Verify that permissions and configuration are correct.' Below this are three buttons: 'Recheck' (with a refresh icon), 'Download Script' (with a download icon), and 'Grant' (with a plus icon). A table below shows the results of the verification check:

Check	Result	Details
Cloud Spanner Restore	Passed	All the required permissions are gra...

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

Step 10. Specify Restore Reason

At the **Reason** step of the wizard, specify a reason for restoring the Cloud Spanner instance. This information will be saved to the session history, and you will be able to reference it later.

The screenshot displays the Veeam Backup for Google Cloud interface during the 'Cloud Spanner Instance Restore' process. The top navigation bar includes the Veeam logo, the product name, the server time (Nov 8, 2023 2:44 PM), the user role (administrator Portal Administrator), and a configuration icon. The main content area is titled 'Cloud Spanner Instance Restore' and features a left-hand navigation pane with options: Instances, Restore Mode, Project, Region, Instance Settings, Encryption, Verification, Reason (highlighted in green), and Summary. The main pane is titled 'Enter reason for this restore operation' and contains a text input field labeled 'Restore reason:' with the text 'evaluating instance restore' entered. At the bottom of the interface, there are three buttons: 'Previous', 'Next', and 'Cancel'.

Step 11. Finish Working with Wizard

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

The screenshot shows the 'Cloud Spanner Instance Restore' wizard in the 'Summary' step. The interface includes a top navigation bar with the Veeam logo, server time (Nov 8, 2023 2:44 PM), user (administrator), and configuration icons. A left sidebar lists steps: Instances, Restore Mode, Project, Region, Instance Settings, Encryption, Verification, Reason, and Summary (highlighted). The main content area is titled 'Review configured settings' and contains a 'Copy to Clipboard' button. Below are sections for 'Project settings' (Project: rnd-qa-monitoring, ID: rnd-qa-monitoring, Service account: veeam-1649186685-sa@rnd-backup-2.iam.gserviceaccount.com), 'General settings' (Restore mode: New location, Instance configuration: us-central1 (Iowa), Regional configuration: Regional, Reason: evaluating instance restore), 'Instances to restore' (1 instance, 3 databases), and 'Validation' (Permission check: Passed, Recheck; Instance settings: Passed). At the bottom are 'Previous', 'Finish', and 'Cancel' buttons.

Performing Database Restore

In case a disaster strikes, you can restore corrupted databases of a Cloud Spanner instance from an image-level backup or a cloud-native snapshot. Veeam Backup for Google Cloud allows you to restore databases to the original location or to a new location.

NOTE

Due to [technical limitations in Google Cloud](#), Veeam Backup for Google Cloud does not support restore to the original location if the source database is still present in the location.

To restore databases of a protected Cloud Spanner instance, do the following:

1. [Launch the Cloud Spanner Database Restore wizard.](#)
2. [Select databases.](#)
3. [Select a project.](#)
4. [Configure target instance settings.](#)
5. [Check required permissions.](#)
6. [Specify a restore reason.](#)
7. [Finish working with the wizard.](#)

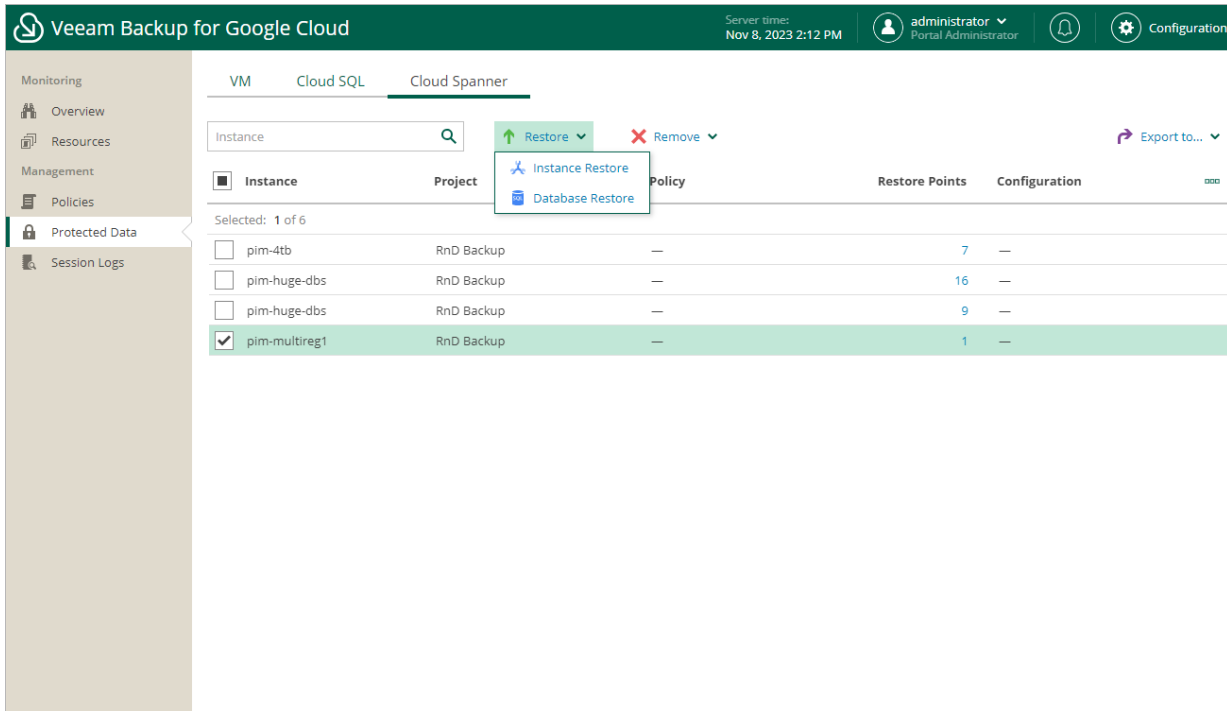
IMPORTANT

Before you start Cloud Spanner database restore, make sure that network settings are configured for each region where worker instances will be deployed during the restore process. For information on how to configure network settings, see [Adding Worker Configurations](#).

Step 1. Launch Database Restore Wizard

To launch the **Cloud Spanner Database Restore** wizard, do the following:

1. Navigate to **Protected Data > Cloud Spanner**.
2. Select the Cloud Spanner instance whose databases you want to restore, and click **Restore > Database Restore**.



The screenshot shows the Veeam Backup for Google Cloud interface. The top navigation bar includes the Veeam logo, the title "Veeam Backup for Google Cloud", the server time "Nov 8, 2023 2:12 PM", and the user "administrator Portal Administrator". The left sidebar shows the navigation menu with "Protected Data" selected. The main content area is titled "Cloud Spanner" and displays a table of instances. The "pim-multireg1" instance is selected, and the "Database Restore" option is highlighted in the "Restore" dropdown menu.

Instance	Project	Policy	Restore Points	Configuration
<input type="checkbox"/> pim-4tb	RnD Backup	—	7	—
<input type="checkbox"/> pim-huge-dbs	RnD Backup	—	16	—
<input type="checkbox"/> pim-huge-dbs	RnD Backup	—	9	—
<input checked="" type="checkbox"/> pim-multireg1	RnD Backup	—	1	—

Step 2. Select Databases

At the **Databases** step of the wizard, click **Add** to select databases to restore, and then choose a restore point that will be used to restore the selected databases. By default, Veeam Backup for Google Cloud uses the most recent valid restore point. However, you can restore the backed-up data to an earlier state.

To select a restore point, do the following:

1. Select a database.
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 Google Cloud provides the following information on each available restore point:

- **Creation Time** – the date when the restore point was created.
- **Destination** – the type of the restore point:
 - § *Snapshot* – a cloud-native snapshot 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 health state of the restore point (the result of the latest health check; applies only to image-level backups).
- **Storage Class** – the storage class of the backup repository where the restore point is stored (applies only to image-level backups).
- **Region** – a region in which the protected Cloud Spanner instance resides.
- **Policy** – a backup policy that created the restore point.

- o **Retention** – a retention configured for the backup policy that created the restore point.

The screenshot shows the Veeam Backup for Google Cloud interface. The top navigation bar includes the Veeam logo, the text 'Veeam Backup for Google Cloud', the server time 'Nov 8, 2023 1:23 PM', and user information 'administrator Portal Administrator'. The main window title is 'Cloud Spanner Database Restore'. On the left, a sidebar contains navigation options: Databases, Project, Instance, Permissions, Reason, and Summary. The 'Databases' section is active, showing a search bar and a table with columns 'Database' and 'Restore Point'. Two databases are listed: 'pim-postgre-db' (selected) and 'pim-spanner-db'. A modal dialog titled 'Choose restore point' is open on the right, displaying a table with columns 'Creation Time', 'Destination', 'State', 'Storage Class', and 'Policy'. One restore point is visible with a creation time of '09/15/2023 7:3...'. At the bottom of the dialog are 'Apply' and 'Cancel' buttons.

Creation Time ↑	Destination	State	Storage Class	Policy
09/15/2023 7:3...	Snapshot	—	—	—

Step 3. Select Project

At the **Project** step of the wizard, select a project that manages a Cloud Spanner instance to which you want to restore the selected databases and specify a service account whose permissions will be used to perform the restore operation. For more information on the required permissions, see [Service Account Permissions](#).

For a project to be displayed in the **Project** drop-down list, it must be added to Veeam Backup for Google Cloud as described in section [Adding Projects and Folders](#). If you have not added the necessary project to Veeam Backup for Google Cloud beforehand, you can do it without closing the **Cloud Spanner Database Restore** wizard. To add a project, click **Add** and complete the **Add Projects and Folders** wizard.

For a service account to be displayed in the list of available accounts, it must be added to Veeam Backup for Google Cloud as described in section [Adding Service Accounts](#), and must be assigned the *Cloud Spanner Instances Restore* operational role as described in section [Adding Projects and Folders](#).

The screenshot shows the 'Specify project' step of the 'Cloud Spanner Database Restore' wizard. The interface includes a top navigation bar with the Veeam logo, 'Veeam Backup for Google Cloud', server time 'Nov 8, 2023 1:25 PM', and user information 'administrator Portal Administrator'. A left sidebar contains navigation links: 'Databases', 'Project' (highlighted), 'Instance', 'Permissions', 'Reason', and 'Summary'. The main content area is titled 'Specify project' and contains the following text: 'Choose a project managing the Cloud Spanner instance to which the selected databases will be restored, and specify a service account that will be used to access the project.' Below this, there are two sections: 'Project' and 'Service account'. The 'Project' section has a dropdown menu with 'RnD Backup (rnd-backup-254612)' selected and a '+ Add' button. The 'Service account' section has the text 'Specify a service account.' and a service account entry: 'Service account: veeam-1649186685-sa@rnd-backup-2.iam.gserviceaccount.com'. At the bottom, there are three buttons: 'Previous', 'Next', and 'Cancel'.

Step 4. Configure Target Instance Settings

At the **Instance** step of the wizard, choose a Cloud Spanner instance that will host the restored databases. To do that, click the link in the **Instance** field, select the necessary Cloud Spanner instance from the **Choose Cloud Spanner instance** list, and click **Apply**. For a Cloud Spanner instance to be displayed in the list of available instances, it must belong to the selected project and be running on a supported database engine.

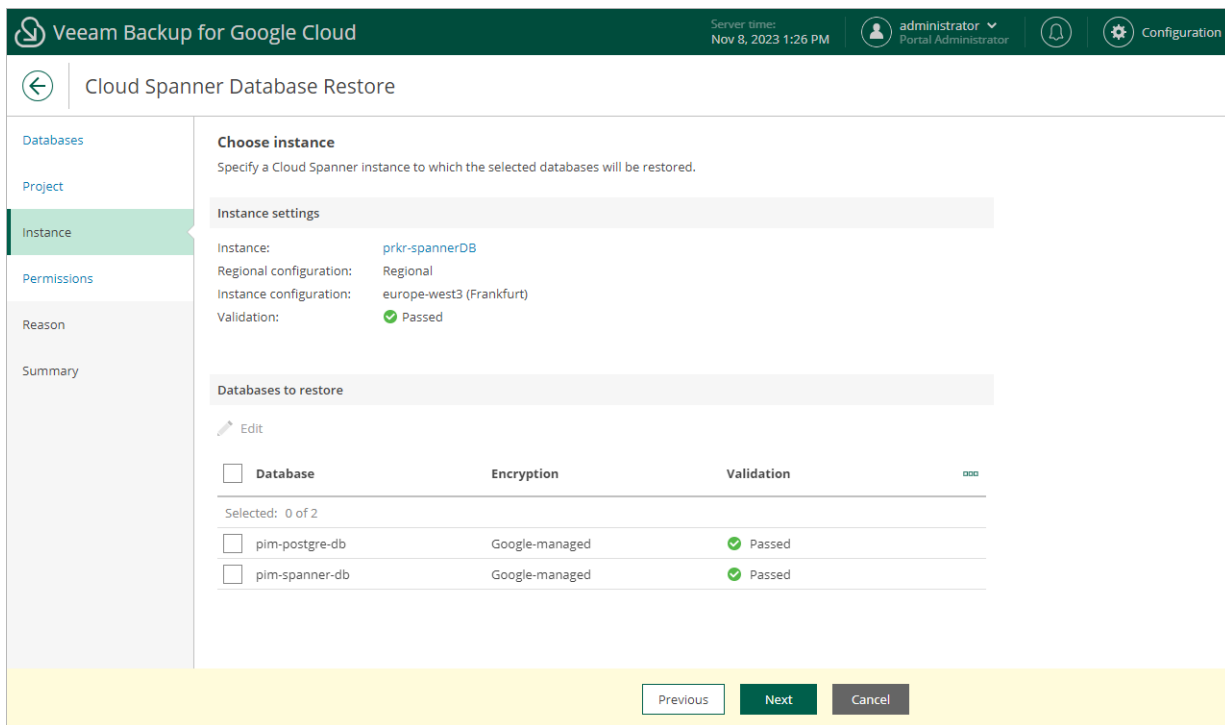
You can also specify new names and choose new encryption schemes for the restored databases. To do that, select a database and click **Edit**.

TIP

Veeam Backup for Google Cloud will perform a number of configuration checks for the selected instance and databases:

- If any of the checks fail to complete successfully for an instance, the wizard will display an error in the **Validation** field.
- If any of the checks fail to complete successfully for a database, the wizard will display an error in the **Validation** column of the **Databases to restore** table.

You can click the link to get more information on an error.



The screenshot shows the 'Cloud Spanner Database Restore' wizard in the 'Instance' step. The interface includes a top navigation bar with the Veeam logo, server time (Nov 8, 2023 1:26 PM), user information (administrator, Portal Administrator), and a 'Configuration' link. A left sidebar contains navigation options: Databases, Project, Instance (selected), Permissions, Reason, and Summary. The main content area is titled 'Choose instance' and includes instructions to specify a Cloud Spanner instance. Below this, 'Instance settings' are displayed: Instance (prkr-spannerDB), Regional configuration (Regional), Instance configuration (europe-west3 (Frankfurt)), and Validation (Passed). A 'Databases to restore' section features an 'Edit' link and a table with columns for Database, Encryption, and Validation. The table lists two databases: pim-postgre-db and pim-spanner-db, both with Google-managed encryption and Passed validation. A 'Selected: 0 of 2' indicator is shown above the table. At the bottom, there are 'Previous', 'Next', and 'Cancel' buttons.

Database	Encryption	Validation
<input type="checkbox"/> pim-postgre-db	Google-managed	✓ Passed
<input type="checkbox"/> pim-spanner-db	Google-managed	✓ Passed

Step 5. Check Required Permissions

At the **Permissions** step of the wizard, Veeam Backup for Google Cloud will verify whether the specified service account has all the necessary permissions required to perform data recovery tasks for the selected project. For more information on the required permissions, see [Service Account Permissions](#).

To see the list of missing permissions that must be granted to the service account in order to perform an operation, click the link in the **Details** column. You can grant the missing permissions to the service account [using the Google Cloud console](#) or instruct Veeam Backup for Google Cloud to do it:

- To grant the missing permissions manually, click **Download Script**. Veeam Backup for Google Cloud will generate a gcloud script that you can run in the Google Cloud console to assign all the necessary permissions to the service account.

The account under which you run the script must have the permissions required both to get and set project IAM policies and to create custom IAM roles (for example, it can have the *iam.securityAdmin* and *iam.roleAdmin* roles assigned). To learn what permissions and roles are required to create custom roles in IAM, see [Google Cloud documentation](#).

- To let Veeam Backup for Google Cloud grant the missing permissions automatically, click **Grant** and then click **Sign in with Google** in the **Grant permissions** window. You will be redirected to the OAuth consent screen authorization page. Sign in using credentials of a Google account that will be used to grant the permissions.

The account under which you sign in to Google Cloud must have the permissions required both to get and set project IAM policies and to create custom IAM roles (for example, it can have the *iam.securityAdmin* and *iam.roleAdmin* roles assigned). To learn what permissions and roles are required to create service account, see [Google Cloud documentation](#).

NOTE

For Veeam Backup for Google Cloud to be able to authorize in Google Cloud, the OAuth consent screen must be configured as described in section [Registering Applications](#). Note that Veeam Backup for Google Cloud does not store in the configuration database the provided Google account credentials and access tokens received during authorization.

To make sure that the missing permissions have been successfully granted, click **Recheck**.

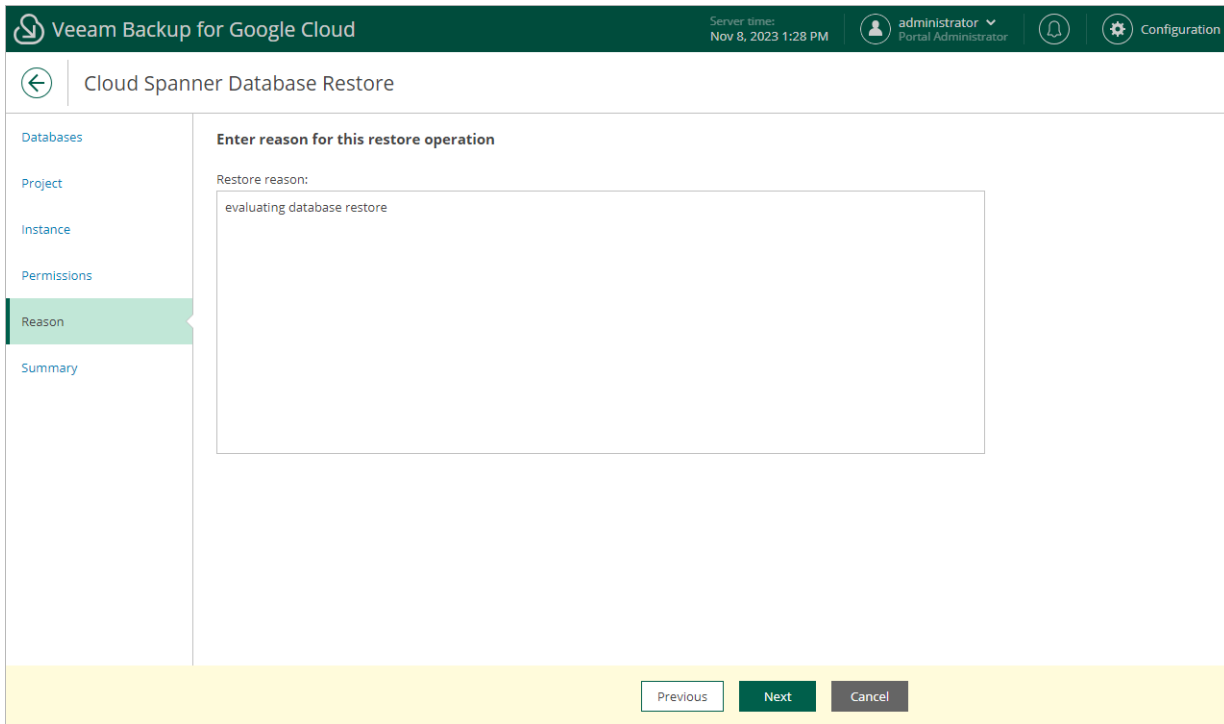
The screenshot shows the Veeam Backup for Google Cloud interface. The top navigation bar includes the Veeam logo, the product name, the server time (Nov 8, 2023 1:27 PM), the user (administrator, Portal Administrator), and a Configuration icon. The main content area is titled 'Cloud Spanner Database Restore'. On the left, there is a sidebar with navigation options: Databases, Project, Instance, Permissions (highlighted), Reason, and Summary. The main area displays a 'Check permissions' section with the instruction 'Verify whether all the required permissions are granted.' Below this are three buttons: 'Recheck' (with a refresh icon), 'Download Script' (with a download icon), and 'Grant' (with a plus icon). A table below shows the results of the check:

Check	Result	Details
Cloud Spanner Restore	Passed	All the required permissions are gra...

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

Step 6. Specify Restore Reason

At the **Reason** step of the wizard, specify a reason for restoring the Cloud Spanner databases. This information will be saved to the session history, and you will be able to reference it later.



The screenshot displays the Veeam Backup for Google Cloud interface. The top navigation bar includes the Veeam logo, the product name, the server time (Nov 8, 2023 1:28 PM), the user (administrator, Portal Administrator), and a Configuration icon. The main content area is titled 'Cloud Spanner Database Restore' and features a left-hand navigation menu with options: Databases, Project, Instance, Permissions, Reason (highlighted), and Summary. The main panel is titled 'Enter reason for this restore operation' and contains a text input field labeled 'Restore reason:' with the text 'evaluating database restore' entered. At the bottom of the interface, there are three buttons: 'Previous', 'Next', and 'Cancel'.

Step 7. Finish Working with Wizard

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

The screenshot shows the 'Cloud Spanner Database Restore' wizard in the 'Summary' step. The interface includes a top navigation bar with the Veeam logo, server time (Nov 8, 2023 1:28 PM), and user information (administrator, Portal Administrator). A left sidebar lists navigation options: Databases, Project, Instance, Permissions, Reason, and Summary (which is highlighted). The main content area is titled 'Review configured settings' and contains a 'Copy to Clipboard' button. It is divided into three sections: 'General settings' (Project name: RnD Backup, Project ID: rnd-backup-254612, Service account: veeam-1649186685-sa@rnd-backup-2.iam.gserviceaccount.com, Reason: evaluating database restore), 'Databases to restore' (Instance: prkr-spannerDB, Regional configuration: regional, Instance configuration: europe-west3 (Frankfurt), Databases: 2 databases), and 'Validation' (Instance settings: Passed, Database settings: Passed, Permission check: Passed). At the bottom, there are three buttons: 'Previous', 'Finish', and 'Cancel'.

Review configured settings

[Copy to Clipboard](#)

General settings

Project name: RnD Backup
Project ID: rnd-backup-254612
Service account: veeam-1649186685-sa@rnd-backup-2.iam.gserviceaccount.com
Reason: evaluating database restore

Databases to restore

Instance: [prkr-spannerDB](#)
Regional configuration: regional
Instance configuration: europe-west3 (Frankfurt)
Databases: [2 databases](#)

Validation

Instance settings: ✔ Passed
Database settings: ✔ Passed
Permission check: ✔ Passed

[Previous](#) [Finish](#) [Cancel](#)

Instant Recovery

Veeam Backup & Replication allows you to use the Instant Recovery feature to restore VM 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

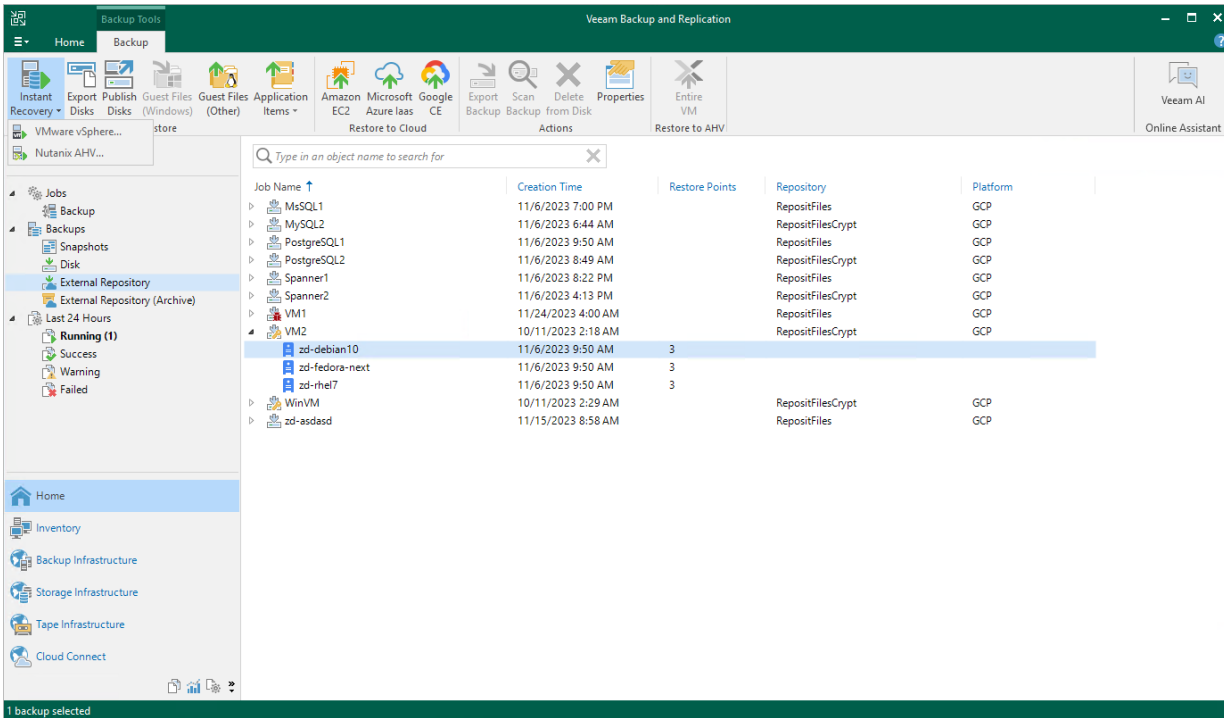
Instant Recovery can be performed only using backup files stored in backup repositories for which you have specified HMAC keys associated with the service accounts that are used to access the repositories. To learn how to specify credentials for repositories, see sections [Creating New Repositories](#) and [Connecting to Existing Appliances](#).

Before you start the recovery operation, make sure that you have added to the backup infrastructure a vCenter Server, a Microsoft Hyper-V server or a Nutanix AHV cluster that will manage the restored VM 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 a VM instance that you want to recover, select the necessary 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 disks of VM instances from image-level backups created by Veeam Backup for Google Cloud and to convert them to the VMDK, VHD or VHDX format. 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 User Guide, section [Disk Export](#).

IMPORTANT

Disk export can be performed only using backup files stored in backup repositories for which you have specified HMAC keys associated with the service accounts that are used to access the repositories. To learn how to specify credentials for repositories, see sections [Creating New Repositories](#) and [Connecting to Existing Appliances](#).

To export disks of a VM instance, 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 a VM instance whose disks 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 disks of backed-up VM 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

Disk publishing can be performed only using backup files stored in backup repositories for which you have specified HMAC keys associated with the service accounts that are used to access the repositories. To learn how to specify credentials for repositories, see sections [Creating New Repositories](#) and [Connecting to Existing Appliances](#).

To publish disks of a VM 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 VM instance whose disks 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 AWS

Veeam Backup & Replication allows you to restore VM instances from image-level backups created with Veeam Backup for Google Cloud to AWS as EC2 instances. You can restore VM instances to any available restore point. For more information, see the Veeam Backup & Replication User Guide, section [Restore to Amazon EC2](#).

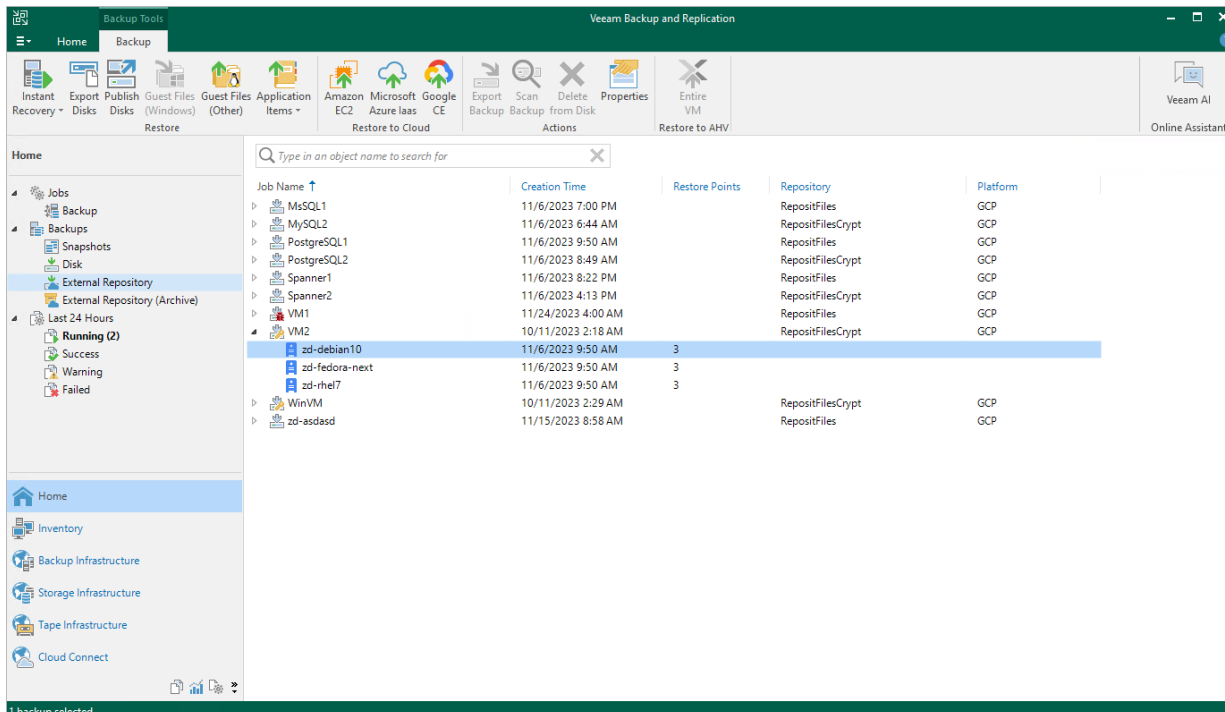
IMPORTANT

Restore to AWS can be performed only using backup files stored in backup repositories for which you have specified HMAC keys associated with the service accounts that are used to access the repositories. To learn how to specify credentials for 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 a VM instance to Amazon EC2, 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 a VM instance that you want to restore, select the necessary instance and click **Amazon EC2** on the ribbon.
4. Complete the **Restore to Amazon EC2** wizard as described in the Veeam Backup & Replication User Guide, section [Restoring to Amazon EC2](#).



Restoring to Microsoft Azure

Veeam Backup & Replication allows you to restore VM instances from image-level backups created with Veeam Backup for Google Cloud to Microsoft Azure as Azure VMs. You can restore VM instances to any available restore point. For more information, see the Veeam Backup & Replication User Guide, section [Restore to Microsoft Azure](#).

IMPORTANT

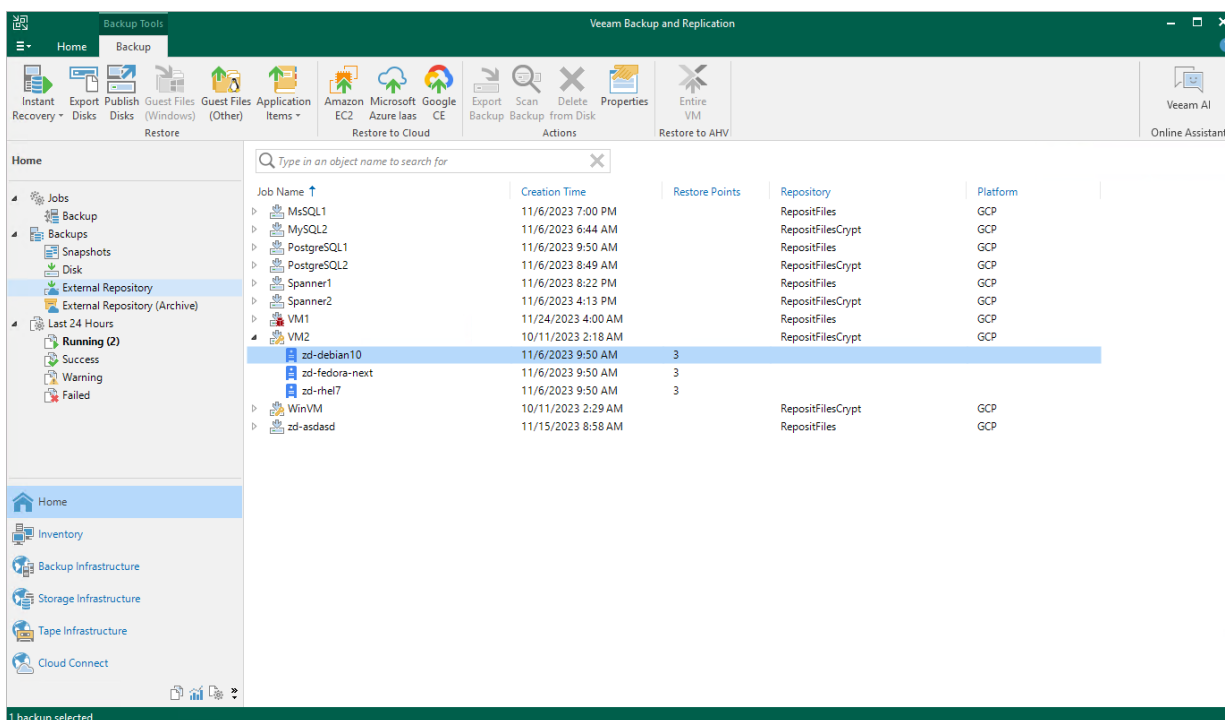
Restore to Microsoft Azure can be performed only using backup files stored in backup repositories for which you have specified HMAC keys associated with the service accounts that are used to access the repositories. To learn how to specify credentials for repositories, see sections [Creating New Repositories](#) and [Connecting to Existing Appliances](#).

Before you start the restore operation:

- Configure the initial settings of an Azure account or Azure Stack account as described in the Veeam Backup & Replication User Guide, section [Configuring Initial Settings](#).
- Check the limitations and prerequisites described in the Veeam Backup & Replication User Guide, section [Before You Begin](#).

To restore a VM instance to Microsoft Azure, 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 a VM 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 Nutanix AHV

Veeam Backup & Replication allows you to restore VM instances from image-level backups created with Veeam Backup for Google Cloud to Nutanix AHV as Nutanix AHV VMs. You can restore VM 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 backup repositories for which you have specified HMAC keys associated with the service accounts that are used to access the repositories. To learn how to specify credentials for repositories, see sections [Creating New Repositories](#) and [Connecting to Existing Appliances](#).

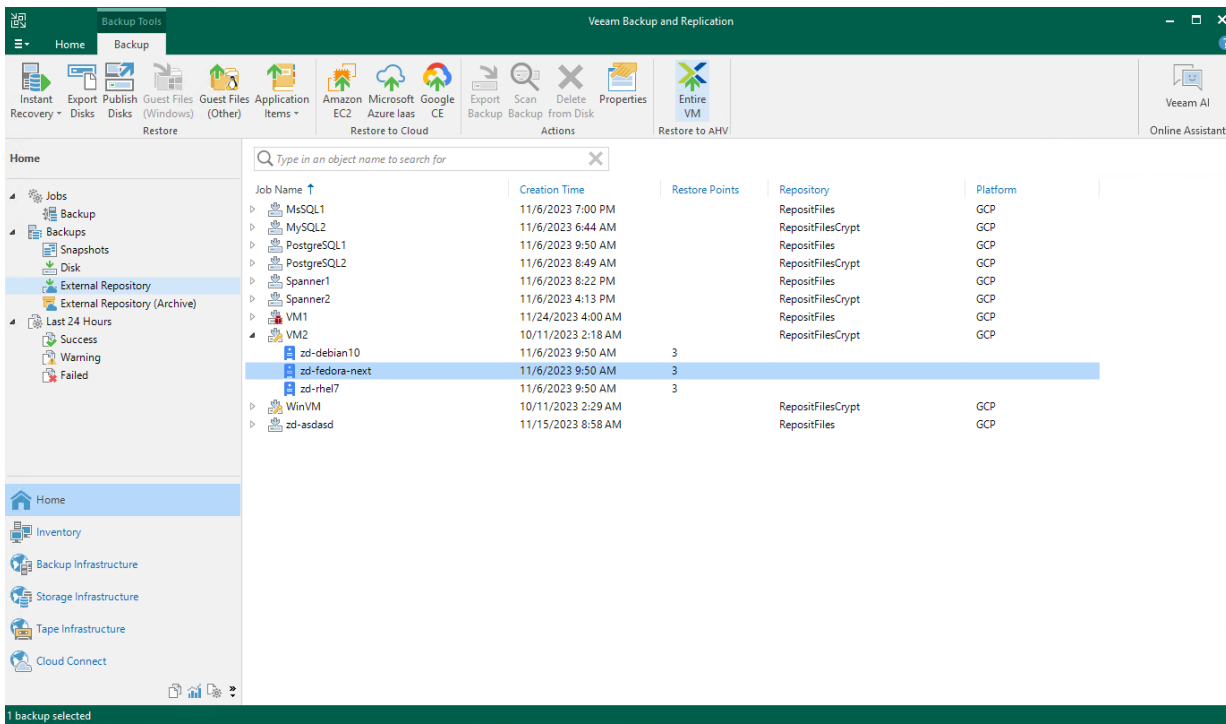
Before you start the restore operation:

- Configure the backup infrastructure as described in the Veeam Backup for Nutanix AHV User Guide, section [Deployment](#).
- If you restore VM instances from standard backups, make sure that these backups 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 VM instances from backups copied to the archive access tier of a [scale-out backup repository](#), make sure that you have retrieved these backups from archive as described in the Veeam Backup & Replication User Guide, section [Retrieving Backup Files](#).

To restore a VM 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 backup policy that protects a VM instance that you want to restore, select the necessary instance and click **Entire VM** on the ribbon.

- 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 Google Cloud comes with an **Overview** dashboard that provides at-a-glance real-time overview of the protected Google Cloud 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 all sessions started for data protection and disaster recovery operations (including system sessions) that completed successfully during the past 24 hours, 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** tab will show only those sessions that have the same status as that clicked in the widget.

For more information on the **Session Logs** tab, see [Viewing Session Statistics](#).

- **Successful Policy Tasks** – displays the number of snapshots, backups and archived backups successfully created by backup policies during a specific time period (the past 24 hours by default), and the number of attempts that were made to create these restore points.

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** tab will show only those sessions during which Veeam Backup for Google Cloud created the same items as that clicked in the widget.

For more information on the **Session Logs** tab, see [Viewing Session Statistics](#).

- **Top Policies** – shows top 8 backup policies for fluctuations in execution time (including retries). For each policy, the widget calculates the growth rate to detect whether it took less or more time for the policy to complete in comparison with the average runtime value for the previous 10 policy launches.
- **Protected Workloads** – displays the number of available Google Cloud resources that got protected by Veeam Backup for Google Cloud 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 backups and archived backups created by Veeam Backup for Google Cloud in storage buckets, the number of snapshots created for the protected resources, and the total size of all VM instance snapshots residing in Google Cloud Storage. The widget also calculates the ratio of the total amount of storage space used in the *Standard* and *Nearline* storage classes to the total amount of storage space used in the *Archive* storage class.
- **Bottlenecks Overview** – is designed to help you avoid possible backup bottlenecks.

The widget analyzes the total amount of time waited to deploy worker instances during data protection operations in different Google Cloud regions, and displays the most problematic region (if any).

The widget also analyzes the amount of disk quota across all regions to detect whether the quota has already been reached in any of the regions, and whether Veeam Backup for Google Cloud failed to deploy a worker instance with the primary profile in that region during a backup or restore process. For more information on machine types of VM instances that operate as worker instances, see [Managing Worker Profiles](#).

The widget also 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, the only way to resolve the issue may be to change the machine type for the backup appliance as described in [Google Cloud documentation](#).

To learn how to resolve a bottleneck, click the **How to resolve?** link in the widget row.

The screenshot displays the Veeam Backup for Google Cloud dashboard with the following sections:

- Sessions in Last 24 Hours:**
 - Failed: 42 ↓
 - Warning: 47 ↓
 - Success: 313 ↑
 - Running now: 0
- Successful Policy Tasks (Last 24 hours):**
 - Snapshots: 169 of 169 (100%)
 - Backups: 163 of 163 (100%)
 - Archives: No data
- Protected Workloads (Last 24 hours):**
 - Virtual machines: 4 of 22 (18%)
 - Cloud SQL instances: 3 of 9 (33%)
 - Cloud Spanner instances: 3 of 11 (27%)
- Storage Usage:**
 - Snapshots: 62 GB (1755 Total)
 - Backups: 21 GB
 - Archives: 5 GB
 - Total: 26 GB
 - Legend: Standard: 21 GB, Nearline: 0 Bytes, Archive: 5 GB
- Top Policies (By duration increase):**

Policy	Duration	Start Time	Percentage
PostgreSQL1	23 min 6 sec	11/06 05:49 PM	+586%
Spanner1	25 min 53 sec	11/16 08:00 AM	+44%
Spanner2	19 min 34 sec	11/16 08:00 AM	+30%
VM2	12 min 30 sec	11/06 05:50 PM	+9%
MsSQL1	22 min 41 sec	11/16 09:00 AM	+0%
zd-asdasd	11 min 45 sec	11/15 04:58 PM	—
- Bottlenecks Overview:**
 - Total workers wait time: 6 min (average: 6 min) on 11/16/2023 8:26 AM at europe-west1 (Belgium). Includes a "How to resolve?" link.
 - Region disk quota: Available
 - Appliance memory sizing: OK

Viewing Session Statistics

For each performed data protection or disaster recovery operation, Veeam Backup for Google Cloud starts a new session and stores its records in the configuration database.

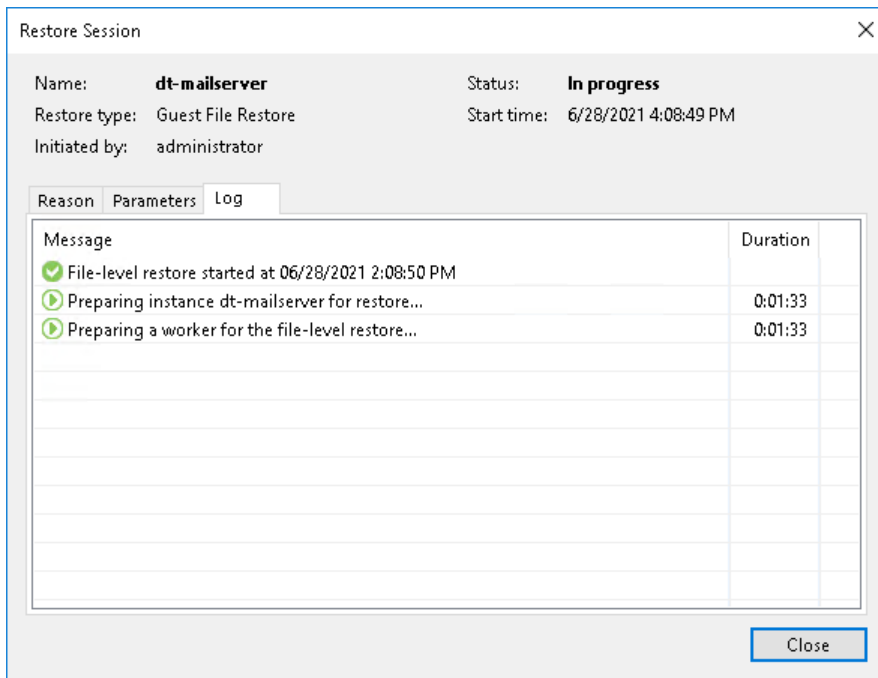
Viewing Session Statistics Using Console

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 the Veeam Backup & Replication User Guide, sections [Viewing Real-Time Statistics](#) and [Viewing Job Session Results](#).

Veeam Backup & Replication also allows you track statistics of most data recovery operations initiated from Veeam Backup for Google Cloud. 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 session.
Alternatively, select the session and click **Statistics** on the ribbon.
- In the Veeam Backup & Replication console, open the **History** view and navigate to **Jobs** or **Restore**. In the working area, double-click the necessary session.
Alternatively, select the session and click **Statistics** on the ribbon.

The opened window will display restore session details such as the name of the Google Cloud resource whose data is being processed, the account under which the session has started, the session status and duration, information on the restore point selected for the 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** tab. To view the full list of tasks executed during an operation, click the link in the **Status** column. To view the full list of Google Cloud resources processed during an operation, click the link in the **Items** column.

TIP

If you want to specify the time period during which Veeam Backup for Google Cloud will keep session records in the configuration database, follow the instructions provided in section [Configuring Global Retention Settings](#).

Veeam Backup for Google Cloud

Server time: May 5, 2022 5:40 PM | wendy_may Administrator | Configuration

Monitoring: Overview, Resources, Management: Policies, Protected Data, Session Logs

Policy: Filtering options (None) All Time

Stop

Type

- Rescan Cloud SQL Instances
- Rescan Prices
- Rescan Repositories
- Rescan Crypto Keys
- Rescan VM Snapshots
- Rescan Networks
- Create Repository
- Rescan VM Instances
- Rescan Regions
- Rescan Crypto Keys
- Rescan Regions
- File-Level Recovery
- File-Level Recovery
- File-Level Recovery
- Rescan Policy Repository
- Cloud SQL Health Check
- Retention
- Retention

File Level Recovery: dr-instance-2-sharedvpc

Session Status

Result	Start Time ↓	End Time	Duration
Failed	05/03/2022 3:13:22 PM	05/03/2022 3:18:59 PM	5 min 37 sec

Session Log

Start Time	Status	Description	Execution Duration
05/03/2022 3:13:23 PM	Success	File-level restore started	—
05/03/2022 3:13:23 PM	Failed	Preparing instance dr-instance-2-sharedvpc for restore...	5 min 36 sec
05/03/2022 3:13:23 PM	Failed	Preparing a worker for the file-level restore...	5 min 36 sec
05/03/2022 3:18:59 PM	Failed	Unable to perform the operation. Unexpected exception: One or more errors occurred. (Failed to perform the operation. Storage does not exist. Failed to restore file from local backup. VFS link: [summary.xml], Target file: [MemFsc://RestoreText_(6ee9004e-3462-425c-a3e8-79eeab31b540)]. CHMOD mask: [0]. Agent failed to process method (DataTransfer.RestoreText). Agent failed to process method (DataTransfer.RestoreText))	—
05/03/2022 3:18:59 PM	Failed	The task has completed	—

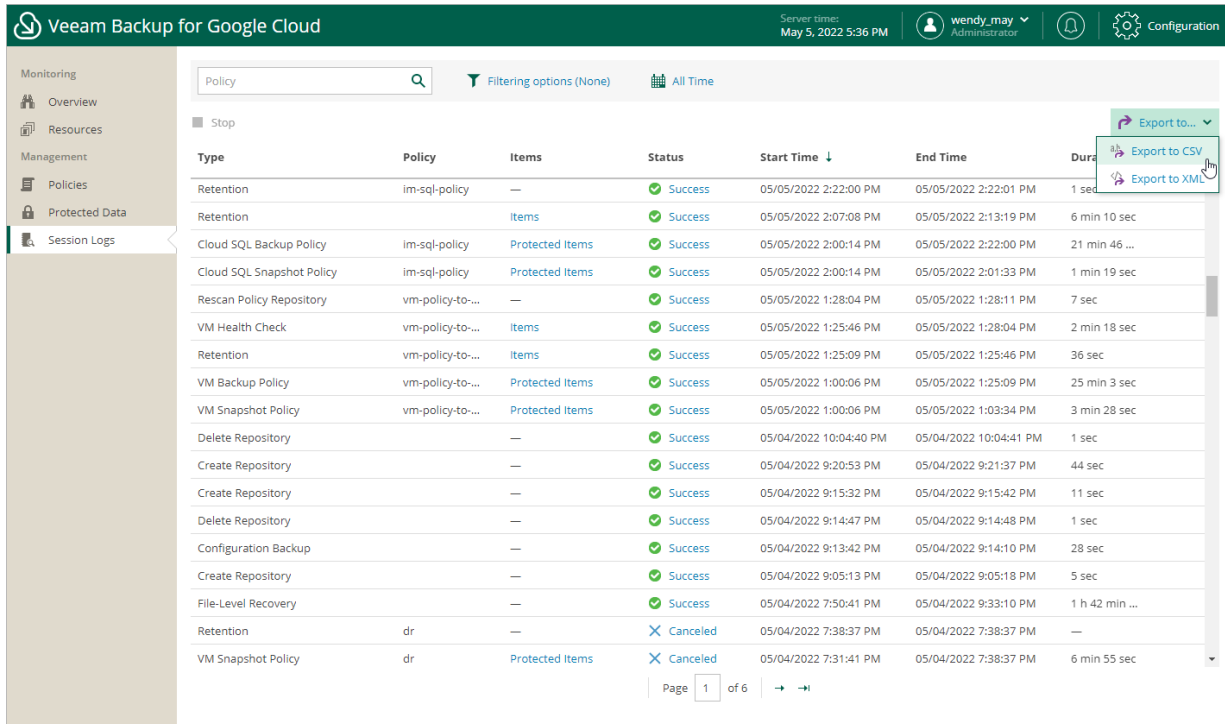
Close

Collecting Object Properties

You can export properties of objects managed by Veeam Backup for Google Cloud as a single .CSV or .XML file. To do that, navigate to the necessary tab and click **Export**. Veeam Backup for Google Cloud 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 Google Cloud will still export all properties of all objects present on the currently opened tab.



The screenshot displays the Veeam Backup for Google Cloud interface. The top navigation bar includes the Veeam logo, the text "Veeam Backup for Google Cloud", the server time "May 5, 2022 5:36 PM", the user name "wendy_may Administrator", and a "Configuration" icon. The left sidebar contains navigation options: Monitoring (Overview, Resources), Management (Policies, Protected Data), and Session Logs. The main content area shows a search bar for "Policy", filtering options, and a table of policies. A dropdown menu is open over the table, showing "Export to...", "Export to CSV", and "Export to XML".

Type	Policy	Items	Status	Start Time ↓	End Time	Duration
Retention	im-sql-policy	—	Success	05/05/2022 2:22:00 PM	05/05/2022 2:22:01 PM	1 sec
Retention		Items	Success	05/05/2022 2:07:08 PM	05/05/2022 2:13:19 PM	6 min 10 sec
Cloud SQL Backup Policy	im-sql-policy	Protected Items	Success	05/05/2022 2:00:14 PM	05/05/2022 2:22:00 PM	21 min 46 ...
Cloud SQL Snapshot Policy	im-sql-policy	Protected Items	Success	05/05/2022 2:00:14 PM	05/05/2022 2:01:33 PM	1 min 19 sec
Rescan Policy Repository	vm-policy-to...	—	Success	05/05/2022 1:28:04 PM	05/05/2022 1:28:11 PM	7 sec
VM Health Check	vm-policy-to...	Items	Success	05/05/2022 1:25:46 PM	05/05/2022 1:28:04 PM	2 min 18 sec
Retention	vm-policy-to...	Items	Success	05/05/2022 1:25:09 PM	05/05/2022 1:25:46 PM	36 sec
VM Backup Policy	vm-policy-to...	Protected Items	Success	05/05/2022 1:00:06 PM	05/05/2022 1:25:09 PM	25 min 3 sec
VM Snapshot Policy	vm-policy-to...	Protected Items	Success	05/05/2022 1:00:06 PM	05/05/2022 1:03:34 PM	3 min 28 sec
Delete Repository		—	Success	05/04/2022 10:04:40 PM	05/04/2022 10:04:41 PM	1 sec
Create Repository		—	Success	05/04/2022 9:20:53 PM	05/04/2022 9:21:37 PM	44 sec
Create Repository		—	Success	05/04/2022 9:15:32 PM	05/04/2022 9:15:42 PM	11 sec
Delete Repository		—	Success	05/04/2022 9:14:47 PM	05/04/2022 9:14:48 PM	1 sec
Configuration Backup		—	Success	05/04/2022 9:13:42 PM	05/04/2022 9:14:10 PM	28 sec
Create Repository		—	Success	05/04/2022 9:05:13 PM	05/04/2022 9:05:18 PM	5 sec
File-Level Recovery		—	Success	05/04/2022 7:50:41 PM	05/04/2022 9:33:10 PM	1 h 42 min ...
Retention	dr	—	Canceled	05/04/2022 7:38:37 PM	05/04/2022 7:38:37 PM	—
VM Snapshot Policy	dr	Protected Items	Canceled	05/04/2022 7:31:41 PM	05/04/2022 7:38:37 PM	6 min 55 sec

Page 1 of 6

Updating Veeam Backup for Google Cloud

Veeam Backup for Google Cloud 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 5.0, you can upgrade backup appliances only in the Veeam Backup & Replication console. To perform upgrade of Veeam Backup for Google Cloud to version 5.0, the backup appliance must be running version 2.0 or later. To upgrade from earlier versions, you must first perform update to Veeam Backup for Google Cloud version 2.0 or later as described in section [Installing Updates](#).

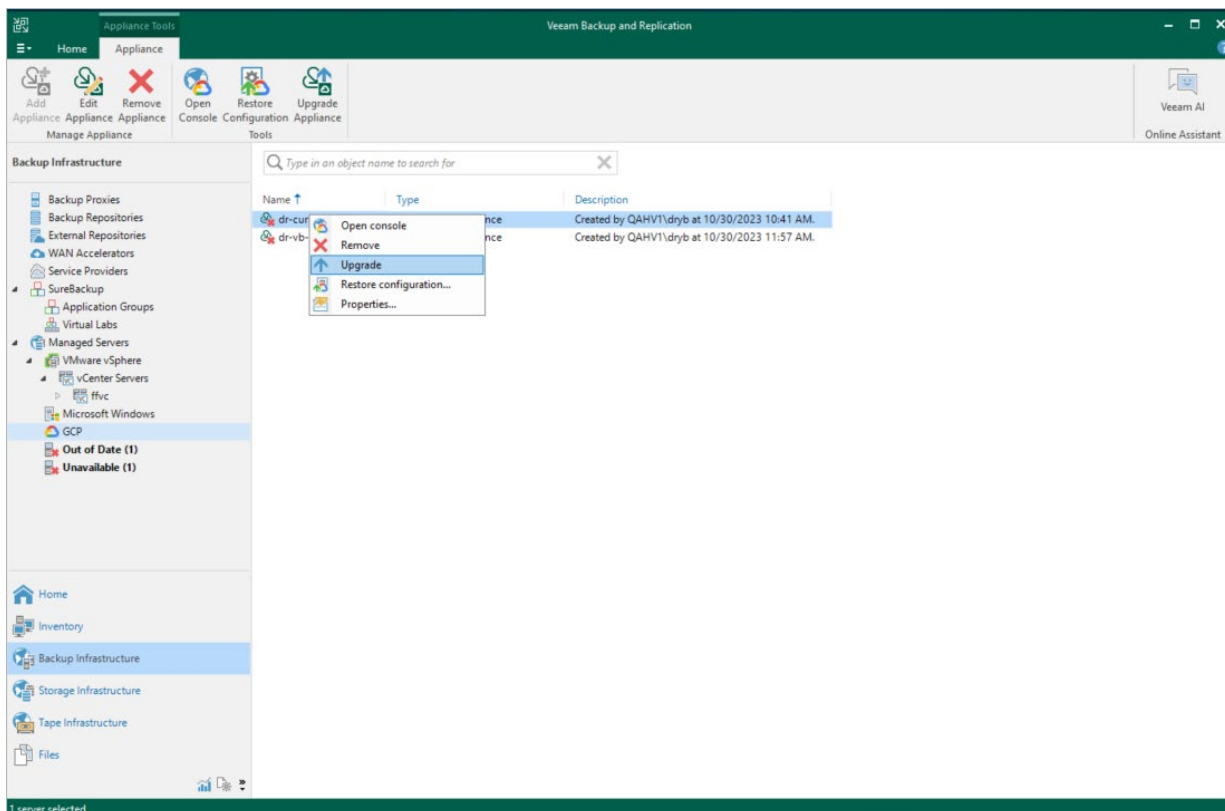
IMPORTANT

Before you upgrade a backup appliance, check whether the Veeam Backup for Google Cloud version is compatible with the current version of Google Cloud Plug-in for Veeam Backup & Replication. For more information, see [System Requirements](#).

Veeam Backup & Replication allows you to download and install new available Veeam Backup for Google Cloud 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**.



Updating Appliances Using Web UI

Veeam Backup for Google Cloud automatically notifies you about newly released package updates available for the operating system running on the backup appliance. However, starting from Veeam Backup for Google Cloud version 5.0, you can use the Veeam Backup for Google Cloud Web UI to install package updates only. To upgrade Veeam Backup for Google Cloud to new versions, follow the instructions provided in section [Upgrading Appliances](#).

Upgrading Appliances

Starting from Veeam Backup for Google Cloud version 5.0, you can upgrade backup appliances only in the Veeam Backup & Replication console. To perform upgrade of Veeam Backup for Google Cloud to version 5.0, the backup appliance must be running version 2.0 or later. To upgrade from earlier versions, you must first perform update to Veeam Backup for Google Cloud version 2.0 or later as described in section [Installing Updates](#).

IMPORTANT

Before you upgrade the backup appliance, make sure that all backup policies are 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 Google Cloud 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 Google Cloud will continue using the license that had been used before you added the Veeam Backup for Google Cloud appliance to the Veeam Backup & Replication infrastructure.

For more information on licensing scenarios, see [Scenarios](#).

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.

IMPORTANT

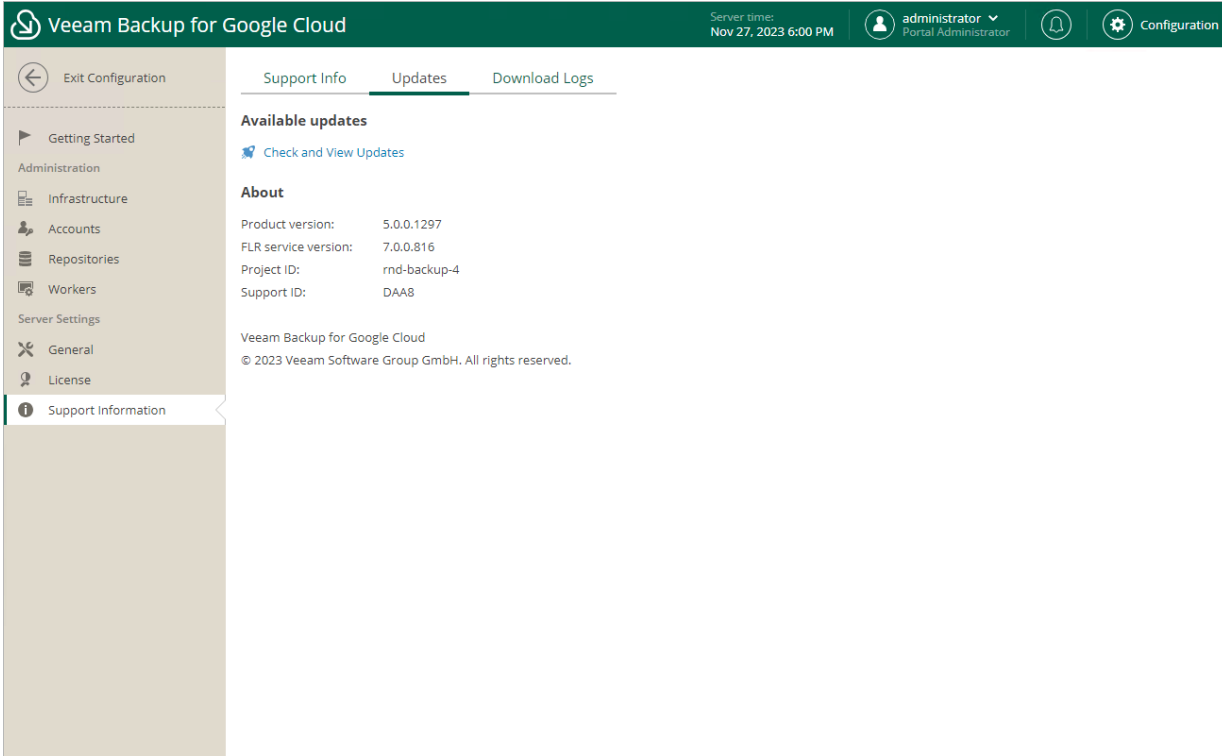
If you remove the backup appliance from the backup infrastructure, you will no longer be able to protect Cloud Spanner resources. For more information, see [Integration with Veeam Backup & Replication](#).

Checking for Updates

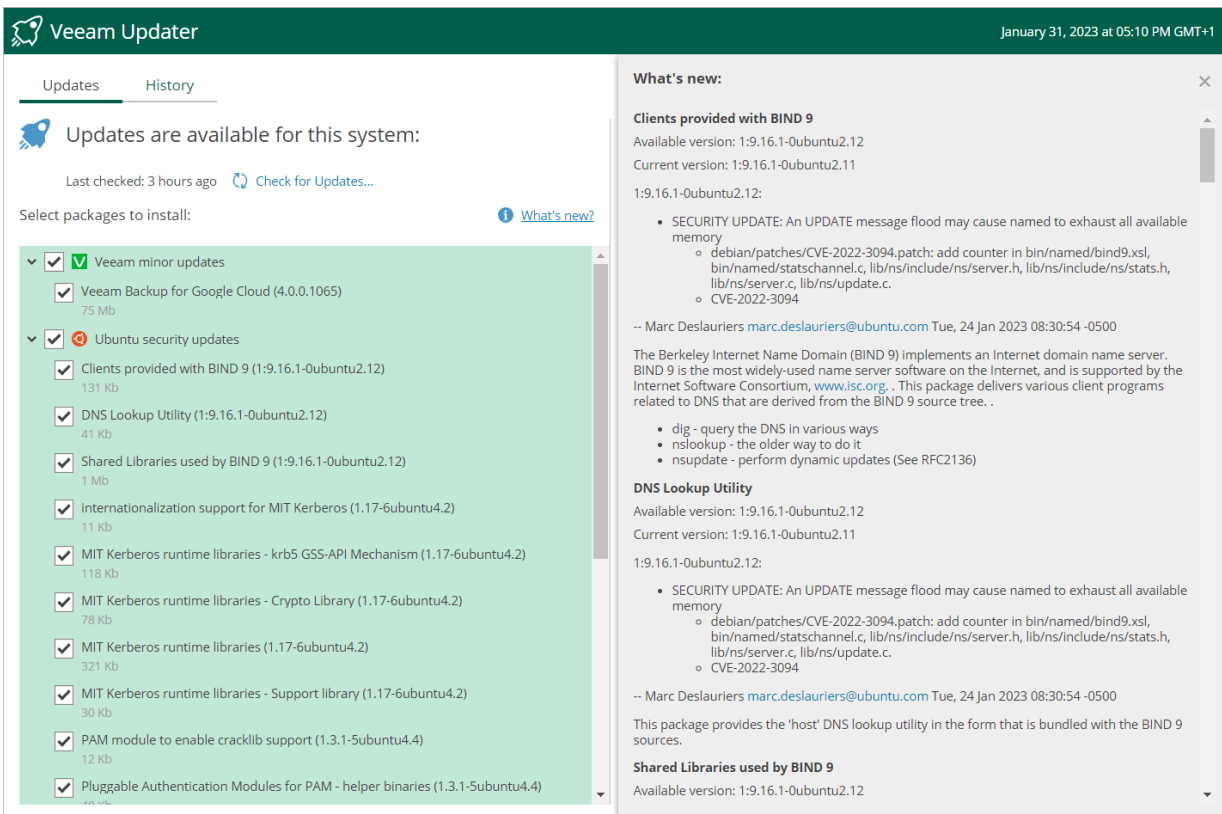
Veeam Backup for Google Cloud automatically notifies you about newly released product versions and package updates available for the operating system running on the backup appliance. However, you can check for the available updates manually if required:

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 Google Cloud will display them on the **Updates** tab of the **Veeam Updater** page. To view detailed information on an update, select the check box next to the update and click **What's new?**



Installing Updates

To download and install new available package updates, 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

- You can update the standalone backup appliance using the Veeam Updater service only. Updating the standalone appliance manually is not supported.
- You can update the backup appliance managed by a Veeam Backup & Replication server from the Veeam Backup & Replication console only, as described in section [Upgrading Appliances Using Console](#). Updating the managed appliance using the Veeam Updater service is not supported.

Installing Updates

IMPORTANT


Before you install a product update, make sure that all backup policies are 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 Google Cloud 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 Google Cloud 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 Google Cloud restarts all services running on the backup appliance, including the Web UI service. That is why Veeam Backup for Google Cloud will log you out when the update process completes.

Scheduling Update Installation

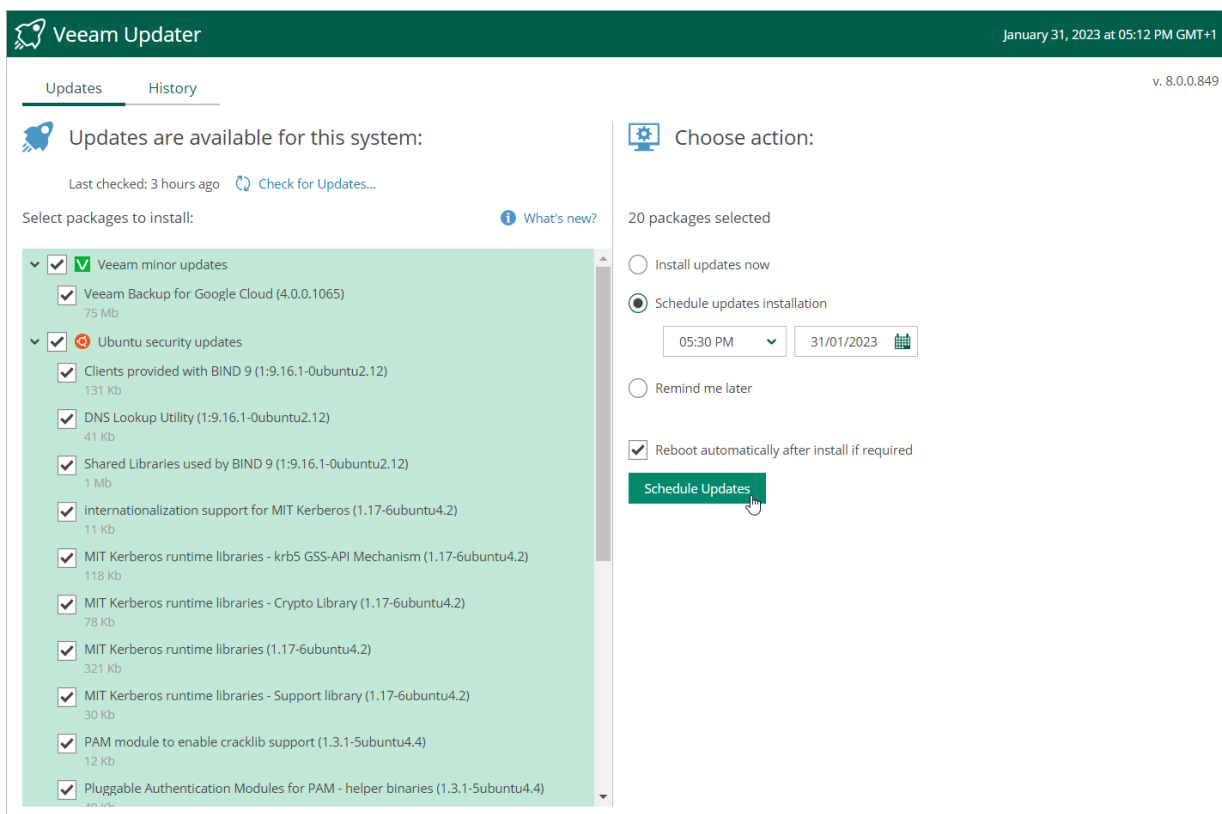
You can instruct Veeam Backup for Google Cloud to automatically download and install available product versions and package updates on a specific date at a specific time:

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 for the update installation, make sure that 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 Google Cloud to reboot the backup appliance if needed.
- c. Click **Schedule Updates**.



Veeam Backup for Google Cloud 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

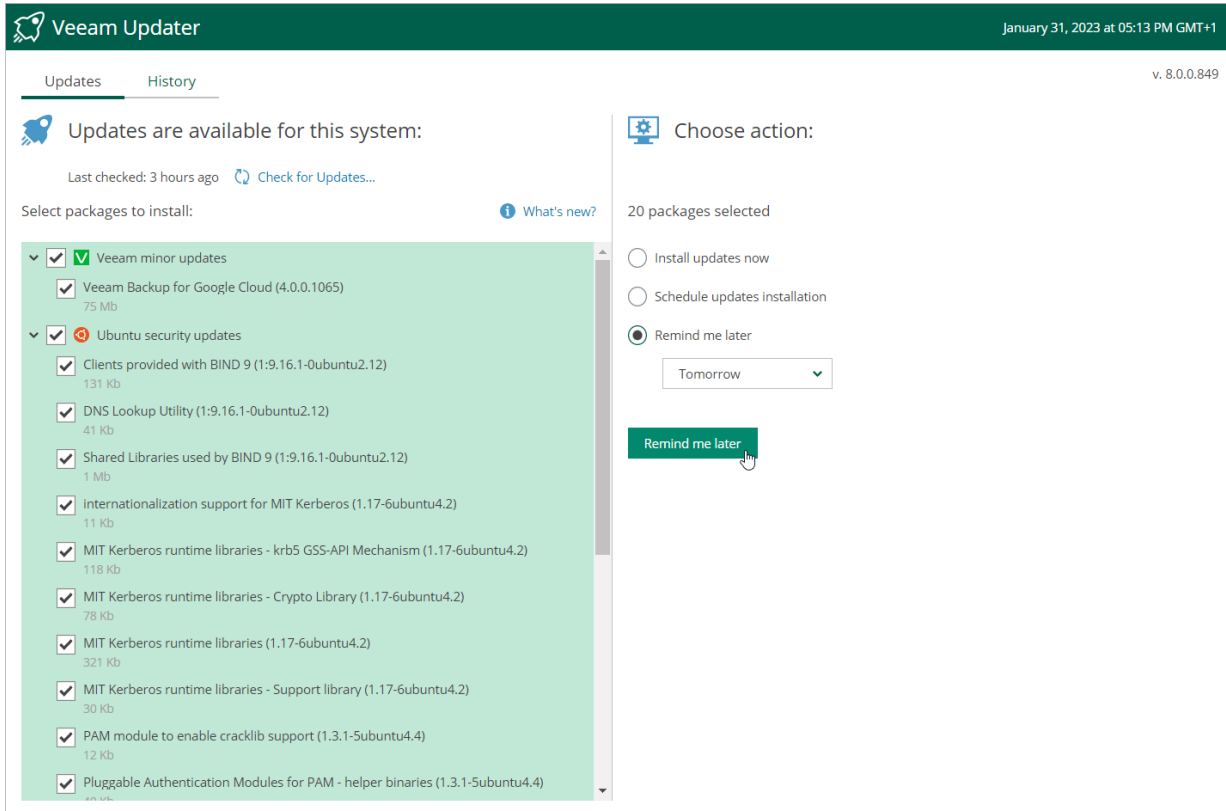
If you have not decided when to install available updates, you can set an update reminder — instruct Veeam Backup for Google Cloud to send an update notification later.

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 Google Cloud will send the reminder on the following Monday.

2. Click Remind me later.



The screenshot shows the Veeam Updater application interface. At the top, the title bar reads "Veeam Updater" on the left and "January 31, 2023 at 05:13 PM GMT+1" on the right. Below the title bar, there are two tabs: "Updates" (active) and "History". The version number "v. 8.0.0.849" is displayed 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: 3 hours ago" status and a "Check for Updates..." button. Below this, it says "Select packages to install:" followed by a "What's new?" link.
- Choose action:** This section shows "20 packages selected" and three radio button options: "Install updates now", "Schedule updates installation", and "Remind me later" (which is selected). Below the "Remind me later" option is a dropdown menu currently set to "Tomorrow". A green "Remind me later" button is highlighted with a mouse cursor.

The list of packages to be installed is as follows:

- Veeam minor updates
 - Veeam Backup for Google Cloud (4.0.0.1065) 75 Mb
- Ubuntu security updates
 - Clients provided with BIND 9 (1:9.16.1-0ubuntu2.12) 131 Kb
 - DNS Lookup Utility (1:9.16.1-0ubuntu2.12) 41 Kb
 - Shared Libraries used by BIND 9 (1:9.16.1-0ubuntu2.12) 1 Mb
 - Internationalization support for MIT Kerberos (1.17-6ubuntu4.2) 11 Kb
 - MIT Kerberos runtime libraries - krb5 GSS-API Mechanism (1.17-6ubuntu4.2) 118 Kb
 - MIT Kerberos runtime libraries - Crypto Library (1.17-6ubuntu4.2) 78 Kb
 - MIT Kerberos runtime libraries (1.17-6ubuntu4.2) 321 Kb
 - MIT Kerberos runtime libraries - Support library (1.17-6ubuntu4.2) 30 Kb
 - PAM module to enable cracklib support (1.3.1-5ubuntu4.4) 12 Kb
 - Pluggable Authentication Modules for PAM - helper binaries (1.3.1-Subuntu4.4)

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**.
3. Switch to the **Updates** tab.
4. Click **Check and View Updates**.
5. 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 Google Cloud 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's a green header with the Veeam logo and 'Veeam Updater' on the left, and the date 'May 5, 2022, 05:02 PM' and version 'v. 6.0.0.684' on the right. Below the header, there are two tabs: 'Updates' and 'History', with 'History' being the active tab. The main content area is titled 'Update sessions history' and features a table with two columns: 'Date' and 'Package'. The 'Date' column lists various dates from April 21, 2022, to May 5, 2022. The 'Package' column lists the names of the updates, such as 'Preparing for the update operation', 'GNU Bourne Again SHell (5.0-6ubuntu1.2)', and 'Veeam Backup for Google Cloud (3.0.0.824)'. A 'Status' column on the right of the package list shows green checkmarks and the word 'Success' for each entry. A 'View Full Log' link is visible above the package list.

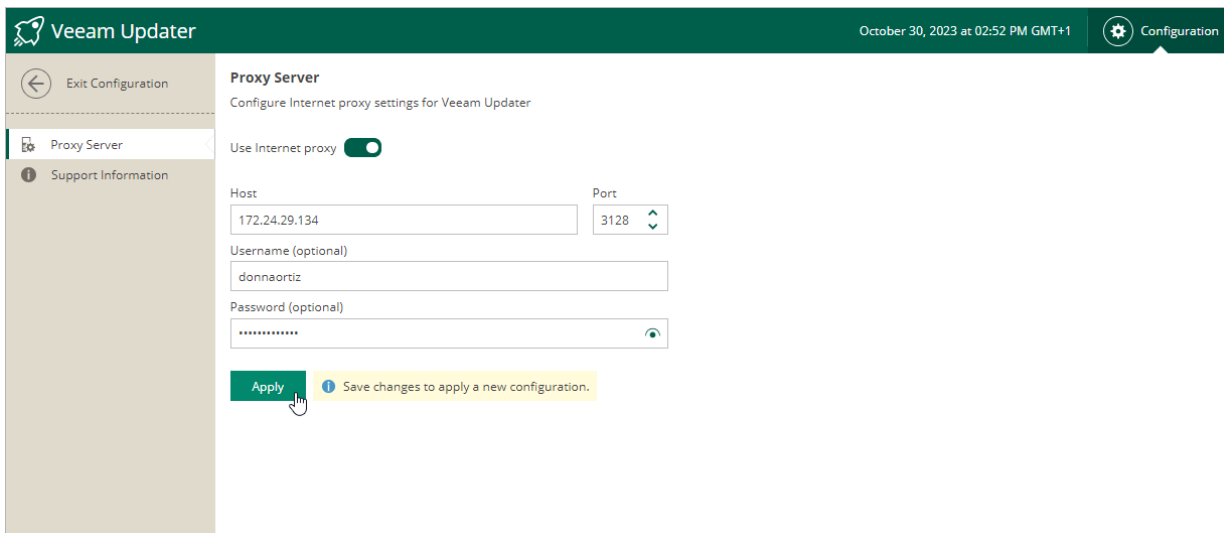
Date	Package	Status
May 5, 2022, 04:45 PM	Preparing for the update operation	Success
May 4, 2022, 03:41 PM	GNU Bourne Again SHell (5.0-6ubuntu1.2)	Success
May 3, 2022, 04:24 PM	command line tool for transferring data with URL syntax (7.68.0-1ubuntu2.10)	Success
May 3, 2022, 01:31 PM	information about the distributions' releases (data files) (0.43ubuntu1.10)	Success
May 2, 2022, 11:18 PM	fast, scalable, distributed revision control system (manual pages) (1:2.25.1-1ubuntu3.4)	Success
May 2, 2022, 12:41 PM	fast, scalable, distributed revision control system (1:2.25.1-1ubuntu3.4)	Success
April 29, 2022, 06:27 PM	easy-to-use client-side URL transfer library (GnuTLS flavour) (7.68.0-1ubuntu2.10)	Success
April 29, 2022, 04:13 PM	easy-to-use client-side URL transfer library (OpenSSL flavour) (7.68.0-1ubuntu2.10)	Success
April 29, 2022, 01:30 PM	SELinux library for manipulating binary security policies (3.0-1ubuntu0.1)	Success
April 28, 2022, 03:07 PM	Secure Sockets Layer toolkit - shared libraries (1.1.1f-1ubuntu2.13)	Success
April 27, 2022, 07:50 PM	Google Cloud Linux kernel headers (5.13.0.1024.29-20.04.1)	Success
April 27, 2022, 04:35 AM	Google Cloud Linux kernel Image (5.13.0.1024.29-20.04.1)	Success
April 26, 2022, 09:15 PM	Linux Kernel Headers for development (5.4.0-109.123)	Success
April 26, 2022, 12:36 AM	Dispatcher service for systemd-networkd connection status changes (2.1-2-ubuntu20.04.3)	Success
April 23, 2022, 04:32 AM	Secure Sockets Layer toolkit - cryptographic utility (1.1.1f-1ubuntu2.13)	Success
April 22, 2022, 04:40 PM	Veeam Backup for Google Cloud (3.0.0.824)	Success
April 22, 2022, 04:31 AM	File level recovery for Veeam backup (5.0.0.580)	Success
April 21, 2022, 04:28 PM		
April 21, 2022, 02:26 PM		

Configuring Web Proxy

To check for available package updates for Veeam Backup for Google Cloud, the Veeam Updater service running on the backup appliance connects to Veeam repositories 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. 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:
 - a. Switch to the **Configuration** page and do the following:
 - 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**.



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 asterisks. 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" on the left, "October 30, 2023 at 02:52 PM GMT+1" in the center, and "Configuration" on the right. The left sidebar shows "Exit Configuration", "Proxy Server", and "Support Information" options.

Getting Technical Support

If you have any questions or issues with Veeam Backup for Google Cloud, 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 components](#)
- The error message or an accurate description of the problem you are facing
- [Log files](#)

Viewing Product Details Using Web UI

To view the product details, do the following:

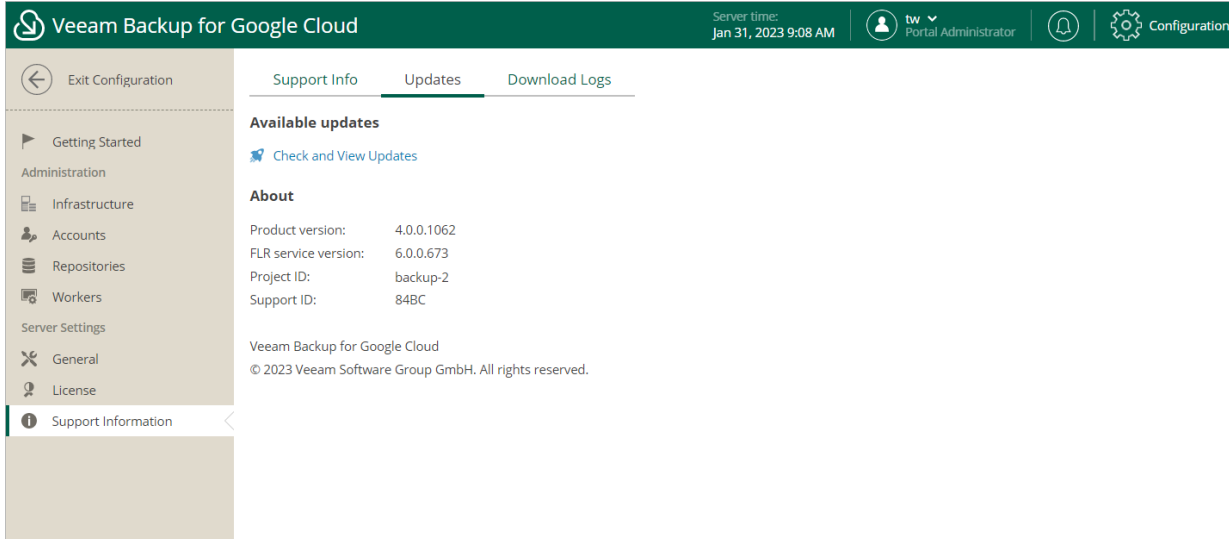
1. Switch to the **Configuration** page.
2. Navigate to **Support Information > Updates**.

The **About** section of the **Updates** tab displays the following information:

- **Product version** — the currently installed version of Veeam Backup for Google Cloud.
- **Project ID** — the unique identification number of the Google Cloud project to which the VM instance running Veeam Backup for Google Cloud belongs.
- **Support ID** — the unique identification number of the Veeam support contract.
- **FLR service version** — the version of the File-Level Recovery Service currently running on the backup appliance.

TIP

You can click the link in the **Available Updates** section of the **Updates** tab to check for, download and install new product versions and available package updates. For more information, see [Updating Veeam Backup for Google](#).



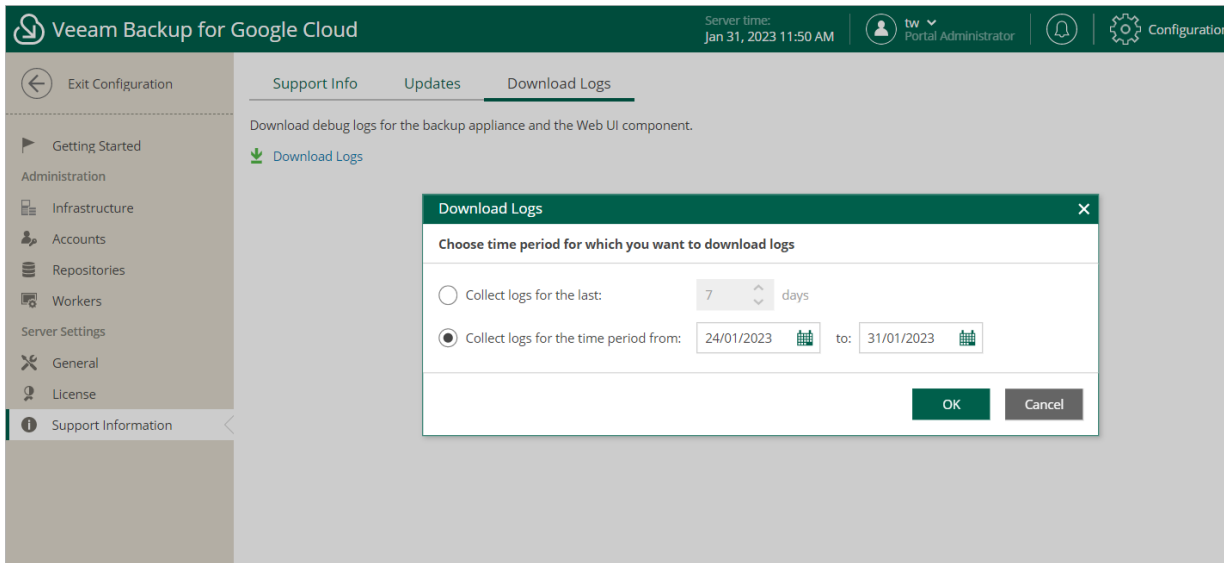
Downloading Product Logs Using Web UI

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 the logs will be collected:
 - o 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 the time period from** option if you want to collect data for a specific period of time in the past.

After you click **OK**, the logs will be saved locally in the default download folder as a single .ZIP archive.



Downloading Product Logs Using 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:

Export logs for these objects:

Export all logs for selected components (may result in a very large log package)

Managed servers:

Server	Components
<input checked="" type="checkbox"/> backupsrv50.tech.lo...	Installer, Mount Server, Transport, Veeam A...
<input checked="" type="checkbox"/> atlanta	

Select All

Clear All

< Previous **Next >** Finish Cancel

Appendix. Configuring Deployment Mode

By default, worker instances launched by Veeam Backup for Google Cloud access protected Google Cloud resources through private virtual networks. The only exception is worker instances deployed during file-level restore operations to access the [File-Level Restore browser](#).

If you do not plan to perform file-level recovery or if you plan to access the browser through private networks only, do the following:

1. Connect to the backup appliance through SSH as described in [Google Cloud documentation](#).
2. Edit the **FlrPerformer** value in the `/opt/veeam/gcpbackup/JobManagerSettings.json` configuration file:

```
"FlrPerformer": {  
    "DisableWorkerPublicIp": true  
}
```

If you want your worker instances to be deployed as [Shielded VMs](#), edit the **Worker** value in the `/opt/veeam/gcpbackup/ServiceSettings.json` configuration file, and restart the `veeambackup` service:

```
"Worker": {  
    "EnableVtpm": true,  
    "EnableIntegrityMonitoring": true  
}
```

If you want your worker instances to be deployed with public IP addresses, add the **SqlWorker** parameter to the `/opt/veeam/gcpbackup/ServiceSettings.json` configuration file:

```
"SqlWorker": {  
    "AllowExternalIp": true  
}
```