




# HYCU R-Cloud



## User Guide

September 2024

# Legal notices

## Copyright notice

© 2024 HYCU. All rights reserved.

This document contains proprietary information, which is protected by copyright. No part of this document may be photocopied, reproduced, distributed, transmitted, stored in a retrieval system, modified or translated to another language in any form by any means, without the prior written consent of HYCU.

## Trademarks

HYCU logos, names, trademarks and/or service marks and combinations thereof are the property of HYCU or its affiliates. Other product names are the property of their respective trademark or service mark holders and are hereby acknowledged.

Amazon Web Services, AWS, Amazon EC2, Amazon S3, and Amazon Cognito are trademarks of Amazon.com, Inc. or its affiliates.

Azure®, Microsoft®, Microsoft Edge™, Microsoft Entra™, and Windows® are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

GCP™, GKE™, Google Chrome™, Google Cloud™, Google Cloud Platform™, Google Cloud Storage™, and Google Compute Engine™ are trademarks of Google LLC.

Kubernetes® is the registered trademark of The Linux Foundation in the United States and/or other countries.

Linux® is the registered trademark of Linus Torvalds in the U.S. and other countries.

Microsoft®, Microsoft Edge™, and Windows® are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

Mozilla and Firefox are trademarks of the Mozilla Foundation in the U.S. and other countries.

SAP HANA® is the trademark or registered trademark of SAP SE or its affiliates in Germany and in several other countries.

## **Disclaimer**

The details and descriptions contained in this document are believed to have been accurate and up to date at the time the document was written. The information contained in this document is subject to change without notice.

HYCU provides this material "as is" and makes no warranty of any kind, expressed or implied, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. HYCU shall not be liable for errors and omissions contained herein. In no event shall HYCU be liable for any direct, indirect, consequential, punitive, special or incidental damages, including, without limitation, damages for loss and profits, loss of anticipated savings, business interruption, or loss of information arising out of the use or inability to use this document, or any action taken based on the information contained herein, even if it has been advised of the possibility of such damages, whether based on warranty, contract, or any other legal theory.

The only warranties for HYCU products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty.

## **Notice**

This document is provided in connection with HYCU products. HYCU may have copyright, patents, patent applications, trademark, or other intellectual property rights covering the subject matter of this document.

Except as expressly provided in any written license agreement from HYCU, the furnishing of this document does not give you any license to these patents, trademarks, copyrights, or other intellectual property on HYCU products. Use of underlying HYCU product(s) is governed by their respective Software License and Support Terms.

**Important:** Please read Software License and Support Terms before using the accompanying software product(s).

HYCU

[www.hycu.com](http://www.hycu.com)

# Contents

1 About R-Cloud (formerly HYCU Protégé)	11
Key features and benefits	11
Data protection environment overview	13
R-Cloud data protection	14
2 Starting with R-Cloud	15
Service pricing	15
Backup and data retention pricing	16
Subscribing to the service	18
Subscribing from AWS Marketplace	19
Subscribing from Google Cloud Marketplace	20
Upgrading from an existing free subscription	21
Signing in to R-Cloud	22
3 Establishing a data protection environment	24
Switching the user interface context	25
Selecting an R-Cloud protection set	26
Setting up targets	27
Backup target types in R-Cloud	27
Setting up an Amazon S3 target	29
Setting up an Azure target	31
Setting up a Google Cloud target	33
Setting up an S3 compatible target	34
Defining your backup strategy	36
Taking advantage of predefined policies	36
Creating custom policies	37

Creating backup windows .....	43
Creating data archives .....	45
Setting default policies .....	48
Setting up automatic policy assignment .....	49
Enabling access to data .....	51
Enabling access to instances .....	51
<b>4 Protecting SaaS applications .....</b>	<b>57</b>
Configuring SaaS application backup options .....	58
Backing up SaaS applications .....	60
Restoring SaaS applications .....	61
<b>5 Protecting applications .....</b>	<b>64</b>
Protecting SAP HANA applications .....	65
Preparing for SAP HANA application protection .....	65
Backing up SAP HANA applications .....	68
Restoring SAP HANA applications .....	70
Protecting Google Kubernetes Engine applications .....	72
Preparing for Google Kubernetes Engine application protection .....	72
Backing up Google Kubernetes Engine applications .....	77
Restoring Google Kubernetes Engine applications .....	78
<b>6 Protecting instances .....</b>	<b>84</b>
Planning instance protection .....	84
Preparing your data protection environment .....	84
Configuring instance backup options .....	86
Backing up instances .....	92
Restoring instances .....	94
Restoring a single instance or its disks .....	95

Restoring multiple instances or disks belonging to multiple instances in a single session .....	121
Restoring individual files or folders .....	132
Restoring files or folders to an instance .....	133
Restoring files or folders to a bucket .....	138
<b>7 Protecting buckets .....</b>	<b>141</b>
Configuring bucket backup options .....	142
Backing up buckets .....	145
Restoring buckets .....	146
<b>8 Performing daily tasks .....</b>	<b>151</b>
Using the R-Cloud dashboard .....	152
Viewing information about entities .....	153
Viewing entity information .....	154
Viewing entity details .....	159
Managing policies .....	163
Viewing policy information .....	163
Creating a policy .....	164
Editing a policy .....	164
Deleting a policy .....	164
Managing targets .....	165
Viewing target information .....	165
Editing targets .....	167
Deactivating and activating targets .....	167
Removing targets .....	168
Checking task statuses .....	169
Viewing events .....	170
Configuring event notifications .....	172

Creating email notifications .....	172
Creating webhook notifications .....	173
Using R-Cloud reports .....	175
Getting started with reporting .....	175
Viewing reports .....	178
Generating reports .....	179
Scheduling reports .....	179
Exporting and importing reports .....	180
Filtering and sorting data .....	181
Filtering data in panels .....	182
Sorting data in panels .....	187
Performing manual backups .....	187
Expiring backups manually .....	188
Exporting the contents of the panel .....	190
Viewing subscription information .....	190
Customizing your R-Cloud web user interface .....	192
<b>9 Customizing R-Cloud .....</b>	<b>193</b>
Managing protection sets .....	194
Creating protection sets .....	195
Editing protection sets .....	196
Deleting protection sets .....	198
Adding cloud accounts .....	199
Adding AWS IAM roles .....	200
Adding Azure service principals .....	201
Adding Google Cloud service accounts .....	202
Adding S3 compatible accounts .....	203
Navigating the HYCU Marketplace .....	204



Managing sources .....	206
Managing AWS accounts .....	206
Managing Google Cloud projects .....	208
Managing R-Cloud modules .....	210
Discovering SaaS services .....	213
Exploring R-Graph .....	215
Managing identity and access .....	221
Managing identity providers .....	222
Managing users .....	223
Managing roles .....	225
Requesting a password reset .....	227
Stopping protection for individual sources .....	227
Excluding instances from synchronization by tagging the instance in AWS or Google Cloud .....	228
<b>10 Troubleshooting .....</b>	<b>230</b>
Known problems and solutions .....	231
Restore of individual files ends with errors or fails .....	231
Inability to change the protection set or to sign in .....	232
Problem with sorting data in the Events panel .....	232
Inability to set up manually created Google Cloud targets .....	232
Assigning a policy to a Google Cloud instance fails .....	233
Snapshot creation fails for instances in a specific Google Cloud project .....	233
Task progress indicator remains at 0% during the backup of a Google Cloud instance .....	234
<b>11 Unsubscribing from R-Cloud .....</b>	<b>235</b>
Stopping service charges .....	235

Preventing account access .....	237
Preventing access to an AWS account .....	238
Preventing access to a Google Cloud account .....	238
Removing the HYCU Managed Service Account permissions .....	239
Canceling your R-Cloud subscription .....	239
Canceling the R-Cloud subscription in the AWS Marketplace .....	240
Canceling the R-Cloud subscription in the Google Cloud Marketplace .....	240
<b>A Resources created by R-Cloud .....</b>	<b>242</b>
<b>B Bulk restore specifications .....</b>	<b>245</b>
Elements of a bulk restore specification .....	245
<b>C Least-privilege permissions used by R-Cloud .....</b>	<b>249</b>
Using a role template for AWS .....	249
AWS permissions required by R-Cloud .....	251
Using a role template for Google Cloud .....	254
Google Cloud permissions required by R-Cloud .....	255
<b>D Deploying a HYCU backup controller .....</b>	<b>259</b>
Deploying a HYCU backup controller to AWS .....	259
Accessing the HYCU web user interface .....	263
Deploying a HYCU backup controller to Google Cloud .....	263
Accessing the HYCU web user interface .....	267

# Chapter 1

## About R-Cloud (formerly HYCU Protégé)

HYCU R-Cloud (R-Cloud), formerly known as HYCU Protégé, is a fully managed backup and recovery service for public clouds and Software as a Service (SaaS) applications that is specifically designed to make data protection as simple and cost-effective as possible, to improve your business agility, and to bring unified security, reliability, performance, and user experience.

The following are the key elements of R-Cloud:

- Service-based backup and recovery
- Improved business agility
- Intuitive user interface
- Low-impact application backup
- Automated target management
- At-a-glance overview of your environment
- Native integration with the platform
- Reduced complexity

## Key features and benefits

The following features make R-Cloud a solution that can transform your business—achieving complete compliance and data protection:

- **Protection against data loss**

Delivers native data protection for instances in Amazon EC2 and Google Cloud, applications running on instances and clusters, Amazon S3 and Google Cloud Storage buckets, and SaaS applications, ensuring data consistency and easy recoverability.

- **Data protection in a few minutes**

Data protection can be enabled in a few minutes after you subscribe to R-Cloud, with no deployment and configuration concerns.

- **SaaS discovery and protection**

In-built SaaS discovery provides new-found visibility into SaaS applications that your organization is using and the data protection status of these applications.

The discovered results are presented via R-Graph – a visual representation of your SaaS data protection environment – enabling you to quickly gain more insight into the status of your SaaS application data protection.

- **Application discovery and protection**

In-built application discovery provides new-found visibility into instances running in cloud environments and clusters, pinpointing where each application is running. The application-specific backup and restore flow ensures that the entire application data is protected and can be recovered to a consistent state and a specific point in time.

- **Predefined policies and options for policy customization**

Simplifies implementation of data protection by providing predefined policies and includes options for policy customization that can address your special data protection needs.

- **Scheduled backups**

Automatic backup scheduling provides data protection based on your recovery point objectives (RPOs).

- **Centralized data protection management and monitoring**

You can join your sources into protection sets to establish centralized data protection management and monitoring.

- **Lower impact on the environment**

Agentless architecture reduces backup load on production instances. In addition, backup windows enable you to completely avoid the impact of backup activity on your production environment during peak hours.

- **Use of data archives**

When you create an archive of data, you ensure your data is isolated from your current activity and safely stored for future reference.

- **Restore of individual files**

A possibility to restore one or more files is an alternative to restoring the entire instance or disk.

- **At-a-glance overview of the data protection environment**

The R-Cloud dashboard helps you to identify potential problems and bottlenecks to improve the performance of your data protection environment.

- **Optimized consumption of storage space**

The HYCU changed block tracking feature slows down the growth of backup data on targets, resulting in significant space savings and consequently reduced storage cost.

- **Integration with your cloud provider billing system**

Cost of data protection is billed by your cloud provider through existing billing or management accounts, without requiring you to provide additional billing information.

- **Business continuity of your data protection environment across different infrastructures**

The SpinUp functionality ensures data resilience by allowing you to migrate protected data between the on-premises and cloud infrastructures. In the event of a disaster in your on-premises environment, the SpinUp functionality provides disaster recovery of data to cloud. For details on the supported on-premises infrastructures and how to employ the SpinUp functionality, see HYCU R-Cloud Hybrid Cloud Edition documentation.

## Data protection environment overview

Before you start protecting data with R-Cloud, make yourself familiar with the following terms related to the data protection environment:

Term	Description
R-Cloud web user interface	An interface for protecting entities and administering the data protection environment.
Sources	Environments for which R-Cloud provides data protection—Google Cloud projects, AWS accounts, and R-Cloud modules.
Protection sets	Groups that join together sources that you add to R-Cloud.

Term	Description
Entities	Resources to which you can assign a policy and for which you therefore provide data protection—SaaS applications, SAP HANA and GKE applications, Amazon EC2 and Google Cloud instances, and Amazon S3 and Google Cloud buckets. Data is always protected at a granular level, allowing you to restore either the whole entities or their parts.
Targets	Storage locations that R-Cloud uses for storing backup data—Amazon S3 buckets, Azure storage accounts, Google Cloud buckets, and S3 compatible buckets. Backup data can also be stored as snapshots.

## R-Cloud data protection

With the R-Cloud data protection solution, you can be confident that your business data is protected, which means that it is backed up in a consistent state, stored to a target, and can be restored.

The AWS accounts, Google Cloud projects, and the R-Cloud modules that you add as sources define the scope of data protection.

R-Cloud enables you to protect SaaS application data, applications, instances, and data in buckets. After you establish your data protection environment, you can enable data protection. After the first backup is successfully completed, you can restore the data if it becomes damaged or corrupted.

# Chapter 2

## Starting with R-Cloud

You can start protecting data after you perform the following tasks:

Task	Instructions
Getting familiar with R-Cloud pricing concepts	<a href="#">“Service pricing”</a> below
Subscribing to R-Cloud	<a href="#">“Subscribing to the service”</a> on page 18
Signing in to the R-Cloud web user interface	<a href="#">“Signing in to R-Cloud”</a> on page 22

### Service pricing

Because R-Cloud utilizes the cloud environment for its service needs, when you enable data protection, you are charged for the backup service, data retention, and the resources that are required for the backup and recovery services.

The total data protection cost is the sum of the following costs:




Figure 2-1: Data protection cost

Cost	Details
Backup and data retention	Cost of backing up data and data retention. For details, see <a href="#">“Backup and data retention pricing”</a> on the next page.
Backup data storage	Cost of storing backup data. The following factors are

Cost	Details
	<p>considered:</p> <ul style="list-style-type: none"> <li>• Backup target type (a snapshot or a target)</li> <li>• Backup frequency</li> <li>• Size of backup data</li> <li>• Backup retention period</li> </ul> <p>If you use a target for storing data, the following is also considered:</p> <ul style="list-style-type: none"> <li>• Use of copies of backup data</li> <li>• Use of data archives, configured archive tiers and their retention periods</li> </ul>
Temporary resources	<p>Cost of temporary resources that R-Cloud creates in the cloud when performing the following tasks:</p> <ul style="list-style-type: none"> <li>• Instance rediscovery after assigning a credential group</li> <li>• Backup of instances</li> <li>• Backup of applications</li> <li>• Backup of buckets</li> <li>• Restore of instances or entire instance disks</li> <li>• Restore of individual files or folders</li> <li>• Restore of applications</li> <li>• Restore of buckets</li> </ul>

An R-Cloud subscription includes a 14-day free trial period. During this time, HYCU does not charge you for the backup and data retention cost. The cost of backup data storage and temporary resources is charged by your cloud provider as usual.

## Backup and data retention pricing

The R-Cloud backup and data retention pricing model provides you with the simplicity and transparency of consumption-based pricing. At the end of your 14-day free trial period, you are billed according to the subscription plan that you select when subscribing to R-Cloud. For details on the subscription plans, see [“R-Cloud subscription plans” on page 18](#).



Pricing for data protection is based on the following (within a monthly billing cycle):

- Capacity of all disks belonging to protected instances and applications
- Size of protected buckets
- Pricing tiers to which protected instances and buckets belong

A pricing tier to which a protected instance, application, or bucket belongs is determined when you assign a policy to the instance, the application, or the bucket. R-Cloud automatically associates the instance, application, or bucket with one of the pricing tiers based on the value of the Backup every option in the policy that defines how frequently data is backed up. For details on policies, see [“Defining your backup strategy” on page 36](#).

Depending on how frequently your data is backed up, each protected instance, application, or bucket belongs to one of the following pricing tiers:

Pricing tier	Data backup frequency (in hours)
platinum	1–3 hours
gold	4–11 hours
silver	12–23 hours
bronze	24 hours or more

### Considerations

- If an instance, an application, or a bucket is deleted from the cloud, but it still has at least one valid restore point available, it is considered protected (its status is Protected deleted) and HYCU automatically associates such an entity with the bronze pricing tier. In the case of instances, it charges you for protecting only the included disks.
- If you unassign a policy from an instance, an application, or a bucket that still has at least one valid restore point available, such an entity is considered protected and HYCU automatically associates it with the bronze pricing tier. In the case of instances, it charges you for protecting only the included disks.
- *Applicable for instances and applications running on them.* If you assign policies to an instance and an application running on the same instance, keep in mind that you will be charged for both protecting the instance and protecting the application.

## R-Cloud subscription plans

R-Cloud offers you the following subscription plans:


- Pay-as-you-go plan  
Select this plan if you want to pay only for what you use for data protection each month.
- Annual subscription plans  
You can choose among different annual subscription plans with token-based pricing.

For more information on pricing and subscription plans for AWS, Google Cloud, and SaaS applications, see your cloud provider marketplace ([AWS Marketplace](#) or [Google Cloud Marketplace](#)) or contact your HYCU sales representative.

## Subscribing to the service

You subscribe to R-Cloud online from your cloud provider marketplace and HYCU then automatically activates the service for you. This is usually done by one user for an entire organization. Depending on your cloud platform, see one of the following sections:

Cloud platform	Instructions
AWS	<a href="#">“Subscribing from AWS Marketplace” on the next page</a>
Google Cloud	<a href="#">“Subscribing from Google Cloud Marketplace” on page 20</a>

 **Note** If you are using the trial version of R-Cloud, you upgrade to a paid R-Cloud subscription from your existing free subscription. For instructions, see [“Upgrading from an existing free subscription” on page 21](#).

## Subscribing from AWS Marketplace

Prerequisites

- You have access to an AWS account.
- Your user account has the `AWSMarketplaceManageSubscriptions` predefined role attached (`arn:aws:iam::aws:policy/AWSMarketplaceManageSubscriptions`).

For details, see AWS documentation.

### Consideration

If you violate the terms of use of R-Cloud, HYCU may temporarily suspend the service for your subscription. Your complete data protection environment is retained for the duration of suspension, but you cannot use the service until the violation is resolved.

### Procedure

1. Open a web browser and go to the [HYCU | AWS Market](#) webpage.
2. Read the solution description, and then click **View purchase options**.
3. On the Configure your contract page, check the displayed contract information, and click **Create contract**. If required, you can modify the contract information.
4. Verify the contract summary, and then click **Pay now**.
5. Click **Setup your account** to complete the registration.
6. On the R-Cloud sign-in webpage select **Create New Subscription** and click **Continue**.
7. On the New subscription page, enter the required information and click **Submit**.
8. R-Cloud is deployed and you are redirected to the R-Cloud sign-in page. Additionally, an email with the sign in link and HYCU account details is sent to you.

HYCU automatically creates a user account for the HYCU Support portal for your subscription and sends you an email notification about it. You can use this account for submitting requests to HYCU Support.

## Subscribing from Google Cloud Marketplace

### Prerequisite

The Google Account you are using has the necessary roles required for purchasing solutions on the Google Cloud Marketplace. For details, see Google Cloud documentation.

## Consideration

If you violate the terms of use of R-Cloud, HYCU may temporarily suspend the service for your subscription. Your complete data protection environment is retained for the duration of suspension, but you cannot use the service until the violation is resolved.

## Procedure

1. Open a web browser and go to the [HYCU | Marketplace - Google Cloud](#) webpage.
2. *Only if using Microsoft Edge.* Enable pop-ups for the \*.cloud.google.com website.
3. Read the solution description, and then click **Subscribe**.
4. On the New HYCU subscription page, in the Subscribe pane, check the displayed billing account information, take note of it, and click **Subscribe**.
5. In the Activate pane, click **Register with HYCU, Inc.**
6. On the R-Cloud sign-in webpage, select **Create New Subscription** and click **Continue**.
7. On the New subscription page, enter the required information and click **Submit**.
8. R-Cloud is deployed and you are redirected to the R-Cloud sign-in page. Additionally, an email with the sign in link and HYCU account details is sent to you.

HYCU automatically creates a user account for the HYCU Support portal for your subscription and sends you an email notification about it. You can use this account for submitting requests to HYCU Support.

## Upgrading from an existing free subscription

### Prerequisites

- You must have the HYCU account ID of your free subscription. For details, see [“Viewing subscription information” on page 190](#).
- *For AWS:* You have access to an AWS account and your user account has the `AWSMarketplaceManageSubscriptions` predefined role attached (`arn:aws:iam::aws:policy/AWSMarketplaceManageSubscriptions`). For details, see AWS documentation.

- *For Google Cloud:* The Google Account you are using has the necessary roles required for purchasing solutions on the Google Cloud Marketplace. For details, see Google Cloud documentation.

### Procedure

1. In the toolbar, click **Upgrade**.
2. Select the cloud provider marketplace offer that you want to use to perform the upgrade:
  - Click **Update with AWS** to use AWS Marketplace.
  - Click **Update with Google** to use Google Cloud Marketplace.
3. On your selected marketplace webpage, complete the subscription procedure:

Selected marketplace	Steps
AWS Marketplace	<ol style="list-style-type: none"> <li>a. Read the solution description, and then click <b>View purchase options</b>.</li> <li>b. On the Configure your contract page, check the displayed contract information, and click <b>Create contract</b>. If required, you can modify the contract information.</li> <li>c. Verify the contract summary, and then click <b>Pay now</b>.</li> </ol>
Google Cloud Marketplace	<ol style="list-style-type: none"> <li>a. Read the solution description, and then click <b>Subscribe</b>.</li> <li>b. On the New HYCU subscription page, in the Subscribe pane, check the displayed billing account information, take note of it, and click <b>Subscribe</b>.</li> <li>c. In the Activate pane, click <b>Register with HYCU, Inc.</b></li> </ol>

4. On the R-Cloud sign-in webpage, select **Upgrade Existing Free Subscription** and click **Continue**.
5. On the Upgrade existing free subscription page, enter your HYCU account ID and click **Submit**.

HYCU upgrades your subscription and sends you an email notification.

# Signing in to R-Cloud

After successfully subscribing to R-Cloud, you can sign in to the R-Cloud web user interface.

## Prerequisites

- You must use a supported web browser. For a list of supported web browsers, see the *HYCU R-Cloud Compatibility Matrix*.
- *Only if you want to protect a Google Cloud project:*
  - In Google Cloud, the Compute Engine default service account must be present on the project that you plan to protect and it must have the Editor role granted. If this service account is not available, you must set up an alternative service account and grant it the Editor role. The name of the service account must be in the following format:  
`hycu-<projectNumber>@<projectId>.iam.gserviceaccount.com`.
  - In Google Compute Engine, your Google Account must have the following roles granted on the projects with instances, applications, and buckets that you plan to protect:
    - Compute Admin (`roles/compute.admin`)
    - Service Account User (`roles/iam.serviceAccountUser`)
  - In the Google Cloud Storage service, your Google Account must have the Storage Admin (`roles/storage.admin`) role granted on the projects whose targets you plan to use for storing data.
  - The Cloud Pub/Sub API must be enabled on the Google Cloud projects with instances, applications, and buckets that you plan to protect.


For details, see Google Cloud documentation.

## Procedure

1. Depending on whether you are signing in for the first time or not, access the R-Cloud sign-in webpage in one of the following ways:
  - If you are accessing the sign-in page for the first time after subscribing: After you subscribe to R-Cloud and R-Cloud is deployed, you are redirected to the sign-in webpage. If you are not redirected, open a web browser and go to the R-Cloud sign-in webpage by using the link in the email that received when you subscribed to R-Cloud. In both cases, the account ID is already part of the URL and you do not need to enter it.


- If you are not accessing the sign-in page for the first time:  
Open the [R-Cloud](#) web page, enter the HYCU account ID you received when you subscribed to R-Cloud, and then click **Next**.

Once you successfully sign in, your account ID is stored by the browser and you no longer need to enter it when you sign-in.

 **Tip** You can set a sign-in alias for your HYCU account. For details, see [“Viewing subscription information” on page 190](#).

2. Depending on how you want to sign-in to R-Cloud, do one of the following:


- *By using dedicated sign-in credentials for HYCU.* Enter your sign-in name and password.

 **Important** You must change the generated temporary password during the first sign-in.

- *By using an identity provider.* Click the preferred identity provider, and then, if required, enter your credentials.

For details on how to integrate R-Cloud with identity providers, see [“Managing identity providers” on page 222](#).

After you sign in to the R-Cloud web user interface, the Dashboard panel appears, and you can start establishing your data protection environment and protecting data.

 **Important** You are automatically signed out of the R-Cloud web user interface after 15 minutes of inactivity and any unsaved changes are lost.

To sign out manually, click  **<EmailAddress>** to open the Session menu, and then click **Sign Out**.

# Chapter 3

## Establishing a data protection environment

After you sign in to R-Cloud, you must establish a data protection environment in which data will be effectively protected.

### Consideration

Keep in mind that the role you have assigned determines what kind of actions you can perform. For details on roles, see [“Managing roles” on page 225](#).

If you have the Administrator role assigned, you can switch between the Subscription and Protection set contexts. Depending on the scope of the tasks that you want to perform, click ▼ next to the name of the currently selected context to switch to another one. The Subscription context enables you to perform administration tasks related to the selected subscription such as adding identity providers, adding or removing users, and changing roles, whereas the Protection set context enables you to perform data protection tasks related to the selected protection set. See [“Managing identity and access” on page 221](#) and [“Managing protection sets” on page 194](#) for details.

### Tasks

Establishing a data protection environment involves the following tasks:

Task	Instructions
1. <i>Only if you plan to use multiple protection sets.</i> Configure a protection set and select it.	<a href="#">“Managing protection sets” on page 194</a> and <a href="#">“Selecting an R-Cloud protection set” on page 26</a>
2. Add the sources (R-Cloud modules, AWS accounts, or Google Cloud projects) to R-Cloud.	<a href="#">“Managing sources” on page 206</a>



Task	Instructions
3. <i>Only if you plan to use manually created targets.</i> Set up targets.	<a href="#">“Setting up targets” on page 27</a>
4. Decide for predefined policies or create custom ones.	<a href="#">“Defining your backup strategy” on page 36</a>
5. <i>Required only in special data protection scenarios.</i> Configure credential groups and assign them to instances.	<a href="#">“Enabling access to data” on page 51</a>


After the data protection environment is established, data protection can be accomplished in several ways to fulfill your particular business needs.

## Switching the user interface context

In the R-Cloud user interface, the scope of tasks you can perform depends on the context you select. You can choose between a subscription context that is used for administration tasks and a protection set context:

User interface context	Description
Subscription	In the Subscription context, only the IAM panel is active. Use this context to perform administration tasks related to your subscription, such as adding identity providers, adding or removing users, or changing roles. See <a href="#">“Managing identity and access” on page 221</a> .
Protection set	In the Protection set context, you select the scope of data protection by selecting a specific protection set.

### Procedure

1. On the toolbar, click  next to the name of the currently selected protection set or subscription.
2. Depending on the context that you want to select, do the following:

User interface context	Instructions
Subscription	From the drop-down menu, select <b>Subscription</b> .
Protection set	From the list of the recently used protection sets, select the preferred protection set. If the preferred protection set is not on the list, click <b>More</b> to see all available protection sets, select the preferred protection set, and then click <b>Confirm</b> .


The R-Cloud web user interface switches the context. Your selection is remembered for the next time you sign in.



## Selecting an R-Cloud protection set

An environment for which R-Cloud provides data protection consists of one or more protection sets that join together sources—AWS accounts, Google Cloud projects, and R-Cloud modules. When you subscribe to R-Cloud, a default protection set is created automatically. Depending on your business needs, you can create additional protection sets and distribute your sources among them, having in mind that you must implement data protection for each protection set individually. For details on managing protection sets, see [“Managing protection sets” on page 194](#).

Selecting an R-Cloud protection set determines your scope of data protection. If no more than one protection set is available in your data protection environment, your data protection scope is always the same and you can safely skip the procedure described in this section.

### Procedure

1. On the toolbar, click  next to the name of the currently selected protection set.
2. From the list of the recently used protection sets, select the protection set with the entities that you want to protect.

 **Note** The currently selected protection set is represented by the  icon.

If the preferred protection set is not on the list, click **More** to see all available protection sets, select the preferred protection set, and then click **Confirm**.

The R-Cloud web user interface switches the context to the selected scope of data protection. Your selection is remembered for the next time you sign in.

## Setting up targets

Targets are locations where protected data is stored. In addition to using targets to store protected data, R-Cloud also allows you to define a snapshot as a location for storing your data.

Setting up targets includes the following tasks:

Task	Instructions
1. Get familiar with backup target types in R-Cloud.	<a href="#">“Backup target types in R-Cloud”</a> below
2. Set up the preferred targets.	Depending on which target type you want to set up, see one of the following sections: <ul style="list-style-type: none"> <li>• <a href="#">“Setting up an Amazon S3 target”</a> on page 29</li> <li>• <a href="#">“Setting up an Azure target”</a> on page 31</li> <li>• <a href="#">“Setting up a Google Cloud target”</a> on page 33</li> <li>• <a href="#">“Setting up an S3 compatible target”</a> on page 34</li> </ul>

## Backup target types in R-Cloud

You can define a target or a snapshot as a location where protected data is stored by selecting the preferred backup target type in the R-Cloud policy. For details on backup target types, see the following sections:

- [“Target”](#) on the next page
- [“Snapshot”](#) on page 29

## Target

Protected data can be stored on targets that you create yourself or R-Cloud creates for you automatically.

**⚠ Caution** Never delete any targets used by R-Cloud because this may result in data loss. Additionally, within targets, ensure that the `hycu/backups/` folders are always kept intact.

### Manually created targets

You can create Amazon S3 buckets, Azure storage accounts, Google Cloud buckets, or S3 compatible buckets and set them up as targets in R-Cloud.

### Automatically created targets

R-Cloud can automatically create Amazon S3 or Google Cloud targets while backing up data. These targets are always created in the same region as the entities that you are protecting, unless you are protecting SaaS applications that do not run in AWS or Google Cloud. In this case, the targets are created in the default region of any AWS account or Google Cloud project that was added to R-Cloud as a source, or in the default region of any AWS account or Google Cloud project in which a manually created target resides.

**ⓘ Important** Automatically created targets are not available if you are protecting the following entities:

- SAP HANA applications
- SaaS applications that are related to R-Cloud modules that support storing data on a staging target

For R-Cloud modules that do not support storing data on a staging target, automatically created targets are available only for storing backup data of the related SaaS applications (and not copies of backup data or archive data).

The same target is used for storing the protected data of multiple entities where possible. You can use these targets also for storing your data.

For the target naming conventions, see [“Resources created by R-Cloud” on page 242](#).

**📄 Note** *Only if you plan to protect SaaS applications.* If the R-Cloud module supports storing data on a staging target, R-Cloud uses a staging target either to temporarily store SaaS application data before it is moved to the target

that you define in the R-Cloud policy, or to store SaaS application data as a snapshot. For details on staging targets and how to add them to R-Cloud, see [“Adding R-Cloud modules” on page 210](#).

## Snapshot

If you are protecting SaaS applications, GKE applications using persistent volumes, or instances, you can define a snapshot as a location for storing backup data. Snapshots that are created by R-Cloud are stored at the following locations:

- *For SaaS applications:* On staging targets or remote storage.
- *For GKE applications using persistent volumes:* On clusters on which the GKE applications are deployed.
- *For AWS instances:* In the AWS account that contains the instances.
- *For Google Cloud instances:* In the Google Cloud project that contains the instances.

If snapshots created by R-Cloud are deleted from any of these locations, you will not be able to restore backup data from this location. However, you can still restore your data from targets if copies of backup data or data archives exist.

For the snapshot naming conventions, see [“Resources created by R-Cloud” on page 242](#).

## Setting up an Amazon S3 target

### Prerequisite

An AWS IAM role must be added to R-Cloud. You can do this in one of the following ways:

- By creating an IAM role as part of adding the AWS account where the target resides to R-Cloud as a source. For instructions, see [“Managing sources” on page 206](#).
- By adding an IAM role to R-Cloud as a cloud account. For instructions, see [“Adding AWS IAM roles” on page 200](#).

## Limitations

- Storing data to a publicly available target is not supported.
- Storing data to a target on which a lifecycle configuration is set is not supported and may result in data loss.
- Only copies of backup data can be stored to a target with the S3 Glacier Flexible Retrieval or S3 Glacier Deep Archive storage tier. Keep in mind that AWS can charge you additionally for premature removal of data if the retention period specified in your policy is shorter than the recommended (minimum) retention period in AWS.


## Considerations

- You can set up the same target in multiple protection sets.
- Storing data to a target that has Object Lock (WORM) enabled is supported.


## Recommendation


The exclude policy is automatically assigned to the bucket that is added to R-Cloud as a target. It is highly recommended that you do not change this default configuration.

### Accessing the Targets panel

To access the Targets panel, in the navigation pane, click  **Targets**. Alternatively, in the Dashboard panel, click the **Targets** widget title.

## Procedure

1. In the Targets panel, click  **Add**.
2. Select **Amazon S3**, and then click **Next**.
3. Depending on the type of your bucket, click one of the following:
  - **General purpose bucket**
  - **Directory bucket**
4. In the Target field, enter the name of the target.
5. In the Size Quota field, specify the amount of storage space that should be used for storing data (in MiB, GiB, or TiB).

 **Important** The specified amount represents a soft limit, therefore actual usage may exceed it.

6. Use the **Enforce quota** switch to stop running backups if this target reaches its size quota. The backups will start running again after you increase the

size quota of this target or assign a different policy to the entities. Such a policy must use a target with the sufficient size quota.

7. From the Storage Class drop-down, select the storage class that you want to use for storing the data.
8. From the Cloud Account drop-down menu, select the IAM role that you want to be used for performing all operations on the target.

By clicking **Add**, you are automatically redirected to the dialog box that enables you to add the preferred cloud account to R-Cloud, if not already added.

9. *Only if you are adding a directory bucket.* In the Region field, enter the region of your bucket (for example, us-east-1).
10. Click **Save**.

The target is added to the list of targets in the Targets panel. For details on managing targets, see [“Managing targets” on page 165](#).

## Setting up an Azure target

### Prerequisites

- An Azure service principal must be added to R-Cloud. For instructions, see [“Adding Azure service principals” on page 201](#).
- *Only if you plan to store data to an Azure target for which immutability (WORM) is enabled.* In Azure, do the following:
  - Set the Enable version-level immutability support option at the storage account level. For details, see Azure documentation.
  - Assign the Storage Blob Data Owner role to the service principal. For a list of all the required roles, see [“Adding Azure service principals” on page 201](#).


### Limitations

- Storing data to a publicly available target is not supported. Therefore, make sure that the Allow Blob anonymous access setting is disabled in Azure.
- Storing data to a target for which a lifecycle management policy is configured is not supported and may result in data loss.


## Considerations


- You can set up the same target in multiple protection sets.
- *Only if you plan to store data to an Azure target for which immutability (WORM) is enabled.* When backing up data, R-Cloud sets the retention period of the immutable blob data to the retention period defined in the R-Cloud policy.

### Accessing the Targets panel

To access the Targets panel, in the navigation pane, click  **Targets**. Alternatively, in the Dashboard panel, click the **Targets** widget title.

## Procedure

1. In the Targets panel, click  **Add**.
2. Select **Azure**, and then click **Next**.
3. In the Target field, enter the name of the target.
4. In the Size Quota field, specify the amount of storage space that should be used for storing data (in MiB, GiB, or TiB).

 **Important** The specified amount represents a soft limit, therefore actual usage may exceed it.

5. Use the **Enforce quota** switch to stop running backups if this target reaches its size quota. The backups will start running again after you increase the size quota of this target or assign a different policy to the entities. Such a policy must use a target with the sufficient size quota.
6. From the Cloud Account drop-down menu, select the Azure service principal that you want to be used for performing all operations on the target.  
By clicking **Add**, you are automatically redirected to the dialog box that enables you to add the preferred cloud account to R-Cloud, if not already added.
7. Click **Save**.

The target is added to the list of targets in the Targets panel. For details on managing targets, see [“Managing targets” on page 165](#).



## Setting up a Google Cloud target

### Prerequisites

- Your HYCU Managed Service Account (HMSA) must have access to the target.
- *Only if you plan to select a service account other than the HMSA for performing all operations on the target.* The service account must have access to the target.

### Limitations

- Storing data to a publicly available target is not supported.
- Storing data to a target on which a lifecycle configuration is set is not supported and may result in data loss.


### Considerations

- You can set up the same target in multiple protection sets.
- Storing data to a target that has Object Lock (WORM) enabled is supported.
- *Only if you plan to select a service account other than the HMSA for performing all operations on the target that will store the copy of backup data.* The service account must have sufficient permissions also for performing operations on the target that will store primary backup data.
- *For Google Cloud targets with a soft delete policy enabled:* R-Cloud will automatically remove the policy from the target to ensure that your data is stored most cost-efficiently.


### Recommendation

The exclude policy is automatically assigned to the bucket that is added to R-Cloud as a target. It is highly recommended that you do not change this default configuration.

#### Accessing the Targets panel

To access the Targets panel, in the navigation pane, click  **Targets**.  
Alternatively, in the Dashboard panel, click the **Targets** widget title.

### Procedure

1. In the Targets panel, click  **Add**.
2. Select **Google Cloud**, and then click **Next**.

3. In the Target field, enter the name of the target.
4. In the Size Quota field, specify the amount of storage space that should be used for storing data (in MiB, GiB, or TiB).

**ⓘ Important** The specified amount represents a soft limit, therefore actual usage may exceed it.

5. Use the **Enforce quota** switch to stop running backups if this target reaches its size quota. The backups will start running again after you increase the size quota of this target or assign a different policy to the entities. Such a policy must use a target with the sufficient size quota.
6. *Only if you want a service account other than the HMSA to be used for performing all operations on the target.* From the Cloud Account drop-down menu, select the preferred service account.

By clicking **Add**, you are automatically redirected to the dialog box that enables you to add the preferred cloud account to R-Cloud, if not already added.

7. Click **Save**.

The target is added to the list of targets in the Targets panel. For details on managing targets, see [“Managing targets” on page 165](#).

## Setting up an S3 compatible target

### Prerequisite

An S3 compatible account must be added to R-Cloud and it must have the required permissions granted to access the target that you are adding (s3:GetObject, s3:PutObject, s3:DeleteObject, s3:DeleteObjectVersion, s3:GetBucketObjectLockConfiguration, s3:ListBucket, s3:GetBucketLocation, and s3:GetBucketVersioning). For details on adding an S3 compatible account, see [“Adding S3 compatible accounts” on page 203](#).


### Limitations

- Storing data to a publicly available target is not supported.
- Storing data to a target on which lifecycle management is enabled is not supported and may result in data loss.

## Considerations


- You can set up the same target in multiple protection sets.
- Storing data to a target that has Object Lock (WORM) enabled is supported.


### Accessing the Targets panel

To access the Targets panel, in the navigation pane, click  **Targets**.

Alternatively, in the Dashboard panel, click the **Targets** widget title.

## Procedure

1. In the Targets panel, click  **Add**.
2. Select **S3 Compatible**, and then click **Next**.
3. In the Target field, enter the name of the target.
4. In the Size Quota field, specify the amount of storage space that should be used for storing data (in MiB, GiB, or TiB).
 

 **Important** The specified amount represents a soft limit, therefore actual usage may exceed it.
5. Use the **Enforce quota** switch to stop running backups if this target reaches its size quota. The backups will start running again after you increase the size quota of this target or assign a different policy to the entities. Such a policy must use a target with the sufficient size quota.
6. In the Account ID field, enter the ID of the S3 compatible account.
7. From the Service Provider drop-down menu, select your S3 compatible service provider.
8. From the Cloud Account drop-down menu, select the S3 compatible account that you want to be used for performing all operations on the target.
 

By clicking **Add**, you are automatically redirected to the dialog box that enables you to add the preferred cloud account to R-Cloud, if not already added.
9. *Only if you are adding an OVHcloud S3 compatible target.* In the Region field, enter the prefix of the region where the target is located (for example, DE).
10. Click **Save**.

The target is added to the list of targets in the Targets panel. For details on managing targets, see [“Managing targets” on page 165](#).

## Defining your backup strategy

R-Cloud enables you to schedule automatic backups to achieve the optimum level of data protection based on your recovery point objective (RPO) and backup retention requirements. Backups can be scheduled to start each time the specific number of minutes, hours, days, weeks, or months has passed.

When defining your backup strategy, consider the specific needs of your environment and the RPO that represents the maximum period of time for which data loss is considered acceptable. For example, setting the RPO to 24 hours means that your business can tolerate losing only data from the last 24 hours.

Decide which of the following policy approaches best suits the needs of your environment:

Policy approach	Description
Applying a predefined policy	You can use any of the predefined policies to simplify the data protection implementation. For details, see <a href="#">“Taking advantage of predefined policies”</a> below.
Creating a custom policy	If none of the predefined policies meets the needs of your environment, you can create a new policy and tailor it to your needs. For details, see <a href="#">“Creating custom policies”</a> on the next page.

If you consider one of the predefined or custom policies satisfies all data protection goals of your environment, you can set such a policy as default. For details, see [“Setting default policies”](#) on page 48.

## Taking advantage of predefined policies

When establishing a data protection environment, you can take advantage of the predefined policies that provide a fast and convenient way of enabling data protection and cover the most common data protection scenarios.

R-Cloud comes with the following predefined policies:

Predefined policy name	Back up data every...	Keep snapshots for...	Keep copies of backup data for...
platinum	2 hours	1 day	1 week
gold	4 hours	1 day	1 week
silver	12 hours	1 day	1 week
bronze	24 hours	2 days	1 week

If you want to exclude entities from backups, you can use the exclude policy.

### Consideration

Predefined policies use automatically created targets for storing backup data. For details on targets, see [“Setting up targets” on page 27](#).

## Creating custom policies

If the needs of your data protection environment are not covered with any of the predefined policies, you can create a new policy and tailor it to your needs. In this case, besides setting the desired RPO, the retention period for the backup data, and the target, you can also enable one or more additional policy options for optimal policy implementation.

If you plan to protect SaaS applications, Google Kubernetes Engine applications, instances, or buckets, you can also enable one or more of the following policy options:

Policy option	Allows you to...
Backup Window	Start all backup tasks within specified time frames to improve efficiency and avoid an overload of your environment. For details, see <a href="#">“Creating backup windows” on page 43</a> .
Copy <sup>ab</sup>	Create a copy of backup data.
Archiving <sup>a</sup>	Preserve your data for future reference. For details, see <a href="#">“Creating data archives” on page 45</a> .
Labels <sup>b</sup>	Set up automatic policy assignment based on the labels or

Policy option	Allows you to...
	tags added to the SaaS applications, the applications in Google Kubernetes Engine, the instances in Google Compute Engine or Amazon EC2, or the buckets in Google Cloud Storage or Amazon S3.

<sup>a</sup> *For GKE applications:* This policy option is available only for applications using persistent volumes.

<sup>b</sup> This policy option is not available for all SaaS applications. For more information, see the [R-Cloud Module Guides](#).

### Prerequisites

- *Only if you plan to select a manually created target.* The target must be set up. For instructions, see [“Setting up targets” on page 27](#).
- *Only if you plan to enable the Backup Window policy option.* A backup window must exist for the selected R-Cloud protection set. For instructions, see [“Creating backup windows” on page 43](#).
- *Only if you plan to enable the Archiving policy option.* A data archive must exist for the selected R-Cloud protection set. For instructions, see [“Creating data archives” on page 45](#).
- *Only if you plan to enable the Labels policy option.*
  - *Google Cloud specifics:* The HYCU Managed Service Account (HMSA) must have the following roles granted on the projects with the instances that you plan to protect, the clusters on which the GKE applications that you plan to protect are deployed, or the buckets that you plan to protect:
    - Compute Admin (`roles/compute.admin`)
    - Service Account User (`roles/iam.serviceAccountUser`)
    - Storage Admin (`roles/storage.admin`)
    - *Required only if protecting GKE applications.* Kubernetes Engine Admin (`roles/container.admin`)


For instructions on how to grant permissions to service accounts, see Google Cloud documentation.

- The labels that you plan to specify in R-Cloud must be added to SaaS applications, to GKE applications in Google Kubernetes Engine as metadata labels, to instances in Google Compute Engine or Amazon EC2 as labels (preferred) or custom metadata tags, or to buckets in Google Cloud Storage or Amazon S3 as bucket labels.

For instructions on how to do this, see the [R-Cloud Module Guides](#), or the Kubernetes, AWS, or Google Cloud documentation.

#### Limitation

- *Only if you plan to use the same target for the backup data and for the data archive.* The same storage class cannot be used for the backup data and for the data archive.

 **Note** If you select the automatically created target when creating your custom policy and the data archive, R-Cloud will use the same target for both.


For details on the available storage classes for targets, see [“Viewing target information” on page 165](#). For details on the automatic storage class selection during archiving, see [“Creating data archives” on page 45](#).

#### Considerations

- R-Cloud automatically associates the resource with one of the pricing tiers based on the value of the Backup every option that you set in the policy. However, if you are storing data as a snapshot and have enabled the Archiving option, the pricing tier is automatically set to bronze regardless of the specified RPO.
- If you want your data to be stored as a snapshot and on a target, make sure to select the Snapshot backup target type and also enable the Copy policy option.
- *Only if you plan to enable the Labels policy option.*
  - Labels that you specify in policies in R-Cloud must be unique within the selected protection set.
  - When matched, the `hycu-policy` custom metadata tag takes precedence over other labels or tags that might be added to the same SaaS application, to the same application in Google Kubernetes Engine, to the same instance in Google Compute Engine or Amazon EC2, or to the same bucket in Google Cloud Storage or Amazon S3. For more information on the `hycu-policy` tag, see [“Setting up automatic policy assignment” on page 49](#).
- *Only if you plan to store data on a target.* Backup and restore speed depend on the region of the chosen target and the regions of the resources that you are protecting. The optimum speed is achieved when the target and the resources reside in the same region.


- Depending on whether your R-Cloud module supports storing data on a staging target, consider the following:
  - If the R-Cloud module supports storing data on a staging target, SaaS application data cannot be stored to automatically created targets.
  - If the R-Cloud module does not support storing data on a staging target, only backup data can be stored to automatically created targets (and not copies of backup data or archive data).

### Accessing the Policies panel

To access the Policies panel, in the navigation pane, click  **Policies**.

Alternatively, in the Dashboard panel, click the **Policies** widget title.

### Procedure

1. In the Policies panel, click  **New**. The New Policy dialog box opens.
2. Enter a name for your policy and, optionally, its description.
3. Enable the required policy options by clicking them (the Backup policy option is mandatory and therefore enabled by default). Depending on what kind of data you plan to protect, the following policy options are available:


Policy option	Instance and GKE application data protection	SAP HANA application data protection	Bucket data protection	SaaS application data protection
<b>Backup Window</b>	✓	×	✓	✓
<b>Copy</b>	✓ <sup>a</sup>	×	✓	✓ <sup>b</sup>
<b>Archiving</b>	✓ <sup>a</sup>	×	✓	✓
<b>Labels</b>	✓	×	✓	✓ <sup>b</sup>

<sup>a</sup> For GKE applications: This policy option is available only for applications using persistent volumes.

<sup>b</sup> This policy option is not available for all SaaS applications. For more information, see the [R-Cloud Module Guides](#).

4. In the Backup section, do the following:
  - a. In the Backup every field, set the RPO (in months, weeks, days, hours, or minutes).



 **Note** You can set the RPO to 30 minutes in the following cases:

- If you are storing data only as a snapshot.
- If you are storing data as a snapshot and have enabled the Archiving option.

For all other cases, the minimum RPO is one hour.

- In the Retention fields, set a retention period (in months, weeks, or days) for the backup data.
- Select one of the following backup target types:

Backup target type	Next target-related step
<p><i>Applicable only if you are protecting SaaS applications, GKE applications using persistent volumes, or instances. <b>Snapshot</b><sup>a</sup></i></p>	<p><i>Only if protecting Google Cloud instances.</i> Under Snapshot Location, select <b>Regional</b> or <b>Multi-regional</b>.</p> <p>For example, if your instance resides in the <code>us-central1-a</code> zone, with the Multi-regional option selected, a snapshot of the instance is replicated to all us regions, whereas with the Regional option selected, a snapshot is stored only in the <code>us-central1</code> region.</p>
<p><b>Target</b></p>	<p>From the Target drop-down menu, select one of the following for storing data:</p> <ul style="list-style-type: none"> <li>• <b>Automatically selected</b> If you select this option, R-Cloud creates a target and uses it for storing the data. If an automatically created target already exists, it is used instead. For details about automatically created targets, see “<a href="#">Backup target types in R-Cloud</a>” on page 27.</li> <li>• Any available target of your choice</li> </ul>

<sup>a</sup> This backup target type is not available for all SaaS applications. For more information, see the [R-Cloud Module Guides](#).

5. Depending on which policy options you have enabled, do the following:

Policy option	Instructions
Backup Window	<p>In the Backup Window section, from the Backup window drop-down menu, select a backup window for backup tasks.</p> <p>If you do not select a backup window, the Always value is shown, which means that your backups are allowed to run at any time.</p>
Copy <sup>ab</sup>	<p>In the Copy section, do the following:</p> <ol style="list-style-type: none"> <li>Set a retention period (in months, weeks, or days) for the copy of backup data.</li> <li>From the Target drop-down menu, select one of the following for storing the copy of backup data: <ul style="list-style-type: none"> <li><b>Automatically selected</b> If you select this option, R-Cloud creates a target and uses it for storing the data. If an automatically created target already exists, it is used instead. For details about automatically created targets, see <a href="#">“Backup target types in R-Cloud” on page 27</a>.</li> <li>Any available target of your choice</li> </ul> <p>When selecting a preferred target for the copy of backup data, make sure that this target is different from the one you selected for the backup.</p> </li> </ol>
Archiving <sup>a</sup>	<p>In the Archiving section, from the Data archive drop-down menu, select a data archive.</p>
Labels <sup>b</sup>	<p>In the Labels section, enter a label key and value, and then click <b>Add</b>. If required, repeat the action as appropriate.</p> <p>For details on automatic policy assignment, see <a href="#">“Setting up automatic policy assignment” on page 49</a>.</p>

<sup>a</sup> For GKE applications: This policy option is available only for applications using persistent volumes.

<sup>b</sup> This policy option is not available for all SaaS applications. For more information, see the [R-Cloud Module Guides](#).

## 6. Click **Save**.

The policy is created and added to the list of policies. For details on managing policies, see [“Managing policies” on page 163](#).


## Creating backup windows

R-Cloud enables you to define time frames when backup tasks are allowed to start. If you use a backup window, the backup tasks are started only within the hours you specify, which improves effectiveness and prevents overloading your data protection environment. For example, you can schedule your backup tasks to run on non-production hours to reduce the load during peak hours.



You can use backup windows with both predefined policies and custom policies.

**ⓘ Important** When defining a backup window, make sure that the RPO specified in the affected policy can be achieved within this backup window. If the RPO is shorter than any time frame during which backups are not allowed to start, this will result in your GKE applications, instances, and buckets not being compliant with backup requirements.

### Accessing the Policies panel


To access the Policies panel, in the navigation pane, click  **Policies**. Alternatively, in the Dashboard panel, click the **Policies** widget title.

### Procedure

1. In the Policies panel, click  **Backup Window**. The Backup Window dialog box opens.
2. Click  **New**.
3. Enter a name for your backup window and, optionally, its description.
4. From the Time zone drop-down menu, select the time zone for the backup window.

**📄 Note** If the time zone that you selected supports daylight saving time, it is enabled by default.

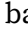

5. Select the days and hours during which backups are allowed to run.

 **Tip** If you click a day label or an hour label, you allow backups to run that whole day or that hourly period for all days of the week. You can also click and drag to quickly select a time frame that includes your preferred days and hours.


The selected time frames are displayed in the Time frames field. If you want to delete any of the selected time frames, pause on it, and then click **x**.

6. Click **Save**.

7. Click **Close**.

You can later edit any of the existing backup windows (click  **Edit** and make the required modifications) or delete the ones that you do not need anymore (click  **Delete**).

After you create a backup window, you can do the following:

- Specify the backup window when creating a new policy. For details, see [“Creating custom policies” on page 37](#).
- Assign the backup window to an existing policy. To do so, select the policy, click  **Edit**, and then make the required modifications.

## Example

You have selected the bronze policy and allowed new backup tasks to run on weekdays from 6 PM to 6 AM (Eastern Time), and on Saturday and Sunday all day long.

**Backup Window > New** ? X

Name: non-production-hours

Description - Optional: weekdays from 6 PM to 6AM, Saturdays and Sundays all day

Time zone: Etc/GMT+5 (UTC-05:00)

00 01 02 03 04 05 06 07 08 09 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24

MON: 00:00-06:00, 18:00-24:00

TUE: 00:00-06:00, 18:00-24:00

WED: 00:00-06:00, 18:00-24:00

THU: 00:00-06:00, 18:00-24:00

FRI: 00:00-06:00, 18:00-24:00

SAT: 00:00-24:00

SUN: 00:00-24:00

Time Frames: Clear All

MON 00:00 - 06:00 X MON 18:00 - 24:00 X TUE 00:00 - 06:00 X TUE 18:00 - 24:00 X WED 00:00 - 06:00 X WED 18:00 - 24:00 X

THU 00:00 - 06:00 X THU 18:00 - 24:00 X FRI 00:00 - 06:00 X FRI 18:00 - 24:00 X SAT 00:00 - 24:00 X SUN 00:00 - 24:00 X

Close Back Save

In this case, the backup tasks can be run every 24 hours at any point of time within the specified time frames.


## Creating data archives

R-Cloud enables you to create archives of your protected data and keep them for a longer period of time. By archiving data, the data is stored for future reference on a daily, weekly, monthly, or yearly basis. Your data is isolated from current activity and safely stored in a secure cloud archive location.



### Prerequisite

*Only if you plan to select a manually created target for the data archive.* The target must be set up. For instructions, see [“Setting up targets” on page 27.](#)

## Accessing the Policies panel


To access the Policies panel, in the navigation pane, click  **Policies**. Alternatively, in the Dashboard panel, click the **Policies** widget title.

## Procedure

1. In the Policies panel, click  **Archiving**. The Archiving dialog box opens.
2. Click  **New**.
3. Enter a name for your data archive and, optionally, its description.
4. Add any of the following archiving options to the list of the enabled options by clicking it:

<b>Daily</b>	Allows you to create a daily archive of data.
<b>Weekly</b>	Allows you to create a weekly archive of data.
<b>Monthly</b>	Allows you to create a monthly archive of data.
<b>Yearly</b>	Allows you to create a yearly archive of data.

5. In the Start at fields, specify the hour and the minute when the archiving task should start.
6. From the Time zone drop-down menu, specify the appropriate time zone.
7. *Only if you have enabled the Weekly, Monthly, and/or Yearly archiving option.* Specify when to archive data.
8. For each enabled archiving option, do the following:
  - a. In the Retention box, set the retention period to be used.

 **Note** Make sure that the retention period is longer than the RPO to prevent the data archive from expiring before a new backup is performed.

- b. From the Target drop-down menu, select one of the following for storing the data archive:
      - **Automatically selected**  
If you select this option, R-Cloud creates a target and uses it for storing the data. If an automatically created target already exists, it is used instead. For details about automatically created targets, see [“Backup target types in R-Cloud” on page 27](#).
      - Any available target of your choice



- c. From the Storage class drop-down menu, select the storage class that you want to use for storing the data archive.

If you select the **Automatically selected** option, the storage class is automatically selected based on the selected target type and the retention period:

Retention period	Assigned storage classes
Less than 1 month	<ul style="list-style-type: none"> <li>• Amazon S3: S3 Standard</li> <li>• Azure: Hot tier</li> <li>• Google Cloud: Standard</li> <li>• S3 Compatible: S3 Standard</li> </ul>
1 month to less than 3 months	<ul style="list-style-type: none"> <li>• Amazon S3: S3 Standard-IA</li> <li>• Azure: Cool tier</li> <li>• Google Cloud: Nearline</li> <li>• S3 Compatible: S3 Standard</li> </ul>
3 months to less than 1 year	<ul style="list-style-type: none"> <li>• Amazon S3: S3 Glacier</li> <li>• Azure: Cold tier</li> <li>• Google Cloud: Coldline</li> <li>• S3 Compatible: S3 Standard</li> </ul>
More than 1 year	<ul style="list-style-type: none"> <li>• Amazon S3: S3 Deep Archive</li> <li>• Azure: Archive tier</li> <li>• Google Cloud: Archive</li> <li>• S3 Compatible: S3 Standard</li> </ul>


For details on storage classes, see the relevant cloud provider documentation.

9. Click **Save**.

You can later edit any of the existing data archives (click  **Edit** and make the required modifications) or delete the ones that you do not need anymore (click  **Delete**). Keep in mind that you cannot modify a target on which an archiving task is in progress.

After you create a data archive, you can do the following:

- Specify a data archive when creating a new policy. For details, see [“Creating custom policies” on page 37](#).

- Include the data archive into an existing policy. To do so, select the policy, click  **Edit**, and then make the required modifications.

## Setting default policies


You can select one of the predefined or custom policies to be the default policy for your data protection environment. When you set the default policy, depending on your choice, the default policy will be assigned to one of the following:

- Only newly discovered resources.
- Both newly discovered resources and all existing resources that do not have an assigned policy yet.

### Consideration


Setting a default policy is overridden by assigning policies automatically. For more information, see [“Setting up automatic policy assignment” on the next page](#).

#### Accessing the Policies panel

To access the Policies panel, in the navigation pane, click  **Policies**.



Alternatively, in the Dashboard panel, click the **Policies** widget title.

### Procedure

1. In the Policies panel, select the policy that you want to set as the default one, and then click  **Set Default**. The Set Default Policy dialog box opens.
2. Depending on the resources to which you want the default policy to be assigned, select one or more check boxes:
  - **Instances**
  - **Applications**
  - **Buckets**
  - **SaaS**

The default policy will be assigned to all newly discovered resources.

3. Enable the **Assign to resources without policy** switch if you want the default policy to be assigned also to all selected resources that do not have an assigned policy yet.
4. Click **Save**.

The default policy is represented by the  icon. If you later decide not to use this policy as the default one, click  **Clear Default**. Keep in mind that by



doing so, you do not unassign this policy from the resources to which it was assigned.

## Setting up automatic policy assignment

You can set up automatic assignment of policies to SaaS applications, Google Kubernetes Engine (GKE) applications, instances, or buckets by using one of the following methods:

### Method 1 (labels, tags, or metadata)

Entities	Instructions
SaaS applications <sup>a</sup>	Add labels or tags to SaaS applications, and then specify the corresponding keys and values in R-Cloud policies. For details, see <a href="#">“Creating custom policies” on page 37</a> .
GKE applications	Add metadata labels to applications in Google Kubernetes Engine, and then specify the corresponding keys and values in R-Cloud policies. For details, see <a href="#">“Creating custom policies” on page 37</a> .
Instances	Add tags to instances in Amazon EC2, or labels (preferred) or custom metadata to instances in Google Compute Engine, and then specify the corresponding keys and values in R-Cloud policies. For details, see <a href="#">“Creating custom policies” on page 37</a> .
Buckets	Add labels to buckets in Google Cloud Storage, or tags to buckets in Amazon S3, and then specify the corresponding keys and values in R-Cloud policies. For details, see <a href="#">“Creating custom policies” on page 37</a> .

<sup>a</sup> Setting up automatic policy assignment is not supported for all SaaS applications. For more information, see the [R-Cloud Module Guides](#).

### Method 2 (the hycu-policy tag)

Entities	Instructions
SaaS applications <sup>a</sup>	Add the hycu-policy tag to SaaS applications, applications in Google Kubernetes Engine, instances in Amazon EC2 or Google Compute Engine, or buckets in Amazon S3 or Google

Entities	Instructions
	Cloud Storage. Use the following key/value pair:
GKE applications	Key: <code>hycu-policy</code>
Instances	Value: <code>&lt;PolicyName&gt;</code>
Buckets	In this case, <code>&lt;PolicyName&gt;</code> is the name of an R-Cloud policy (for example, <code>gold</code> ).

<sup>a</sup> Setting up automatic policy assignment is not supported for all SaaS applications. For more information, see the [R-Cloud Module Guides](#).

The corresponding policies are automatically assigned to the SaaS applications, GKE applications, instances, or buckets during the next entity synchronization in R-Cloud.

### Prerequisites

- All relevant prerequisites that apply also for manual policy assignment must be fulfilled. For details, see [“Backing up SaaS applications” on page 60](#), [“Backing up Google Kubernetes Engine applications” on page 77](#), [“Backing up instances” on page 92](#), or [“Backing up buckets” on page 145](#).
- *For Google Kubernetes Engine applications:* The resource objects for which you want to set up automatic policy assignment must be deployed as applications (the resource object of `kind: Application` must be defined in the application deployment).


### Considerations

- Assigning policies automatically takes precedence over assigning policies manually or setting a default policy. This means that the label, the tag, or the metadata added to the preferred SaaS application, GKE application, instance, or bucket defines which policy is assigned to it, even if the same entity already has an assigned policy.
- If you want to assign a new policy to a SaaS application, a GKE application, an instance, or a bucket for which automatic policy assignment has been set up, do one of the following:
  - Define new tags, labels, or metadata as described in this section.
  - Assign the policy to the entity as described in [“Backing up SaaS applications” on page 60](#), [“Backing up Google Kubernetes Engine applications” on page 77](#), [“Backing up instances” on page 92](#), or [“Backing up buckets” on page 145](#).

up buckets” on page 145. In this case, the manually assigned policy will not be overridden by the automatically assigned one again.

## Enabling access to data

In the following data protection scenarios, you must enable access to data by assigning credential groups to instances in R-Cloud:

Guest OS	Data protection scenario
any	<p>You plan to use the default view when restoring individual files or folders.</p> <p> <b>Note</b> If access to the data is enabled at the time of the restore (and not before the backup), individual files or folders can only be restored by using the filesystem view.</p> <p>For details about the views that are available during the restore, see “Restoring individual files or folders” on page 132.</p>
Linux	<ul style="list-style-type: none"> <li>You plan to protect SAP HANA applications.</li> <li>You plan to use pre-snapshot or post-snapshot scripts and run them with a user account that you specify.</li> </ul>
Windows	You plan to use pre-snapshot or post-snapshot scripts.

## Enabling access to instances

To enable access to instances, you must perform the following tasks:

Task	Instructions
1. Configure the port settings on the instances.	“Configuring port settings on instances” on the next page
2. Configure credential groups.	“Configuring credential groups” on the next page
3. Assign credential groups to the instances.	“Assigning credential groups” on page 55

## Configuring port settings on instances

The following table lists the inbound ports that you must open on each instance by configuring and applying a network firewall rule:

Guest OS	Network service protocol	Port	Transport protocol
Linux	SSH	22	N/A
Windows	WinRM	5986	HTTPS
		5985	HTTP

For instructions on how to configure and apply the network firewall rule, see AWS or Google Cloud documentation.

**Note** For Google Cloud instances: Optionally, you can make the network firewall rule more restrictive so that it allows network traffic only from legitimate sources and to legitimate targets. To do so, add `hycu-network-tag` to the network firewall rule.

## Configuring credential groups

### Prerequisites

- A user account with sufficient privileges must be configured within each instance:
  - For Windows: User from the Administrators group
  - For Linux: User with sudo privileges and the NOPASSWD option set

For details on how to do this, see AWS or Google Cloud documentation.

- For Linux instances:
  - Ensure the following within the instance:
    - The specified user account must be a member of the sudo user group.
    - The following line must be included in the `/etc/sudoers` file:

```
<UserName> ALL=(ALL) NOPASSWD:
/opt/hycu/tmp/discoverLinuxMountPointDiskMapping.sh*,
/opt/hycu/tmp/hycuflr*, /usr/bin/mount, /usr/bin/umount,
/usr/bin/mkdir, /usr/bin/rmdir, /usr/bin/rm
```

- Only if you want R-Cloud to access the instance by using a specific user account with password authentication. The SSH server must be configured

to allow password authentication for signing-in on to the instance.


- *For Ubuntu 22.04 instances that have RSA key-based authentication configured:*

You must add the `PubkeyAcceptedKeyTypes=+ssh-rsa` parameter to the `/etc/ssh/sshd_config` file, and then restart the SSH service by running the `systemctl restart ssh.service` command.



### Limitation

*Only if you use the SSH protocol with public key authentication.* If keys are generated with PuttyKeyGen or ssh-keygen using the legacy PEM format, only DSA and RSA keys are supported.

#### Accessing the Instances panel

To access the Instances panel, in the navigation pane, click  **Instances**.

### Procedure

1. In the Instances panel, select the instance to which you want to assign a credential group.
2. Click  **Credentials**. The Credential Groups dialog box opens.
3. Click  **New**.
4. In the Credential group name field, enter a name for the credential group.
5. From the Protocol drop-down menu, select one the following protocol options:



Protocol option	Instructions
<b>Automatic</b>	<p>Select this option if you want R-Cloud to automatically select a protocol for accessing the instance—the SSH protocol (TCP port 22) or the WinRM protocol (TCP port 5985, HTTP transport)—, and then enter the user name and password of a user account that has required permissions to access the instance.</p> <p>Use the following format for the user name:</p> <ul style="list-style-type: none"> <li>• <b>Linux:</b> <code>&lt;LocalOrDomainUserName&gt;</code></li> <li>• <b>Windows:</b> <code>&lt;LocalUserName&gt;</code>, <code>&lt;Domain&gt;\&lt;DomainUserName&gt;</code>,</li> </ul>

Protocol option	Instructions			
SSH	<code>&lt;DomainUserName&gt;@&lt;Domain&gt;</code>			
	<p>Select this option if you want to use the SSH protocol for accessing the instance, and then do the following:</p> <ol style="list-style-type: none"> <li>a. In the Port field, enter the SSH server port number.</li> <li>b. From the Authentication drop-down menu, select the type of authentication you want to be used, and then provide the required information: <table border="1" data-bbox="564 685 1323 1841" style="width: 100%; border-collapse: collapse;"> <tr> <td data-bbox="564 685 807 925" style="width: 30%; text-align: center; vertical-align: top;"><b>Password authentication</b></td> <td data-bbox="807 685 1323 925">Enter the user name (in the <code>&lt;LocalOrDomainUserName&gt;</code> format) and password of a user account that has required permissions to access the instance.</td> </tr> <tr> <td data-bbox="564 925 807 1841" style="width: 30%; text-align: center; vertical-align: top;"><b>Public key authentication</b></td> <td data-bbox="807 925 1323 1841"> <p>Do the following:</p> <ol style="list-style-type: none"> <li>i. Enter the user name (in the <code>&lt;LocalOrDomainUserName&gt;</code> format) and password of a user account that has required permissions to access the instance.</li> <li>ii. Click <b>Browse</b>. Browse for and then select the file with the private key and click <b>Open</b>. For information on how to obtain the private key, see Google Cloud or AWS documentation.</li> <li>iii. <i>Only if the private key is encrypted.</i> Enter the private key passphrase.</li> </ol> <div style="border-left: 2px solid purple; padding-left: 10px; margin-top: 10px;"> <p><b>ⓘ Important</b> This selection is mandatory in cases where the SSH server is configured to use public key authentication.</p> </div> </td> </tr> </table> </li> </ol>	<b>Password authentication</b>	Enter the user name (in the <code>&lt;LocalOrDomainUserName&gt;</code> format) and password of a user account that has required permissions to access the instance.	<b>Public key authentication</b>
<b>Password authentication</b>	Enter the user name (in the <code>&lt;LocalOrDomainUserName&gt;</code> format) and password of a user account that has required permissions to access the instance.			
<b>Public key authentication</b>	<p>Do the following:</p> <ol style="list-style-type: none"> <li>i. Enter the user name (in the <code>&lt;LocalOrDomainUserName&gt;</code> format) and password of a user account that has required permissions to access the instance.</li> <li>ii. Click <b>Browse</b>. Browse for and then select the file with the private key and click <b>Open</b>. For information on how to obtain the private key, see Google Cloud or AWS documentation.</li> <li>iii. <i>Only if the private key is encrypted.</i> Enter the private key passphrase.</li> </ol> <div style="border-left: 2px solid purple; padding-left: 10px; margin-top: 10px;"> <p><b>ⓘ Important</b> This selection is mandatory in cases where the SSH server is configured to use public key authentication.</p> </div>			

Protocol option	Instructions
<b>WinRM</b>	<p>Select this option to use the WinRM protocol for instance access and to enable the credential group adjustment for the actual WinRM server configuration.</p> <ol style="list-style-type: none"> <li>From the Transport drop-down menu, select the transport protocol of the WinRM server in the instance.</li> <li>In the Port field, enter the WinRM server port number.</li> <li>Enter the user name (in the <code>&lt;LocalOrDomainUserName&gt;</code> format (for Google Cloud) or <code>&lt;localuser&gt;</code>, <code>&lt;domain&gt;\&lt;user&gt;</code>, or <code>&lt;user&gt;@&lt;domain&gt;</code> format for AWS) and the password of a user account that has required permissions to access the instance.</li> </ol>

6. Click **Save**.

The name of the credential group appears in the list of credential groups in the Credential Groups dialog box.

You can also edit any of the existing credential groups (select a credential group, click  **Edit**, and then make the required modifications) or delete the ones that you do not need anymore (select a credential group, and then click  **Delete**).


## Assigning credential groups

You can assign credential groups to instances by using the R-Cloud web user interface or by using labels or metadata tags. Depending on how you want to assign the credential groups to the instances, see the following sections:

- [“Assigning credential groups by using the R-Cloud web user interface” below](#)
- [“Assigning credential groups by using labels or metadata tags” on the next page](#)

### Assigning credential groups by using the R-Cloud web user interface

#### Procedure

1. In the Instances panel, select the instances to which you want to assign a credential group.
2. Click  **Credentials**. The Credential Groups dialog box opens.

- From the list of credential groups, select the credential group that you want to assign to the selected instances, and then click **Assign**.

The name of the assigned credential group appears in the Credential group column of the Instances panel. R-Cloud performs instance and application discovery after you assign the credentials to the instance. The Discovery status in the Instances and Applications panels is updated accordingly.

**Tip** If several instances share the same user name and password, you can use multiple selection to assign the same credential group to them.

To unassign a credential group from an instance, in the Instances panel, select the instance, click **Credentials**, and then click **Unassign**.

### Assigning credential groups by using labels or metadata tags

You can assign a credential group to an instance by adding the `hycu-credential-group` tag to the instance in Amazon EC2 or Google Compute Engine as a label or a metadata tag. Use the following name/value pair:

Name	Value
<code>hycu-credential-group</code>	<p><code>&lt;CredentialGroupName&gt;</code>            In this case, <code>&lt;CredentialGroupName&gt;</code> is the name of the credential group that you want to assign to the instance.</p>


The credential group is automatically assigned to the instance during the next instance synchronization in R-Cloud.




# Chapter 4

## Protecting SaaS applications

R-Cloud enables you to protect your SaaS application data with fast and reliable backup and restore operations.

For a list of SaaS applications for which R-Cloud provides data protection through R-Cloud modules, see the HYCU Marketplace. To access the Marketplace panel, in the navigation pane, click  **Marketplace**. For more information on the HYCU Marketplace, see “[Navigating the HYCU Marketplace](#)” on page 204.



 **Tip** You can also search for a SaaS application in the HYCU Marketplace by entering its name, or part of the name in the Search field or filter the R-Cloud modules by platform or by category.

### Prerequisite

Before you start protecting data, you must be familiar with all the prerequisites, limitations, considerations, and recommendations described for each SaaS application individually in the [R-Cloud Module Guides](#).

### Consideration

Keep in mind that the role you have assigned determines what kind of actions you can perform. For details on roles, see “[Managing roles](#)” on page 225.

 **Note** R-Cloud performs automatic synchronization of SaaS applications at periodic intervals. However, you can at any time update the list of SaaS applications also manually by clicking  **Refresh**.

For details on how to efficiently protect SaaS application data, see the following sections:

- “[Configuring SaaS application backup options](#)” on the next page
- “[Backing up SaaS applications](#)” on page 60
- “[Restoring SaaS applications](#)” on page 61


# Configuring SaaS application backup options

Before you start protecting SaaS applications, you can adjust SaaS application protection to the needs of your data protection environment by configuring backup options.

**ⓘ Important** Configuring backup options is not supported for all types of SaaS applications. Additionally, the list of available backup options varies depending on the type of your SaaS application.


Backup option	Description
Exclude resources	Enables you to specify one or more resources to be excluded from the backup.
Options	Enables you to use backup options specific to each SaaS application or SaaS application resource (for example, if you are protecting Google Cloud SQL, you can set the offload option that enables R-Cloud to delegate the export operation to a separate temporary instance).
Temporary instance configuration	<p>Enables you to specify the source, the region, and the subnet where you want R-Cloud to create a temporary instance during the backup. If the specified source is an AWS account, you can also select a security group for the temporary instance.</p> <p><b>ⓘ Important</b> If you do not configure this backup option, R-Cloud by default creates the temporary instance in your AWS account or Google Cloud project after you set up a target in R-Cloud or add a source to R-Cloud.</p>


## Accessing the SaaS panel

To access the SaaS panel, in the navigation pane, click  **SaaS**.

## Procedure

1. In the SaaS panel, select the SaaS application or the resource for which you want to configure backup options.

2. Click  **Configuration**. The SaaS Configuration dialog box opens.
3. Depending on what you want to do, perform the required action:

I want to...	Instructions
Exclude resources from the backup.	On the Exclude Resources tab, select the resources that you want to exclude from the backup.
Use a backup option specific to my SaaS application or resource.	On the Options tab, specify which of the available backup options you want to use and provide the required information.
Specify the source, the region, the subnet, and the security group for a temporary instance.	<p>On the Temporary Instance Configuration tab, do the following:</p> <ol style="list-style-type: none"> <li>a. From the Source drop-down menu, select the source for the temporary instance. <ul style="list-style-type: none"> <li> <b>Important</b> If the type of the source that you select for the temporary instance differs from the source where the target specified in the R-Cloud policy resides, this may result in data egress charges.</li> </ul> </li> <li>b. From the Region drop-down menu, select the preferred region.</li> <li>c. From the Subnet drop-down menu, select the preferred subnet.</li> <li>d. <i>For AWS accounts:</i> Optionally, from the Security Group drop-down menu, select the preferred security group. By default, the temporary instance is created in the default security group of the preferred subnet.</li> </ol>

4. Click **Save**.

# Backing up SaaS applications

With R-Cloud, you can back up your SaaS application data securely and efficiently.


## Limitations

- A policy that has Snapshot specified as the backup target type can be assigned only to specific SaaS applications and/or resources. For information whether such a policy can be assigned to your SaaS application and/or its resources, see the [R-Cloud Module Guides](#).
- *Only if one or more SaaS applications that you plan to back up reside on a different cloud platform than the target defined in the policy.* The RPO of the assigned policy must not be shorter than 24 hours.
- *Only if your R-Cloud module supports storing data on a staging target.* The staging target that you add to R-Cloud when adding an R-Cloud module, and the target that is defined in the policy that is assigned to the related SaaS application must reside on the same cloud platform.
- Depending on whether your R-Cloud module supports storing data on a staging target, the following limitations apply:
  - If the R-Cloud module supports storing data on a staging target, SaaS application data cannot be stored to automatically created targets.
  - If the R-Cloud module does not support storing data on a staging target, only backup data can be stored to automatically created targets (and not copies of backup data or archive data).

## Consideration


*Only if your R-Cloud module supports storing data on a staging target.* If you use an automatically created staging target, the HMSA must be configured to perform all operations on the target specified in the policy that is assigned to the related SaaS application. Alternatively, the same cloud account must be configured to perform all operations on both targets (the staging target and the target specified in the policy that is assigned to the related SaaS application).


### Accessing the SaaS panel

To access the SaaS panel, in the navigation pane, click  **SaaS**.

## Procedure

1. Select the SaaS applications and/or resources that you want to back up.

 **Note** If you want to narrow down the list of displayed SaaS applications, use the filtering options as described in [“Filtering and sorting data” on page 181](#).

2. Click  **Set Policy**. The Set Policy dialog box opens.
3. From the list of available policies, select the preferred policy.
4. Click **Assign** to assign the policy to the selected SaaS applications and/or resources.

After you assign a policy to a SaaS application, a backup task starts immediately. Subsequent backups are scheduled according to the values defined in the policy.

If required, you can also perform a manual backup of individual SaaS applications and/or resources at any time. For details, see [“Performing manual backups” on page 187](#).

## Restoring SaaS applications

R-Cloud enables you to restore an entire SaaS application or its resources to a specific point in time.

### Prerequisite

*Only if you plan to restore your SaaS application data to a different source.* Your SaaS application must support restoring data to a different source and more than one source (that is, more than one R-Cloud module of the same type) must be added to R-Cloud.

### Limitation


If your data is stored as a snapshot, you cannot restore it to a different source.

### Considerations

- Only one restore task can run at the same time for a SaaS application or its resource.


- *Only if a SaaS application resource is deleted from cloud.* If the deleted resource has at least one valid restore point available in R-Cloud, it is considered protected and its status is Protected deleted.


### Accessing the SaaS panel


To access the SaaS panel, in the navigation pane, click  **SaaS**.

### Procedure


1. In the SaaS panel, click the SaaS application or the resource that you want to restore. The Detail view appears at the bottom of the screen.

 **Note** The Detail view appears only if you click a SaaS application. Selecting the check box before the name of the SaaS application does not open the Detail view.

2. In the Detail view, select the preferred restore point.
3. Click  **Restore**. The Restore dialog box opens.
4. *Only if the selected SaaS application supports restoring data to a different source and more than one source has been added to R-Cloud.*
  - a. Depending on where you want to restore your SaaS application data, do one of the following:
    - Select **Restore to same source** to restore your SaaS application or its resource to the original R-Cloud module.
    - Select **Restore to different source** to restore your SaaS application or its resource to a different R-Cloud module of the same type.
  - b. Click **Next**.

 **Important** For information whether your SaaS application supports restoring data to a different source, see the [R-Cloud Module Guides](#).

5. Select at what level you want to restore your SaaS application or resource (for example, at an instance, database, attachment, or story level), and then click **Next**.

 **Important** The list of restore options varies depending on the type of your SaaS application. For details about restore options for your SaaS application, see the [R-Cloud Module Guides](#).

6. Select the SaaS application data that you want to restore, and then click **Next**.

7. From the Restore from drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among which you can select:
  - **Automatic:** This option ensures the fastest and most cost-effective restore.
  - **Backup (Snapshot)**
  - **Backup (Target)**
  - **Copy**
  - **Archive—(daily, weekly, monthly, yearly)**
8. *Only if restore options specific to your SaaS application are available.* Specify which of the available restore options you want to use and provide the required information.
9. Click **Restore**.

# Chapter 5

## Protecting applications

R-Cloud enables you to protect your Google Cloud application data with fast and reliable backup and restore operations. After you prepare your application for data protection and back it up, you can choose to restore either the whole application or only specific application items. For a list of supported applications, see the *HYCU R-Cloud Compatibility Matrix*.

### Prerequisite

Cloud Resource Manager API, Compute Engine API, Cloud Identity and Access Management API, Cloud Billing API, and Cloud Storage API must be enabled on the Google Cloud projects that contain the instances and Google Kubernetes Engine clusters on which the applications that you want to protect are running. For instructions on how to enable APIs, see Google Cloud documentation.

### Considerations

- Keep in mind that the role you have assigned determines what kind of actions you can perform. For details on roles, see [“Managing roles” on page 225](#).
- R-Cloud uses an external IP address to access Google Cloud APIs if Private Google Access is disabled on subnets. If your data protection environment requires the use of an internal IP address, make sure Private Google Access is enabled on subnets. For details, see Google Cloud documentation.

Depending on what type of Google Cloud applications you plan to protect, follow the required instructions:

I plan to protect...	Instructions
SAP HANA applications	<a href="#">“Protecting SAP HANA applications” on the next page</a>
Google Kubernetes Engine applications	<a href="#">“Protecting Google Kubernetes Engine applications” on page 72</a>



# Protecting SAP HANA applications

Protecting SAP HANA application data consist of the following tasks:

Task	Instructions
Preparing SAP HANA applications for data protection, which includes enabling access to application data and configuring backup options.	<a href="#">“Preparing for SAP HANA application protection”</a> below
Backing up SAP HANA applications.	<a href="#">“Backing up SAP HANA applications”</a> on page 68
Restoring SAP HANA application data.	<a href="#">“Restoring SAP HANA applications”</a> on page 70

## Preparing for SAP HANA application protection

Before you start protecting SAP HANA applications, you must prepare your environment for application data protection. Preparing your environment for SAP HANA application data protection includes the following tasks:

Task	Instructions
1. <i>Mandatory</i> . Make sure R-Cloud can access applications that you want to protect.	<a href="#">“Enabling access to application data”</a> below
2. <i>Optional</i> . Configure SAP HANA application backup options.	<a href="#">“Configuring SAP HANA application backup options”</a> on page 67


### Enabling access to application data

After you assign credentials to instances as described in [“Enabling access to data” on page 51](#), the process of application discovery starts automatically. When the application discovery task completes, the discovered applications are listed in the Applications panel.

Each discovered application has one of the following statuses:


Discovery status	Description
✔	R-Cloud can access discovered applications that you want to protect with instance credentials. However, if your applications require database-level authentication, you must make sure to provide also application-specific credentials before you can start protecting your data. In this case, follow the procedure described in this section. Otherwise, you can continue with protecting application data as described in <a href="#">“Backing up SAP HANA applications”</a> on page 68.
✘	<p>The instance credentials do not have proper permissions and R-Cloud cannot access applications. To enable R-Cloud to access the applications, reassign credentials to instances so that they have proper permissions. For instructions on how to assign credentials to an instance, see <a href="#">“Enabling access to data”</a> on page 51.</p> <p>After the discovery status of your application is ✔, make sure to provide also application-specific credentials if your application requires database-level authentication. In this case, follow the procedure described in this section.</p>

### Accessing the Applications panel

To access the Applications panel, in the navigation pane, click 

#### **Applications.**

#### Procedure

1. In the Applications panel, select the applications that you want to back up.
2. Click  **Configuration**. The Application Configuration dialog box opens.
3. On the Credentials tab, make sure the **Use instance credentials** switch is disabled, and then enter credentials for a user account with required permissions and access to the applications.
4. Click **Save**.

You can continue with protecting application data as described in [“Backing up SAP HANA applications”](#) on page 68.

You can later unassign the credentials from an instance or delete the instance credentials that you do not need anymore. For details, see [“Enabling access to](#)

data” on page 51. Keep in mind that you can do this only if the discovered applications running on the instance do not have assigned policies or available restore points. Therefore, before unassigning or deleting credentials, make sure to unassign policies or mark restore points as expired.


## Configuring SAP HANA application backup options

Before you start protecting SAP HANA applications, you can adjust application protection to the needs of your data protection environment by configuring backup options.

### Backup options

Backup option	Description
Temporary instance configuration	Enables you to specify the region, the zone, and the subnet where you want R-Cloud to create a temporary instance during the backup. By default, the temporary instance is created in the original project of the application.
Backups	Enables you to configure the backup chain length. In this case, a new backup chain is started when the number of the full and subsequent incremental backups in a backup chain exceeds the value you specify. The default value is 7.

### Accessing the Applications panel


To access the Applications panel, in the navigation pane, click .

#### **Applications.**

### Prerequisites

- *Only if you plan to configure backup options for multiple applications.* All applications must have the same values set for each option that you plan to configure.
- *Only if you plan to specify the temporary instance location and subnet.* VPC Network Peering must be configured. For details on how to configure VPC Network Peering, see Google Cloud documentation.

## Procedure

1. In the Applications panel, select the applications for which you want to configure backup options.
2. Click  **Configuration**. The Application Configuration dialog box opens.
3. Depending on what you want to do, perform the required action:

I want to...	Instructions
Specify the temporary instance location and subnet.	<p>On the Temporary instance configuration tab, do the following:</p> <ol style="list-style-type: none"> <li>a. From the Region drop-down menu, select the preferred region.</li> <li>b. From the Zone drop-down menu, select the preferred zone.</li> <li>c. From the Subnet drop-down menu, select the preferred subnet. By default, the temporary instance is created in the default subnet of the preferred region and zone.</li> </ol>
Configure the backup chain length.	On the Backups tab, in the Backup chain length field, specify when you want a new backup chain to be started.

4. Click **Save**.

## Backing up SAP HANA applications

With R-Cloud, you can back up your SAP HANA application data securely and efficiently.

### Prerequisites

- *Only if you plan to back up applications running on instances that use Shared VPC networks.* Your user account or the HYCU Managed Service Account (HMSA) must be granted the following permissions in the Shared VPC host project: `compute.firewalls.list`, `compute.networks.list`, `compute.networks.get`, `compute.subnetworks.list`, `compute.subnetworks.use`, and `compute.subnetworks.get`.
- The minimum required SAP HANA privileges of the configured SAP HANA database user must be `BACKUP ADMIN` and `CATALOG READ`.

- The configured SAP HANA database user must have access permissions to all databases that are being backed up.
- *For SAP HANA systems with the same SID:* A separate target must be configured for each SAP HANA system.


### Limitation

SAP HANA application data can be stored only to Google Cloud targets.

### Considerations

- Application data can be stored only to manually created targets, and not to automatically created targets or as a snapshot.
- Before each backup task, the Backint agent is configured to use the service account that you specified when setting up the target for storing backup data. If you are using the default instance service account, the access scope for storage must be Read Write. For details on Cloud API access scopes, see Google Cloud documentation.
- During each backup task, R-Cloud activates also the automatic backup of logs and backup catalogs using the Backint agent.
- *Only if you have set up SAP HANA system replication.* You can assign the policy only to the primary system. In the event of a failover, after the secondary system takes over from the primary system, make sure to assign the policy to the new primary system.


### Accessing the Applications panel


To access the Applications panel, in the navigation pane, click 

#### **Applications.**

### Procedure

1. In the Applications panel, select the applications that you want to back up.

 **Tip** To narrow down the list of displayed applications, you can use the filtering options as described in “[Filtering and sorting data](#)” on [page 181](#).

2. Click  **Set Policy**. The Set Policy dialog box opens.
3. From the list of available policies, select the preferred policy.
4. Click **Assign** to assign the policy to the selected applications.

After you assign a policy to an application, a backup task starts immediately. Subsequent backups are scheduled according to the values defined in the policy.

If required, you can also perform a manual backup of any application at any time. For details, see [“Performing manual backups” on page 187](#).

## Restoring SAP HANA applications

R-Cloud enables you to restore either a whole application or only individual application items to a specific point in time.

### Prerequisites

- The instance to which you are restoring application data must be up and running.
- *Only if you plan to restore applications running on instances that use Shared VPC networks.* Your user account or the HYCU Managed Service Account (HMSA) must be granted the following permissions in the Shared VPC host project: `compute.firewalls.list`, `compute.networks.list`, `compute.networks.get`, `compute.subnetworks.list`, `compute.subnetworks.use`, and `compute.subnetworks.get`.
- *Only if you are restoring SAP HANA tenant databases without a system database.*
  - Tenant databases that you want to restore must exist.
  - A system database must be online and tenant databases must be stopped. For details on how to stop the tenant databases, see SAP HANA documentation.

### Limitation

You can restore application data only to an instance that belongs to the currently selected protection set and on which an SAP HANA application has already been discovered.

### Considerations


- When restoring data, the automatic backup of backup catalogs using the Backint agent is disabled until the next backup task.
- *Only if you plan to enable the Clear logs option for the selected restore point.* Any subsequent restore using a restore point belonging to the same backup chain will also require the Clear logs option to be enabled.

- After restoring only a system database, make sure to start all the tenant databases.

### Recommendation

After restoring data, it is recommended to perform a full backup of data.


#### Accessing the Applications panel


To access the Applications panel, in the navigation pane, click .

#### **Applications.**

### Procedure

1. In the Applications panel, click the application that you want to restore. The Detail view appears at the bottom of the screen.

 **Note** The Detail view appears only if you click an application. Selecting the check box before the name of the application will not open the Detail view.

2. In the Detail view, select the preferred restore point, and then click  **Restore**. The Application Restore dialog box opens.
3. From the Project drop-down menu, select the project that contains the instance to which you want to restore application data. By default, the original project of the instance on which the application is running is selected.
4. From the Zone drop-down menu, select the zone that contains the instance to which you want to restore application data. By default, the original zone of the instance on which the application is running is selected.
5. From the Instance drop-down menu, select the instance to which you want to restore application data.
6. Select the **Databases** check box if you want to restore the whole application or, from the list of databases that are available for the restore, select the ones that you want to restore.
7. Specify a point in time to which you want to restore application data. The databases will be restored to the state they were in at the specified time.
8. Enable the **Clear logs** switch if you want to initialize the log area. This option is by default disabled if you are restoring application data to the same instance and enabled if you are restoring application data to a different instance.
9. Click **Restore**.

# Protecting Google Kubernetes Engine applications

Protecting Google Kubernetes Engine (GKE) application data consist of the following tasks:

Task	Instructions
Preparing GKE applications for data protection, which includes applying labels on resource objects, discovering applications, and configuring backup options.	<a href="#">“Preparing for Google Kubernetes Engine application protection” below</a>
Backing up GKE applications.	<a href="#">“Backing up Google Kubernetes Engine applications” on page 77</a>
Restoring GKE application data.	<a href="#">“Restoring Google Kubernetes Engine applications” on page 78</a>

## Preparing for Google Kubernetes Engine application protection

Before you start protecting your Google Kubernetes Engine (GKE) applications, you must prepare your environment for application data protection.

### Prerequisite

The HYCU Managed Service Account (HMSA) must have the Compute Admin, Service Account User, Storage Admin, and Kubernetes Engine Admin roles granted on the projects with the Kubernetes clusters on which the GKE applications that you plan to protect are deployed.

For instructions on how to grant permissions to service accounts, see Google Cloud documentation.

### Limitations

- Protecting applications running on GKE clusters that were created by using the Autopilot mode of operation is not supported.



- R-Cloud does not support protecting applications that are configured in a subnet where Google Private Access is enabled and that are at the same time running on one of the following clusters:
  - A public GKE cluster without an internal IP address.
  - A private GKE cluster with the selected Access control plane using its external IP address option without an internal IP address.
- *For applications using volumes:* Only GCE persistent disk volumes and CSI volumes are supported.

Preparing your environment for GKE application data protection includes the following tasks:

Task	Instructions
1. <i>Mandatory.</i> Make sure appropriate labels are applied on all resource objects.	<a href="#">“Applying labels on resource objects” below</a>
2. <i>Mandatory.</i> Make sure your GKE applications are discovered in R-Cloud.	<a href="#">“Discovering applications” on the next page</a>
3. <i>Optional.</i> Configure GKE application backup options.	<a href="#">“Configuring GKE application backup options” on the next page</a>

## Applying labels on resource objects


To ensure that your GKE applications are successfully discovered and protected, appropriate metadata labels must be applied on the following:

- *Resource objects:* Make sure the following is defined:
  - `app.kubernetes.io/name: <AppName>` label in the `.yaml` file of the resource object
 

**Note** Specifying this label is recommended by R-Cloud. However, you can also use `app: <AppName>`.
  - Namespace in the metadata of the resource object
- *Persistent volume objects:* By applying labels, you ensure that persistent volumes can be discovered and linked to Google Compute Engine disks, which is required for zone/region identification:

### Example

```
topology.kubernetes.io/zone=us-east-1c
topology.kubernetes.io/zone=us-east-1c__us-east-1b (for
replicated disks)
topology.kubernetes.io/region=us-east-1
```


 **Note** For persistent volumes that use a Container Storage Interface (CSI) provider, the zone/region is specified in the volume handle (for example, volumeHandle: projects/<ProjectID>/zones/<Zone>/disks/<DiskName>).

The following deprecated Kubernetes labels are also supported:

```
failure-domain.beta.kubernetes.io/region=<RegionName>
failure-domain.beta.kubernetes.io/zone=<ZoneName>
```

For details on labels, see [Kubernetes documentation](#).

## Discovering applications

After you enable the HMSA, the process of application discovery starts automatically. When the application discovery task completes, the discovered applications are listed in the Applications panel. An automatic application synchronization task is performed every 15 minutes. You can update the application list manually at any time by navigating to the Applications panel and clicking  **Refresh**.

### Consideration

Before a GKE application can be discovered, the Kubernetes cluster on which it is deployed must be discovered by R-Cloud. This is an automated task that is performed every 15 minutes.

## Configuring GKE application backup options

You can adjust GKE application protection to the needs of your data protection environment by configuring application backup options.

### Backup options

Backup option	Description
Pre/post scripts	Enables you to specify the pre-snapshot and post-snapshot scripts to perform necessary actions before and/or after the snapshot of an application is created.

Backup option	Description
Temporary instance configuration	Enables you to specify the region, the zone, and the subnet where you want R-Cloud to create a temporary instance during the backup. By default, the temporary instance is created in the project of the GKE cluster on which the application is running.

### Prerequisites

- *Only if you plan to use pre-snapshot and post-snapshot scripts.*
  - The script must be located in a bucket to which the HMSA has access.
  - The `#!/usr/bin/env python3` header must be specified in the script.
  - The following line of code must be present in the script:


```
os.environ['GOOGLE_APPLICATION_CREDENTIALS'] =
'/tmp/hycu/serviceAccount.json'
```

- *Only if you plan to configure backup options for multiple applications.* All applications must have the same values set for each option that you plan to configure.

### Limitations

- You cannot specify a different subnet for the temporary instance if you are protecting applications running on a private GKE cluster with the disabled Access control plane using its external IP address option.
- *Only if you plan to use pre-snapshot and post-snapshot scripts.*
  - Only Python scripts are supported.
  - For making API calls, you can use only the following Python libraries:
    - `googleapiclient` for Google Cloud API calls
    - `kubernetes` for Kubernetes API calls


#### Accessing the Applications panel

To access the Applications panel, in the navigation pane, click .

#### **Applications.**

### Procedure

1. In the Applications panel, select the applications for which you want to configure backup options.

2. Click  **Configuration**. The Application Configuration dialog box opens.
3. Depending on what you want to do, provide the required information:
  - *Only if specifying the pre-snapshot and post-snapshot scripts.* On the Pre/post scripts tab, specify the scripts to perform necessary actions before and/or after the snapshot of the application is created:

- In the Pre-snapshot script field, enter the path to the script that R-Cloud will run just before it creates the snapshot of the application.
- In the Post-snapshot script field, enter the path to the script that R-Cloud will run immediately after it creates the snapshot of the application.

**ⓘ Important** When entering the path to the script, make sure to enter it correctly, including lowercase and uppercase letters, as the path is case sensitive. You must specify the path in the following format:

```
gs://bucket-name/script.py parameter1 parameter2 ...
```

**Example** The following is an example of the first lines of a pre-snapshot script:

```
#!/usr/bin/env python3
import os
import kubernetes

os.environ['GOOGLE_APPLICATION_CREDENTIALS'] =
'/tmp/hycu/serviceAccount.json'
```

- *Only if specifying the temporary instance location and subnet.* On the Temporary instance configuration tab, provide the following information:
    - a. From the Region drop-down menu, select the preferred region.
    - b. From the Zone drop-down menu, select the preferred zone.
    - c. From the Subnet drop-down menu, select the preferred subnet. By default, the temporary instance is created in the default subnet of the preferred region and zone.
4. Click **Save**.

## Backing up Google Kubernetes Engine applications

With R-Cloud, you can back up your GKE application data securely and efficiently.

### Prerequisite

*Only if you plan to back up applications running on clusters that use Shared VPC networks.* Your user account or the HYCU Managed Service Account (HMSA) must be granted the following permissions in the Shared VPC host project: `compute.firewalls.list`, `compute.networks.list`, `compute.networks.get`, `compute.subnetworks.list`, `compute.subnetworks.use`, and `compute.subnetworks.get`.


### Limitation

GKE application data can be stored only to Google Cloud targets.

### Consideration

*Only if the target defined in the policy that you plan to assign to GKE applications has a service account other than the HMSA specified.* R-Cloud uses this service account to access the GKE applications that you are backing up.


#### Accessing the Applications panel


To access the Applications panel, in the navigation pane, click 

**Applications.**

### Procedure

1. In the Applications panel, select the applications that you want to back up.

 **Tip** To narrow down the list of displayed applications, you can use the filtering options as described in [“Filtering and sorting data” on page 181](#).

2. Click  **Set Policy**. The Set Policy dialog box opens.
3. From the list of available policies, select the preferred policy.
4. Click **Assign** to assign the policy to the selected applications.

After you assign a policy to an application, a backup task starts immediately. Subsequent backups are scheduled according to the values defined in the policy.

If required, you can also perform a manual backup of any application at any time. For details, see [“Performing manual backups” on page 187](#).

## Restoring Google Kubernetes Engine applications

R-Cloud enables you to restore a whole application or only individual application items to a specific point in time.

### Prerequisites

*Only if you plan to specify post-restore scripts.*

- The script must be located in a bucket to which the HMSA has access.
- The `#!/usr/bin/env python3` header must be specified in the script.
- The following line of code must be present in the script:

```
os.environ['GOOGLE_APPLICATION_CREDENTIALS'] =
'/tmp/hycu/serviceAccount.json'
```

### Limitations


- Using the Restore Storage option is supported only for applications using persistent volumes.
- If your data is stored as a snapshot, you cannot restore it to a different source.
- *Only if you plan to specify post-restore scripts.*
  - Only Python scripts are supported.
  - For making API calls, you can use only the following Python libraries:
    - `googleapiclient` for Google Cloud API calls
    - `kubernetes` for Kubernetes API calls

Depending on how you want to restore data, do one of the following:

I want to...	Restore option	Instructions
Restore application storage together with all resource objects that are associated with the application to the original or a different location.	Restore Whole Application	<a href="#">“Restoring a whole application” on the next page</a>
Restore application storage to the	Restore Storage	<a href="#">“Restoring</a>

I want to...	Restore option	Instructions
original or a different location.		storage” on page 80
Restore specific resource objects to the original or a different location.	Restore Resource Objects	“Restoring resource objects” on page 82

### Accessing the Applications panel

To access the Applications panel, in the navigation pane, click 


#### Applications.


## Restoring a whole application

You can restore a whole application to its original or a different location by restoring application storage together with all resource objects that are associated with the application.

### Procedure

1. In the Applications panel, click the application that you want to restore. The Detail view appears at the bottom of the screen.

 **Note** The Detail view appears only if you click an application. Selecting the check box before the name of the application will not open the Detail view.

2. In the Detail view, select the preferred restore point, and then click  **Restore**. The Application Restore dialog box opens.
3. Select **Restore Whole Application**, and then click **Next**.
4. From the Restore from drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among which you can select:
  - **Automatic**: This option ensures the fastest and most cost-effective restore.
  - **Backup (Snapshot)**
  - **Backup (Target)**
  - **Copy**
  - **Archive—(daily, weekly, monthly, yearly)**
5. From the Target cluster drop-down menu, select the cluster to which you want to restore the application. You can select only among the clusters that

are in the same region as the application. By default, the original cluster of the application is selected.

6. From the Target namespace drop-down menu, select the namespace to which you want to restore the application. The original namespace of the application is preselected.
7. Use the **Keep original configuration** switch if you want to keep the existing resource object configuration. If you disable the switch, the resource object configuration will be overwritten (including persistent volumes).
8. *Optional.* In the Post-restore script field, enter the path to the script that R-Cloud should run after the restore.

**ⓘ Important** When entering the path to the script, make sure to enter it correctly, including lowercase and uppercase letters, as the path is case sensitive. You must specify the path in the following format:  
 gs://<PathtoBucket>/script.py parameter1 parameter2 ...

**Example** The following is an example of the first lines of a post-restore script:

```
#!/usr/bin/env python3
import os

os.environ['GOOGLE_APPLICATION_CREDENTIALS'] =
'/tmp/hycu/serviceAccount.json'
```

9. Click **Restore**.

## Restoring storage

You can restore data that was stored on one or more disks at backup time to the same or a different location by restoring one or more persistent volume claims.


**ⓘ Important** You cannot restore an application by restoring its storage. For instructions on how to restore a whole application, see [“Restoring a whole application” on the previous page](#).


### Procedure

1. In the Applications panel, click the application whose storage you want to restore. The Detail view appears at the bottom of the screen.

**📄 Note** The Detail view appears only if you click an application. Selecting the check box before the name of the application will not open



- the Detail view.
2. In the Detail view, select the preferred restore point, and then click  **Restore**. The Application Restore dialog box opens.
  3. Select **Restore Storage**, and then click **Next**.
  4. From the Restore from drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among which you can select:
    - **Automatic**: This option ensures the fastest and most cost-effective restore.
    - **Backup (Snapshot)**
    - **Backup (Target)**
    - **Copy**
    - **Archive—(daily, weekly, monthly, yearly)**
  5. From the Target cluster drop-down menu, select the cluster to which you want to restore storage. You can select only among the clusters that are in the same region as the application. By default, the original cluster of the application is selected.
  6. From the Target namespace drop-down menu, select the namespace to which you want to restore storage. The original namespace of the application is preselected.
  7. *Optional*. In the Post-restore script field, enter the path to the script that R-Cloud should run after the restore.

 **Important** When entering the path to the script, make sure to enter it correctly, including lowercase and uppercase letters, as the path is case sensitive. You must specify the path in the following format:


`gs://<PathtoBucket>/script.py parameter1 parameter2 ...`

**Example** The following is an example of the first lines of a post-restore script:

```
#!/usr/bin/env python3
import os

os.environ['GOOGLE_APPLICATION_CREDENTIALS'] =
'/tmp/hycu/serviceAccount.json'
```


8. Select the persistent volume claims that you want to restore.

 **Tip** Select the **Disks** check box to restore all persistent volume claims.

9. Use the **Keep original volumes** switch if you want to keep the original persistent volumes. If you disable the switch, the original volumes will be overwritten by the restored ones.
10. Click **Restore**.



## Restoring resource objects

You can restore specific resource objects to their original or a different location.

 **Caution** Restoring resource objects must be performed in the correct order, taking into account the dependencies among different resource objects.

### Procedure

1. In the Applications panel, click the application whose resource objects you want to restore. The Detail view appears at the bottom of the screen.
 

 **Note** The Detail view appears only if you click an application. Selecting the check box before the name of the application will not open the Detail view.
2. In the Detail view, select the preferred restore point, and then click  **Restore**. The Application Restore dialog box opens.
3. Select **Restore Resource Objects**, and then click **Next**.
4. From the Restore from drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among which you can select:
  - **Automatic**: This option ensures the fastest and most cost-effective restore.
  - **Backup (Snapshot)**
  - **Backup (Target)**
  - **Copy**
  - **Archive—(daily, weekly, monthly, yearly)**
5. From the Target cluster drop-down menu, select the cluster to which you want to restore resource objects. The original cluster of the application is preselected.

6. *Optional.* In the Post-restore script field, enter the path to the script that R-Cloud should run after the restore.

ⓘ **Important** When entering the path to the script, make sure to enter it correctly, including lowercase and uppercase letters, as the path is case sensitive. You must specify the path in the following format:  
`gs://<PathtoBucket>/script.py parameter1 parameter2 ...`

**Example** The following is an example of the first lines of a post-restore script:

```
#!/usr/bin/env python3
import os

os.environ['GOOGLE_APPLICATION_CREDENTIALS'] =
'/tmp/hycu/serviceAccount.json'
```

7. Click **Next**.
8. From the list of available resource objects, select the ones that you want to restore.
9. Click **Restore**.

# Chapter 6

## Protecting instances

R-Cloud enables you to protect your instance data with fast and reliable backup and restore operations. After you back up an instance, you can choose to restore the entire instance, disks, or individual files.

For details on how to protect instance data efficiently, see the following sections:

- [“Planning instance protection” below](#)
- [“Backing up instances” on page 92](#)
- [“Restoring instances” on page 94](#)
- [“Restoring individual files or folders” on page 132](#)

## Planning instance protection

Before performing a backup, get familiar with the prerequisites, limitations, considerations, and recommendations that are general for all data protection environments and those that are specific for your data protection environment needs.

- [“Preparing your data protection environment” below](#)
- [“Configuring instance backup options” on page 86](#)

## Preparing your data protection environment

### Prerequisites

- *For AWS:* Before backing up an instance, make sure the following prerequisites are met in the location of the instance:
  - In a VPC without public IPs or in subnets without public IPs, you must create the following VPC endpoints:

- Interface endpoints: Amazon EC2 (ec2), AWS Security Token Service (sts), Amazon SQS (sqs), and Amazon SNS (sns)
- Gateway endpoint for Amazon S3

For details on how to enable Amazon VPC endpoints, see AWS documentation.

- The security group that the instance belongs to must have an inbound firewall rule for port 443 (HTTPS), source IP 0.0.0.0/0 and an outbound firewall rule for port 443 (HTTPS), destination IP 0.0.0.0/0.

For instructions on how to configure and apply the network firewall rule, see AWS documentation.

- *For Google Cloud:*
  - The HYCU Managed Service Account (HMSA) must have the Compute Admin, Service Account User, and Storage Admin roles granted on the projects with the instances that you plan to protect. For instructions on how to grant permissions to service accounts, see Google Cloud documentation.
  - Cloud Resource Manager API, Compute Engine API, Cloud Identity and Access Management API, Cloud Billing API, and Cloud Storage API must be enabled on the Google Cloud projects that contain the instances that you want to protect. For instructions on how to enable APIs, see Google Cloud documentation.
  - *Only if you plan to back up and restore instances that use Shared VPC networks.* Your user account or the HYCU Managed Service Account (HMSA) must be granted the following permissions in the Shared VPC host project: `compute.firewalls.list`, `compute.networks.list`, `compute.networks.get`, `compute.subnetworks.list`, `compute.subnetworks.use`, and `compute.subnetworks.get`.

## Limitations

- Instance memory is not protected.
- Crash consistency of backup data is guaranteed only for each disk individually.
- *For Google Cloud:* Local SSDs are not protected.

## Considerations

- Keep in mind that the role you have assigned determines what kind of actions you can perform. For details on roles, see [“Managing roles”](#) on

[page 225](#).

- Backup data (including copies of backup data and data archives) that R-Cloud creates is crash-consistent, but it may not always be application-consistent. If pre-snapshot scripts are not provided, the application consistency of backup data is limited to the following:
  - Applications that store their data on a single disk.
  - AWS instances and applications that comply with the prerequisites for creating a Windows Volume Shadow Copy Service (VSS) snapshot. For more information about Windows VSS snapshot prerequisites, see [“Backing up instances” on page 92](#).
  - Google Cloud instances and applications that comply with the restrictions for creating a Windows Volume Shadow Copy Service (VSS) snapshot. For details, see Google Cloud documentation.
- *For Google Cloud:* R-Cloud uses an external IP address to access Google Cloud APIs if Private Google Access is disabled on subnets. If your data protection environment requires the use of an internal IP address, make sure Private Google Access is enabled on subnets. For details, see Google Cloud documentation.

### Recommendation

*Only if you delete an instance from Google Cloud.* If an instance that you delete from Google Cloud still has at least one valid restore point available in R-Cloud, it is considered protected and its status is Protected deleted. If you create a new instance with the same name, project, and zone in Google Cloud, R-Cloud will recognize this instance as the old one during instance synchronization and change its status from Protected deleted to Protected. Using the restore points of such an instance for a restore could result in data corruption. Therefore, it is recommended that you create the new instance with a different name, project, or zone, or that you mark the restore points of the old instance as expired before performing a restore. For details on marking restore points as expired, see [“Expiring backups manually” on page 188](#).


## Configuring instance backup options

Before you start protecting instances, you can adjust instance protection to the needs of your data protection environment by configuring backup options. You can configure backup options for a single instance or for multiple instances at the same time.

## Backup options

Backup option	Description
Pre/post scripts	Enables you to specify the pre-snapshot and post-snapshot scripts to perform necessary actions before and/or after the snapshot of an instance is created. For example, if the instance hosts a database management system, you may want to put the database offline before the snapshot is created to ensure an application-consistent backup and bring the database back online when the snapshot creation completes. For details, see <a href="#">“Setting pre/post scripts”</a> below.
Exclude from backup	Enables you to specify any disk to be excluded from the instance backup. For details, see <a href="#">“Excluding disks from the backup”</a> on page 90.  <div style="border-left: 2px solid purple; padding-left: 10px; margin-left: 20px;"> <p><b>ⓘ Important</b> This backup option can be configured only for a single instance.</p> </div>
Temporary instance configuration	Enables you to specify the region, the zone, and the subnet where you want R-Cloud to create a temporary instance during the backup. For AWS instances, you can also specify a security group. By default, the temporary instance is created in the source of the original instance. For details, see <a href="#">“Configuring a temporary instance”</a> on page 91.

## Accessing the Instances panel

To access the Instances panel, in the navigation pane, click  **Instances**.

## Setting pre/post scripts

### Prerequisites

- Access to the instance file system must be enabled. For instructions, see [“Enabling access to data”](#) on page 51.
- A script must be available in an accessible folder.
- The user account must have permissions to run a script on the instance with the assigned credentials.

## Consideration


The scripts are run from the home directory of the user account that R-Cloud uses for running the scripts.


Depending on the operating system on the instance, the following user accounts are used:

- *For AWS instances:* The user account that is assigned to the instance in R-Cloud through a credential group.
- *For Google Cloud instances running Linux:*
  - *The instance has no credential group assigned in R-Cloud:* The HYCU Managed Service Account (HMSA).
  - *The instance has a credential group assigned:* The user account specified in the credential group.
- *For Google Cloud instances running Windows:* The user account that is assigned to the instance in R-Cloud by means of a credential group.

## Procedure

1. In the Instances panel, select the instances for which you want to set pre/post scripts.

 **Note** If you want to narrow down the list of displayed instances, use the filtering options as described in “[Filtering and sorting data](#)” on [page 181](#).

2. Click  **Configuration**.
3. Depending on whether you want to set a pre/post-snapshot script for a single instance or multiple instances, on the Pre/Post Scripts tab, do the following:



- For a single instance:

Field	Instructions
Pre-snapshot Script	<p>Enter the path to the script that R-Cloud should run before it creates a snapshot of the instance. For example:</p> <ul style="list-style-type: none"> <li>◦ Windows: %USERPROFILE%\quiesce_db.bat</li> <li>◦ Linux: bash /home/&lt;UserName&gt;/freeze_db.sh</li> </ul>
Post-snapshot Script	<p>Enter the path to the script that R-Cloud should run after it creates a snapshot of the instance. For example:</p> <ul style="list-style-type: none"> <li>◦ Windows: %USERPROFILE%\resume_db.bat</li> <li>◦ Linux: bash /home/&lt;UserName&gt;/thaw_db.sh</li> </ul>

- For multiple instances:

Field	Instructions
Pre-snapshot Script	<p>a. Specify the path to the script that R-Cloud should run before it creates the snapshots of the instances. To do so, choose one of the following:</p> <ul style="list-style-type: none"> <li>◦ If you want to use a new script, select <b>Add New</b>, enter the path to the script, and then click <b>Save</b>. The following are examples of the scripts: <ul style="list-style-type: none"> <li>◦ Windows: %USERPROFILE%\quiesce_db.bat</li> <li>◦ Linux: bash /home/&lt;UserName&gt;/freeze_db.sh</li> </ul> </li> <li>◦ If any of the selected instances already have a pre-snapshot script set and you want to use the same script for all other selected instances, select the preferred script.</li> </ul> <p>b. <i>Only if any of the selected instances already have a pre-snapshot script set.</i> Select the <b>Override these</b></p>

Field	Instructions
	<p><b>instances</b> check box if you want the specified script to be used for all the selected instances.</p>
Post-snapshot Script	<p>a. Specify the path to the script that R-Cloud should run after it creates the snapshots of the instances. To do so, choose one of the following:</p> <ul style="list-style-type: none"> <li>◦ If you want to use a new script, select <b>+</b> <b>Add New</b>, enter the path to the script, and then click <b>Save</b>.</li> </ul> <p>The following are examples of the scripts:</p> <ul style="list-style-type: none"> <li>◦ Windows:           <pre>%USERPROFILE%\resume_db.bat</pre> </li> <li>◦ Linux:           <pre>bash /home/&lt;UserName&gt;/thaw_db.sh</pre> </li> <li>◦ If any of the selected instances already have a post-snapshot script set and you want to use the same script for all other selected instances, select the preferred script.</li> </ul> <p>b. <i>Only if any of the selected instances already have a post-snapshot script set.</i> Select the <b>Override these instances</b> check box if you want the specified script to be used for all the selected instances.</p>

4. Click **Save**.

## Excluding disks from the backup


### Consideration

If you exclude the boot disk from the backup, you cannot restore the entire instance, but only its disks.

### Procedure

1. In the Instances panel, select the instance for which you want to exclude disks from the backup.

 **Note** If you want to narrow down the list of displayed instances, use the filtering options as described in “[Filtering and sorting data](#)” on [page 181](#).

2. Click  **Configuration**.
3. On the Exclude from Backup tab, select the disks that you want to exclude from the backup.
4. Click **Save**.


## Configuring a temporary instance


### Prerequisites

- *Only if you are protecting Google Cloud instances and you plan to specify a different subnet for the temporary instance.* If you plan to use pre-snapshot and post-snapshot scripts, or back up instances for which the restore of individual files is allowed, VPC Network Peering must be configured. For details on how to configure VPC Network Peering, see Google Cloud documentation.
- *Only if you plan to configure the temporary instance for multiple instances.* The selected instances must have the same values set for the location and the subnet. In the case of AWS instances, the same value must be set also for the security group.

### Procedure

1. In the Instances panel, select the instances for which you want to configure the temporary instance.

 **Note** If you want to narrow down the list of displayed instances, use the filtering options as described in [“Filtering and sorting data” on page 181](#).


2. Click  **Configuration**.
3. On the Temporary Instance Configuration tab, select the following:
  - a. From the Region drop-down menu, select the preferred region.
  - b. From the Zone drop-down menu, select the preferred zone.
  - c. From the Subnet drop-down menu, select the preferred subnet. By default, the temporary instance is created in the default subnet of the preferred region and zone.
  - d. *For AWS instances:* Optionally, from the Security Group drop-down menu, select the preferred security group. By default, the temporary instance is created in the default security group of the preferred subnet.
4. Click **Save**.

# Backing up instances

With R-Cloud, you can back up your instance data securely and efficiently.

## Prerequisites

- *For AWS instances whose data you plan to back up from one AWS account to a target in another AWS account:* The ID of the AWS account to which the instance belongs must be specified as the `Principal` element in the role trust policy of the IAM role within the AWS account that contains the target. For details on policies and permissions in IAM, see AWS documentation.
- *For AWS instances running Windows for which you want to ensure application consistency of backup data:*
  - You must create an IAM role for VSS-enabled snapshots and attach it to the instance. For details on how to create an IAM role for VSS-enabled snapshots, see AWS documentation.
  - All attached disks must be online.

 **Note** You can check if a VSS snapshot was successfully created for the instance in the backup task summary and report.

- *Only if you plan to use pre-snapshot or post-snapshot scripts.* Access to data must be enabled. For details, see [“Enabling access to data” on page 51](#).

## Limitations


- *For AWS instances:* If an instance has private access configured, you cannot store its backup data on a Google Cloud target. For more information about private access, see AWS documentation for VPC endpoints.
- *For Google Cloud instances:* If an instance has private access configured, you cannot store its backup data on an Amazon S3 target. For more information about private access, see Google Cloud documentation for Private Google Access.

## Considerations


- When backing up an instance with multiple disks, R-Cloud performs a parallel backup. In the Tasks panel, you can view details on the backup progress, including the progress of backing up each individual disk.
- When backing up an instance, R-Cloud can reduce the size of the backup data that will be transferred to the target by using data compression:

- If backing up the instance to the selected target for the first time, the data is compressed by default.
- If backing up the instance to a target that stores the previous backup data of the instance, the backup data is not compressed. This is done to avoid performing a full backup due to the changed properties of the compressed data.
- *Only if you plan to restore individual files or folders.* To use the default view when restoring individual files or folders, access to data must be enabled and instance discovery must complete successfully before the backup is made. If the default view is not available, you can restore the files or folders by using the filesystem view. For instructions on how to enable access to data, see [“Enabling access to data” on page 51](#).


### Accessing the Instances panel


To access the Instances panel, in the navigation pane, click  **Instances**.

### Procedure


1. Select the instances that you want to back up. You can update the instance list by clicking  **Refresh**. In a protection set with a large number of sources, the update may take a while.

To narrow down the list of displayed instances, use the filtering options as described in [“Filtering and sorting data” on page 181](#).

2. Click  **Set Policy**. The Set Policy dialog box opens.
3. From the list of available policies, select the preferred policy.
4. Click **Assign** to assign the policy to the selected instances.

 **Important** *Only if one or more instances that you want to back up reside on a different cloud platform than the target defined in the policy. Click **Assign Anyway** if you want to assign the policy to the instance. This may result in data transfer fees. To cancel assigning the policy, click **Cancel**.*

When you assign a policy to an instance, a backup task starts immediately. Subsequent backups are scheduled according to the values defined in the policy.

 **Note** The first backup task may be delayed if a backup of the instance already exists.

You can also perform a manual backup of individual instances at any time. For details, see [“Performing manual backups” on page 187](#).

# Restoring instances

R-Cloud enables you to restore an entire instance or its individual disks to a specific point in time, or multiple instances or disks belonging to multiple instances in a single session. In the event of a disaster in your environment, you can also restore instance data to a different cloud platform by using the Move Instance restore option.

## Prerequisites

- *Only if you plan to specify post-restore scripts.*
  - Access to the instance file system must be enabled. For instructions, see [“Enabling access to data” on page 51](#).
  - A script must be available in an accessible folder.
  - The user account must have permissions to run a script on the instance.
- *For AWS:* Before restoring an instance to the same or a different VPC, subnet, or region, make sure the following prerequisites are met in the location of the instance:
  - In a VPC without public IPs or in subnets without public IPs, you must create the following VPC endpoints:
    - Interface endpoints: Amazon EC2 (ec2), AWS Security Token Service (sts), Amazon SQS (sqs), and Amazon SNS (sns)
    - Gateway endpoint for Amazon S3

For details on how to enable Amazon VPC endpoints, see AWS documentation.

- The security group that the instance belongs to must have an inbound firewall rule for port 443 (HTTPS), source IP 0.0.0.0/0 and an outbound firewall rule for port 443 (HTTPS), destination IP 0.0.0.0/0.

For instructions on how to configure and apply the network firewall rule, see AWS documentation.

## Limitation

If your data is stored as a snapshot, you cannot restore it to a different source.


## Considerations

- Only one restore task can run at the same time for the instance.
- *Only if you plan to specify post-restore scripts.* The scripts are run from the home directory of the user account that R-Cloud uses for running the scripts.
- Depending on the operating system on the instance, the following user accounts are used:
  - *For AWS instances:* The user account that is assigned to the instance in R-Cloud through a credential group.
  - *For Google Cloud instances running Linux:*
    - *The instance has no credential group assigned in R-Cloud:* The HYCU Managed Service Account (HMSA).
    - *The instance has a credential group assigned:* The user account specified in the credential group.
  - *For Google Cloud instances running Windows:* The user account that is assigned to the instance in R-Cloud by means of a credential group.

Depending on what you plan to restore, see one of the following sections:

I plan to restore...	Instructions
A single instance or its disks.	<a href="#">“Restoring a single instance or its disks” below</a>
Multiple instances or disks belonging to multiple instances in a single session.	<a href="#">“Restoring multiple instances or disks belonging to multiple instances in a single session” on page 121</a>

### Accessing the Instances panel

To access the Instances panel, in the navigation pane, click  **Instances**.

## Restoring a single instance or its disks

When you restore a single instance or its disks, you can select among the following options:

Option	Description	Instructions
Restore Instance	Enables you to restore an instance and its disks to the original location	<a href="#">“Restoring an instance” on the</a>

Option	Description	Instructions
	with the same settings.	<a href="#">next page</a>
Restore Disks	Enables you to restore disks and attach them to the same instance.	<a href="#">“Restoring disks” on the next page</a>
Clone Instance	Enables you to restore an instance and its disks by creating a clone of the instance.	<a href="#">“Cloning an instance” on page 98</a>
Clone Disks	Enables you to restore disks by creating their clones and attaching them to the same or a different instance, or by creating their clones in the same or a different source and zone and leaving them unattached.	<a href="#">“Cloning disks” on page 108</a>
Move Instance	Enables you to move an instance by restoring it to a different cloud platform.	<a href="#">“Moving an instance” on page 110</a>
Move Disks	Enables you to move disks by restoring them and attaching them to an instance on a different cloud platform, or by restoring them to a different cloud platform and leaving them unattached.	<a href="#">“Moving disks” on page 119</a>

## Restoring an instance

You can restore an instance and its disks to the original location with the same settings. In this case, you replace the original instance with the restored one.


### Consideration


Any data changes after the last successful backup are not protected and therefore cannot be restored.


### Procedure

1. In the Instances panel, click the instance that you want to restore to open the Detail view.




 **Note** The Detail view appears only if you click an instance. Selecting the check box before the name of the instance does not open the Detail view.

2. In the Detail view that appears at the bottom of the screen, select the preferred restore point.
3. Click  **Restore Instance**. The Restore Options dialog box opens.
4. Select **Restore Options**, and then click **Next**.
5. Select **Restore Instance**, and then click **Next**.
6. From the Restore from drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among which you can select:
  - **Automatic**: This option ensures the fastest and most cost-effective restore.
  - **Backup (Snapshot)**
  - **Backup (Target)**
  - **Copy**
  - **Archive—(daily, weekly, monthly, yearly)**
7. From the Disks drop-down menu, select the instance disks that you want to restore.

 **Note** All disks of the instance are preselected for the restore. With such selection, the entire instance is restored. The boot disk is restored even if you do not select it.

8. *Optional*. In the Post-restore script field, enter the path to the script or a command that R-Cloud should run on the instance after the restore.

 **Note** You can enter any command that the command-line interface of your instance supports.


9. Click **Restore**.


## Restoring disks


You can restore disks and attach them to the same instance. In this case, you replace the original disks with the restored ones.

### Procedure


1. In the Instances panel, click the instance whose disks you want to restore to open the Detail view.

 **Note** The Detail view appears only if you click an instance. Selecting the check box before the name of the instance will not open the Detail view.

2. In the Detail view that appears at the bottom of the screen, select the preferred restore point.
3. Click  **Restore Instance**. The Restore Options dialog box opens.
4. Select **Restore Options**, and then click **Next**.
5. Select **Restore Disks**, and then click **Next**.
6. From the list of disks that are available for the restore, select the ones that you want to restore, and then click **Next**.

 **Note** If you select the boot disk, the instance will be shut down and restarted when the disks are restored.

7. From the Restore from drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among which you can select:
  - **Automatic**: This option ensures the fastest and most cost-effective restore.
  - **Backup (Snapshot)**
  - **Backup (Target)**
  - **Copy**
  - **Archive—(daily, weekly, monthly, yearly)**
8. *Optional.* In the Post-restore script field, enter the path to the script or a command that R-Cloud should run after the restore on the instance to which the restored disks are attached.

 **Note** You can enter any command that the command-line interface of your instance supports.

9. Click **Restore**.

## Cloning an instance

You can clone an instance by restoring it to its original or a new location with custom settings. In this case, you create a new instance containing the restored data alongside the original instance. When cloning an instance, you can change the following properties: the selection of the backed-up disks, the destination source, region, and zone, and the instance network configuration.

For details on how to clone AWS and Google Cloud instances, see the following sections:

- “Cloning an AWS instance” below
- “Cloning a Google Cloud instance” on page 103


## Cloning an AWS instance


### Limitations

- You cannot restore instances that belong to a deleted AWS account. Such instances are not listed in the Instances panel of the R-Cloud web user interface.
- You cannot restore an instance to a different source or AWS region from a snapshot.
- *For instances running Windows:* Using post-restore scripts is not supported.


### Procedure

1. In the Instances panel, click the instance that you want to restore to open the Detail view.


 **Note** The Detail view appears only if you click an instance. Selecting the check box before the name of the instance will not open the Detail view.


2. In the Detail view that appears at the bottom of the screen, select the preferred restore point.
3. Click  **Restore Instance**. The Restore Options dialog box opens.
4. Select **Clone Options**, and then click **Next**.
5. Select **Clone Instance**, and then click **Next**.
6. In the New instance name field, specify a new name for the instance.
7. From the Restore from drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among which you can select:
  - **Automatic:** This option ensures the fastest and most cost-effective restore.
  - **Backup (Snapshot)**
  - **Backup (Target)**
  - **Copy**
  - **Archive—(daily, weekly, monthly, yearly)**
8. *Only if the original operating system image was not found.* From the Image drop-down menu, select the operating system image you want to use. To use a custom image, select **Use custom image** and enter the image AMI ID.

9. *Optional.* In the Post-restore script field, enter the path to the script or a command that R-Cloud should run on the restored instance after the restore.


 **Note** You can enter any command that the command-line interface of your instance supports.

10. From the Destination source drop-down menu, select the source to which you want to restore the instance. The original source of the instance is preselected. You can choose from sources that belong to the currently selected protection set and that your user account can access.
11. From the Destination region and Destination zone drop-down menus, select the AWS region and zone to which you want to restore the instance. The original region and zone of the instance are preselected.
12. Under Disk name, do the following:
- a. Select the instance disks that you want to restore.


 **Note** All disks of the instance are preselected for the restore. With such selection, the entire instance is restored. The boot disk is restored even if you do not select it.

- b. Edit the disks as required. For each selected disk, do the following:
  - i. Click  **Edit Disk**.
  - ii. *Only if you do not want R-Cloud to automatically generate a name for the restored disk device or disk.* Do the following:
    - I. In the New device name field, enter a name for the restored disk device.
    - II. In the New disk name field, enter a name for the restored disk.
  - iii. If you want to change the disk type, from the Disk type drop-down menu, select one of the available disk types for the restored disk. By default, the original disk type is selected.
 

The list shows only the disk types that match the required disk size and can include the following disk types: General Purpose SSD, Previous Generation Volume, and Provisioned IOPS SSD.

If you selected Provisioned IOPS SSD or General Purpose SSD, enter the IOPS number.
  - iv. *Only if you want to add labels to the restored disk.*
    - I. Click **Advanced**.
    - II. Click  **Manage**. The Custom Metadata dialog box opens.

III. Enter a key and a value, and then click **Add** for each label that you want to add.

 **Note** If the selected disk already has one or more labels added, they are listed under Labels. If you want to delete any of the added labels, click **X Delete** next to it.

v. Click **Save**.


13. Under Network interfaces, review the list of networks that the original instance was configured in at the time of backup. The list shows the following for each such network:

- VPC ID
- Subnet ID



For each configured network interface, you can separately adjust its public and private IP address types. By default, the public IP address configuration of the original instance is kept.

#### Modifying network settings


If you want to modify network settings, you can add an additional network interface, edit an existing network interface, or delete a network interface:


- Click **Add network interface** to add a network interface or click  **Edit** next to the network interface that you want to edit, and then follow these steps:
  - a. From the Subnet drop-down menu, select the subnet.
  - b. From the Security group drop-down menu, select the security group.
  - c. In the Public address type field, select the public IP address for the network interface. You can select among the following options:


Option	Description
None	The network interface does not use a public IP address.  This option is preselected if the network interface of the original instance did not use a public IP address.
Auto-assign	The network interface uses an automatically allocated public IP address.



Option	Description
	<p>This option is preselected if the network interface of the original instance used a public IP address.</p> <p> <b>Note</b> Auto-assign will not work if the Auto-assign public IPv4 address on a subnet option is set to No or if more than one network interface is specified.</p>
Elastic IP (Reserved)	The network interface uses an elastic public IP address that was reserved in Amazon EC2 in advance.
Elastic IP (New)	<p>The network interface uses a new elastic public IP address.</p> <p> <b>Note</b> Allocation of the IP address in Amazon EC2 is performed at the very beginning of the restore. If the allocation fails, the restore task is terminated without being logged.</p>

- d. In the Private address type field, select the private IP address for the network interface. You can select between the following options:

Option	Description
Auto-assign	<p>The network interface uses an automatically allocated private IP address.</p> <p>This option is selected by default.</p>
Custom	<p>The network interface uses a private IP address that is defined by you.</p> <p> <b>Important</b> Use of this option might result in IP address conflicts.</p>

- e. Click **Add** or **Save**.
- Click  **Delete** next to the network interface that you want to delete. Keep in mind that you cannot restore the instance without a network interface.

14. *Only if you want to add tags to the restored instance.*
  - a. Click **Advanced**.
  - b. For each custom metadata tag that you want to add, click  **Manage**.
  - c. Enter a key and a value, and then click **Add**.

 **Note** If the selected instance already has one or more custom metadata tags added, they are listed under Custom metadata. If you want to delete any of the added custom metadata tags, click  **Delete** next to it.

15. Click **Restore**.

## Cloning a Google Cloud instance

### Limitation

You cannot restore instances that belong to a deleted Google Cloud project. Such instances are not listed in the Instances panel of the R-Cloud web user interface.


### Considerations


*Only if you plan to replicate disks.*

- The boot disk cannot be replicated.
- Standard persistent disks smaller than 200 GiB cannot be replicated.
- Regional disks can be replicated only across two zones in the same region. One of these zones must be the same as the zone of the target instance.
- If the region or zone of the target instance changes, all regional disks are automatically converted to zonal disks. In this case, the procedure of replicating the disks must be performed again.


### Procedure

1. In the Instances panel, click the instance that you want to restore to open the Detail view.


 **Note** The Detail view appears only if you click an instance. Selecting the check box before the name of the instance will not open the Detail view.


2. In the Detail view that appears at the bottom of the screen, select the preferred restore point.
3. Click  **Restore Instance**. The Restore Options dialog box opens.

4. Select **Clone Options**, and then click **Next**.
5. Select **Clone Instance**, and then click **Next**.
6. In the New instance name field, specify a new name for the instance.
7. From the Restore from drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among which you can select:
  - **Automatic**: This option ensures the fastest and most cost-effective restore.
  - **Backup (Snapshot)**
  - **Backup (Target)**
  - **Copy**
  - **Archive—(daily, weekly, monthly, yearly)**
8. *Optional*. In the Post-restore script field, enter the path to the script or a command that R-Cloud should run on the restored instance after the restore.



 **Note** You can enter any command that the command-line interface of your instance supports.

9. From the Destination source drop-down menu, select the source to which you want to restore the instance. The original source of the instance is preselected. You can choose from sources that belong to the currently selected protection set and that your user account can access.
10. From the Destination region and Destination zone drop-down menus, select the Google Cloud region and zone to which you want to restore the instance. The original region and zone of the instance are preselected.
11. Under Disk name, do the following:
  - a. Select the instance disks that you want to restore.

 **Note** All disks of the instance are preselected for the restore. With such selection, the entire instance is restored. The boot disk is restored even if you do not select it.

- b. Edit the disks as required. For each selected disk, do the following:
  - i. Click  **Edit Disk**.
  - ii. *Only if you do not want R-Cloud to automatically generate a name for the restored disk device or disk.* Do the following:
    - I. In the New device name field, enter a name for the restored disk device.
    - II. In the New disk name field, enter a name for the restored disk.




- iii. If you want to change the disk type, from the Disk type drop-down menu, select one of the available disk types for the restored disk (Standard persistent disk, Balanced persistent disks, or SSD persistent disk). By default, the original disk type is selected.
  - iv. If you want to replicate data between two zones in the region of the instance, make sure the **Replicate this disk within region** check box is selected, and then, from the Target zone drop-down menu, select to which zone you want to replicate data. If the selected disk was regional at backup time, the two zones across which the disk is replicated are shown, otherwise, a list of all zones in the region of the instance is shown.
  - v. *Only if you want to add labels to the restored disk.*
    - I. Click **Advanced**.
    - II. Click  **Manage**. The Custom Metadata dialog box opens.
    - III. Enter a key and a value, and then click **Add** for each label that you want to add.
      -  **Note** If the selected disk already has one or more labels added, they are listed under Labels. If you want to delete any of the added labels, click **X Delete** next to it.
  - vi. Click **Save**.
12. Under Network Interfaces, review the list of networks that the original instance was configured in at the time of backup. The list shows the following for each such network:
- Network type: Subnetwork for VPC networks and shared VPC networks, Legacy for legacy networks
  - Subnetwork name (for VPC networks and shared VPC networks) or network name (for legacy networks)
  - *Only in case of a shared VPC network.* Name of the host project of the network

For each configured network interface, you can separately adjust its public and private IP address types. By default, the public IP address configuration of the original instance is kept.

#### Modifying network settings



If you want to modify network settings, you can add an additional network interface, edit an existing network interface, or delete a network interface:

- Click **Add network interface** to add a network interface or click  **Edit** next to the network interface that you want to edit, and then follow these steps:
  - a. From the Destination network drop-down menu, select the destination network.
  - b. In the Public address type field, select the public IP address for the network interface. You can select among the following options:


Option	Description
None	The network interface does not use a public IP address.
	This option is preselected if the network interface of the original instance did not use a public IP address.
Ephemeral	The network interface uses an automatically allocated public IP address.
	This option is preselected if the network interface of the original instance used a public IP address.
Static (Reserved)	The network interface uses a static public IP address that was reserved in Google Compute Engine in advance.
Static (New)	The network interface uses a static public IP address that is allocated at the time of the restore. If the allocation fails, the instance is assigned a temporary public IP address. Such fallback also sets the restore task status to Done with errors.

- c. In the Private address type field, select the private IP address for the network interface. You can select between the following options:

Option	Description
Ephemeral (Automatic)	The network interface uses an automatically allocated private IP address.  This option is selected by default for the


Option	Description
	preselected network interfaces.
Ephemeral (Custom)	The network interface uses a private IP address that is defined by you.   <b>Important</b> Use of this option might result in IP address conflicts.
Static (Reserved)	<i>Not available for legacy networks.</i> The network interface uses a static private IP address that was reserved in Google Compute Engine or in advance.
Static (New)	<i>Not available for legacy networks.</i> The network interface uses a new static private IP address that is defined by you.   <b>Note</b> Allocation of the IP address in Google Compute Engine is performed at the very beginning of the restore. If the allocation fails, the restore task is terminated without being logged.

d. Click **Add** or **Save**.



- Click  **Delete** next to the network interface that you want to delete. Keep in mind that you cannot restore the instance without a network interface.

13. *Only if you want to add tags and/or labels to the restored instance.*

a. Click **Advanced**.

b. For each label, network tag, or custom metadata tag that you want to add, click  **Manage**.

c. Enter a key and a value, and then click **Add**.

 **Note** If the selected instance already has one or more labels, network tags, and/or custom metadata tags added, they are listed under Labels, Network tags, or Custom metadata. If you want to delete any of the added labels, network tags, and/or custom metadata tags, click  **Delete** next to it.

14. Click **Restore**.

## Cloning disks

You can create clones of disks by restoring them and attaching them to the same or a different instance, or by restoring them to the same or a different source, region, or zone and leaving them unattached. In this case, the original disks will not be overwritten.

### Limitations


- You can attach the restored disks only to an instance that is running the same operating system as the original instance and that belongs to the same protection set as the original instance.
- You cannot restore disks to a different source or region from a snapshot.
- *Only if you plan to restore disks to a different source and leave them unattached.* The default network must be set for the source to which you plan to restore the disks, or the source to which you plan to restore the disks must have the same network as the instance whose disks you plan to restore.


### Considerations



- For details on how the restored disks are named, see [“Resources created by R-Cloud” on page 242](#).
- *Only if you are restoring Google Cloud instance disks and you plan to replicate disks.*
  - The boot disk cannot be replicated.
  - Standard persistent disks smaller than 200 GiB cannot be replicated.
  - Regional disks can be replicated only across two zones in the same region. One of these zones must be the same as the zone of the destination instance.
  - If the region or zone of the destination instance changes, all regional disks are automatically converted to zonal disks. In this case, the procedure of replicating the disks must be performed again.

### Procedure

1. In the Instances panel, click the instance whose disks you want to restore to open the Detail view.



 **Note** The Detail view appears only if you click an instance. Selecting the check box before the name of the instance will not open the Detail view.

2. In the Detail view that appears at the bottom of the screen, select the preferred restore point.
3. Click  **Restore Instance**. The Restore Options dialog box opens.
4. Select **Clone Options**, and then click **Next**.
5. Select **Clone Disks**, and then click **Next**.
6. From the list of disks that are available for the restore, select the ones that you want to restore, and then click **Next**.
7. From the Restore from drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among which you can select:
  - **Automatic**: This option ensures the fastest and most cost-effective restore.
  - **Backup (Snapshot)**
  - **Backup (Target)**
  - **Copy**
  - **Archive—(daily, weekly, monthly, yearly)**
8. Select the source, the region, and the zone that contain the instance to which you want to attach the restored disks.
9. From the Destination instance drop-down menu, select the instance to which you want to attach the restored disks. If you do not want to attach the disks to an instance, select **None (Leave unattached)**.
10. *Optional*. In the Post-restore script field, enter the path to the script or a command that R-Cloud should run after the restore on the instance to which the restored disks are attached.
 

 **Note** You can enter any command that the command-line interface of your instance supports.
11. Edit the disks as required. For each selected disk, do the following:
  - a. Click  **Edit Disk**.
  - b. *Only if you do not want R-Cloud to automatically generate a name for the restored disk device or disk.* Do the following:
    - i. In the New device name field, enter a name for the restored disk device.
    - ii. In the New disk name field, enter a name for the restored disk.
  - c. If you want to change the disk type, from the Disk type drop-down menu, select one of the available disk types for the restored disk. By default, the original disk type is selected. The following disk types are available:

- For Google Cloud: Standard persistent disk, Balanced persistent disks, and SSD persistent disk.
- For AWS: General Purpose SSD, Previous Generation Volume, and Provisioned IOPS SSD.

If you selected Provisioned IOPS SSD or General Purpose SSD, enter the IOPS number.

- d. *Only if you are restoring Google Cloud instance disks.* If you want to replicate data between two zones in the region of the instance, make sure the **Replicate this disk within region** check box is selected, and then, from the Target Zone drop-down menu, select to which zone you want to replicate data. If the selected disk was regional at backup time, the two zones across which the disk is replicated are shown, otherwise, a list of all zones in the region of the instance is shown.
- e. *Only if you want to add labels to the restored disk.*
  - i. Click **Advanced**.
  - ii. Click  **Manage**. The Custom Metadata dialog box opens.
  - iii. Enter a key and a value, and then click **Add** for each label that you want to add.
    -  **Note** If the selected disk already has one or more labels added, they are listed under Advanced. If you want to delete any of the added labels, click **X Delete** next to it.
- f. Click **Save**.

12. Click **Restore**.

## Moving an instance

You can move instances across different cloud platforms (AWS and Google Cloud) by restoring them to the preferred platform. In this case, the original instance will be kept.

For details on how to move instances to AWS and Google Cloud, see the following sections:

- [“Moving an instance to AWS” below](#)
- [“Moving an instance to Google Cloud” on page 115](#)


### Moving an instance to AWS


#### Limitation


You cannot restore an instance to a different cloud platform from a snapshot.

## Procedure


1. In the Instances panel, click the instance that you want to restore to open the Detail view.

 **Note** The Detail view appears only if you click an instance. Selecting the check box before the name of the instance will not open the Detail view.


2. In the Detail view that appears at the bottom of the screen, select the preferred restore point.
3. Click  **Restore Instance**. The Restore Options dialog box opens.
4. Select **Move Options**, and then click **Next**.
5. Select **Move Instance**, and then click **Next**.
6. From the Restore from drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among which you can select:
  - **Automatic**: This option ensures the fastest and most cost-effective restore.
  - **Backup (Snapshot)**
  - **Backup (Target)**
  - **Copy**
  - **Archive—(daily, weekly, monthly, yearly)**
7. From the Destination source drop-down menu, select the source to which you want to restore the instance. You can choose from sources that belong to the currently selected protection set and that your user account can access.
8. Click **Next**.
9. In the New instance name field, enter the name for the instance.
10. *Only if the original operating system image was not found.* From the Image drop-down menu, select the operating system image you want to use. To use a custom image, select **Use custom image** and enter the image AMI ID.
11. *Optional.* In the Post-restore script field, enter the path to the script or a command that R-Cloud should run on the restored instance after the restore.
 



 **Note** You can enter any command that the command-line interface of your instance supports.
12. From the Destination region and Destination zone drop-down menus, select the AWS region and zone to which you want to restore the instance.

13. In the vCPU cores field, enter the number of virtual CPUs for the restored instance multiplied by the number of cores per virtual CPU.
14. In the Memory field, set the amount of memory (in GiB) for the restored instance. The default value is the amount of memory in GiB of the original instance.
15. From the Instance type drop-down menu, select the instance type for the restored instance.

 **Note** The list shows instance types that match the specified number of virtual CPUs and amount of memory, and the boot type of the instance you are moving to cloud (BIOS or UEFI). If no instance type exactly corresponds to the specified values, the closest matches are shown.


16. Under Disk name, do the following:
  - a. Select the disks that you want to restore.

 **Note** All disks of the instance are preselected for the restore. With such selection, the entire instance is restored. The boot disk is restored even if you do not select it.

- b. Edit the disks as required. For each selected disk, do the following:
      - i. Click  **Edit Disk**.
      - ii. *Only if you do not want R-Cloud to automatically generate a name for the restored disk device or disk. Do the following:*
        - I. In the New device name field, enter a name for the restored disk device.
        - II. In the New disk name field, enter a name for the restored disk.
      - iii. If you want to change the disk type, from the Disk type drop-down menu, select one of the available disk types for the restored disk. By default, the original disk type is selected.  
  
The list shows only the disk types that match the required disk size and can include the following disk types: General Purpose SSD, Previous Generation Volume, and Provisioned IOPS SSD.  
  
If you selected Provisioned IOPS SSD or General Purpose SSD, enter the IOPS number.
      - iv. *Only if you want to add labels to the restored disk.*
        - I. Click **Advanced**.
        - II. Click  **Manage**. The Custom Metadata dialog box opens.



- III. Enter a key and a value, and then click **Add** for each label that you want to add.

 **Note** If the selected disk already has one or more labels added, they are listed under Labels. If you want to delete any of the added labels, click **X Delete** next to it.


- v. Click **Save**.

17. Under Network interfaces, review the list of networks that are available in the selected AWS zone.


For each configured network interface, you can separately adjust its public and private IP address types. By default, the public IP address configuration of the original instance is kept.


### Modifying network settings

If you want to modify network settings, you can add an additional network interface, edit an existing network interface, or delete a network interface:


- Click **Add network interface** to add a network interface or click  **Edit** next to the network interface that you want to edit, and then follow these steps:
  - a. From the Subnet drop-down menu, select the subnet.
  - b. From the Security group drop-down menu, select the security group.
  - c. In the Public address type field, select the public IP address for the network interface. You can select among the following options:



Option	Description
None	<p>The network interface does not use a public IP address.</p> <p>This option is preselected if the network interface of the original instance did not use a public IP address.</p>
Auto-assign	<p>The network interface uses an automatically allocated public IP address.</p> <p>This option is preselected if the network interface of the original instance used a public IP address.</p>

 **Note** Auto-assign will not work if the Auto-


Option	Description
	assign public IPv4 address on a subnet option is set to No or if more than one network interface is specified.
Elastic IP (Reserved)	The network interface uses an elastic public IP address that was reserved in Amazon EC2 in advance.
Elastic IP (New)	The network interface uses a new elastic public IP address.   <b>Note</b> Allocation of the IP address in Amazon EC2 is performed at the very beginning of the restore. If the allocation fails, the restore task is terminated without being logged.

- d. In the Private address type field, select the private IP address for the network interface. You can select between the following options:

Option	Description
Auto-assign	The network interface uses an automatically allocated private IP address.  This option is selected by default.
Custom	The network interface uses a private IP address that is defined by you.   <b>Important</b> Use of this option might result in IP address conflicts.

- e. Click **Add** or **Save**.
- Click  **Delete** next to the network interface that you want to delete. Keep in mind that you cannot restore the instance without a network interface.
18. *Only if you want to add tags to the restored instance.*
- Click **Advanced**.
  - For each custom metadata tag that you want to add, click  **Manage**.

- c. Enter a key and a value, and then click **Add**.

 **Note** If the selected instance already has one or more custom metadata tags added, they are listed under Custom metadata. If you want to delete any of the added custom metadata tags, click **X Delete** next to it.

19. Click **Restore**.


### Moving an instance to Google Cloud


#### Limitation

You cannot restore an instance to a different cloud platform from a snapshot.


#### Procedure

1. In the Instances panel, click the instance that you want to restore to open the Detail view.


 **Note** The Detail view appears only if you click an instance. Selecting the check box before the name of the instance will not open the Detail view.

2. In the Detail view that appears at the bottom of the screen, select the preferred restore point.
3. Click  **Restore Instance**. The Restore Options dialog box opens.
4. Select **Move Options**, and then click **Next**.
5. Select **Move Instance**, and then click **Next**.
6. From the Restore from drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among which you can select:
  - **Automatic**: This option ensures the fastest and most cost-effective restore.
  - **Backup (Snapshot)**
  - **Backup (Target)**
  - **Copy**
  - **Archive—(daily, weekly, monthly, yearly)**
7. From the Destination source drop-down menu, select the source to which you want to restore the instance. You can choose from sources that belong to the currently selected protection set and that your user account can access.
8. Click **Next**.
9. In the New instance name field, specify a new name for the instance.


10. *Optional.* In the Post-restore script field, enter the path to the script or a command that R-Cloud should run on the restored instance after the restore.


 **Note** You can enter any command that the command-line interface of your instance supports.


11. From the Destination region and Destination zone drop-down menus, select the Google Cloud region and zone to which you want to restore the instance.
12. In the vCPU cores field, enter the number of virtual CPUs for the restored instance multiplied by the number of cores per virtual CPU.
13. In the Memory field, set the amount of memory (in GiB) for the restored instance.
14. From the Instance type drop-down menu, select the instance type for the restored instance.



 **Note** The list shows instance types that match the specified number of virtual CPUs and amount of memory, and the boot type of the instance you are moving to cloud (BIOS or UEFI). If no such match exists, you can select the custom machine type. For more information about machine types, see Google Cloud documentation.

15. Under Disk name, do the following:
  - a. Select the instance disks that you want to restore.

 **Note** All disks of the instance are preselected for the restore. With such selection, the entire instance is restored. The boot disk is restored even if you do not select it.

- b. Edit the disks as required. For each selected disk, do the following:
      - i. Click  **Edit Disk**.
      - ii. *Only if you do not want R-Cloud to automatically generate a name for the restored disk device or disk.* Do the following:
        - I. In the New device name field, enter a name for the restored disk device.
        - II. In the New disk name field, enter a name for the restored disk.
      - iii. If you want to change the disk type, from the Disk type drop-down menu, select one of the available disk types for the restored disk (Standard persistent disk, Balanced persistent disks, or SSD persistent disk). By default, the original disk type is selected.
      - iv. *Only if you want to add labels to the restored disk.*

- I. Click **Advanced**.
- II. Click  **Manage**. The Custom Metadata dialog box opens.
- III. Enter a key and a value, and then click **Add** for each label that you want to add.

 **Note** If the selected disk already has one or more labels added, they are listed under Labels. If you want to delete any of the added labels, click  **Delete** next to it.


- v. Click **Save**.

16. Under Network interfaces, review the list of networks that are available in the selected Google Cloud zone.

For each configured network interface, you can separately adjust its public and private IP address types. By default, the public IP address configuration of the original instance is kept.

#### Modifying network settings

If you want to modify network settings, you can add an additional network interface, edit an existing network interface, or delete a network interface:





- Click **Add network interface** to add a network interface or click  **Edit** next to the network interface that you want to edit, and then follow these steps:
  - a. From the Destination network drop-down menu, select the destination network.
  - b. In the Public address type field, select the public IP address for the network interface. You can select among the following options:

Option	Description
None	<p>The network interface does not use a public IP address.</p> <p>This option is preselected if the network interface of the original instance did not use a public IP address.</p>
Ephemeral	<p>The network interface uses an automatically allocated public IP address.</p> <p>This option is preselected if the network interface of the original instance used a public IP address.</p>

Option	Description
Static (Reserved)	The network interface uses a static public IP address that was reserved in Google Compute Engine in advance.
Static (New)	The network interface uses a static public IP address that is allocated at the time of the restore. If the allocation fails, the instance is assigned a temporary public IP address. Such fallback also sets the restore task status to Done with errors.

- c. In the Private address type field, select the private IP address for the network interface. You can select between the following options:

Option	Description
Ephemeral (Automatic)	<p>The network interface uses an automatically allocated private IP address.</p> <p>This option is selected by default for the preselected network interfaces.</p>
Ephemeral (Custom)	<p>The network interface uses a private IP address that is defined by you.</p> <p><b>ⓘ Important</b> Use of this option might result in IP address conflicts.</p>
Static (Reserved)	<i>Not available for legacy networks.</i> The network interface uses a static private IP address that was reserved in Google Compute Engine or in advance.
Static (New)	<p><i>Not available for legacy networks.</i> The network interface uses a new static private IP address that is defined by you.</p> <p><b>📄 Note</b> Allocation of the IP address in Google Compute Engine is performed at the very beginning of the restore. If the allocation fails, the restore task is terminated without being logged.</p>

- d. Click **Add** or **Save**.
  - Click  **Delete** next to the network interface that you want to delete. Keep in mind that you cannot restore the instance without a network interface.
17. *Only if you want to add tags and/or labels to the restored instance.*
- a. Click **Advanced**.
  - b. For each label, network tag, or custom metadata tag that you want to add, click  **Manage**.
  - c. Enter a key and a value, and then click **Add**.
-  **Note** If the selected instance already has one or more labels, network tags, and/or custom metadata tags added, they are listed under Labels, Network tags, or Custom metadata. If you want to delete any of the added labels, network tags, and/or custom metadata tags, click  **Delete** next to it.
18. Click **Restore**.



## Moving disks

You can move disks by restoring them and attaching them to an instance on a different cloud platform, or by restoring them to a different cloud platform and leaving them unattached. In this case, the original disks will be kept.



### Limitation

You cannot restore disks to a different cloud platform from a snapshot.

### Procedure

1. In the Instances panel, click the instance whose disks you want to restore to open the Detail view.
  -  **Note** The Detail view appears only if you click an instance. Selecting the check box before the name of the instance will not open the Detail view.
2. In the Detail view that appears at the bottom of the screen, select the preferred restore point.
3. Click  **Restore Instance**. The Restore Options dialog box opens.
4. Select **Move Options**, and then click **Next**.
5. Select **Move Disks**, and then click **Next**.


6. From the Restore from drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among which you can select:
  - **Automatic:** This option ensures the fastest and most cost-effective restore.
  - **Backup (Snapshot)**
  - **Backup (Target)**
  - **Copy**
  - **Archive—(daily, weekly, monthly, yearly)**
7. From the Destination source drop-down menu, select the source to which you want to restore the instance. You can choose from the sources that belong to the currently selected protection set and that your user account can access.
8. Click **Next**.
9. From the list of disks that are available for the restore, select the ones that you want to restore, and then click **Next**.
10. Select the region and the zone that contain the instance to which you want to attach the restored disks.
11. From the Destination instance drop-down menu, select the instance to which you want to attach the restored disks. If you do not want to attach the disks to an instance, select **None (Leave unattached)**.
12. *Optional.* In the Post-restore script field, enter the path to the script or a command that R-Cloud should run after the restore on the instance to which the restored disks are attached.
 


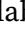
 **Note** You can enter any command that the command-line interface of your instance supports.
13. Edit the disks as required. For each selected disk, do the following:
  - a. Click  **Edit Disk**.
  - b. *Only if you do not want R-Cloud to automatically generate a name for the restored disk device or disk.* Do the following:
    - i. In the New device name field, enter a name for the restored disk device.
    - ii. In the New disk name field, enter a name for the restored disk.
  - c. If you want to change the disk type, from the Disk type drop-down menu, select one of the available disk types for the restored disk. By default, the original disk type is selected. The following disk types are available:



- For Google Cloud: Standard persistent disk, Balanced persistent disks, and SSD persistent disk.
- For AWS: General Purpose SSD, Previous Generation Volume, and Provisioned IOPS SSD.

If you selected Provisioned IOPS SSD or General Purpose SSD, enter the IOPS number.

- d. *Only if you want to add labels to the restored disk.*
  - i. Click **Advanced**.
  - ii. Click  **Manage**. The Custom Metadata dialog box opens.
  - iii. Enter a key and a value, and then click **Add** for each label that you want to add.

 **Note** If the selected disk already has one or more labels added, they are listed under Advanced. If you want to delete any of the added labels, click  **Delete** next to it.

- e. Click **Save**.

14. Click **Restore**.

## Restoring multiple instances or disks belonging to multiple instances in a single session

When you restore multiple instances or disks belonging to multiple instances in a single session, you can select among the following options:

Option	Description	Instructions
Restore Instances	Enables you to restore multiple instances by creating clones of the instances.	“Restoring multiple instances in a single session” on the next page
Restore Disks	Enables you to restore multiple disks on multiple instances at once.	“Restoring multiple disks in a single session” on page 130

Option	Description	Instructions
Restore from JSON	Enables you to upload an existing restore specification to R-Cloud and use it to restore multiple instances or disks.	<a href="#">“Restoring multiple instances or disks from a JSON file” on page 131</a>

## Restoring multiple instances in a single session

You can restore multiple instances by using a restore specification that you create in the R-Cloud web user interface. After the restore specification is created, you can use it immediately or further modify it according to your needs. You can also download the restore specification and use it the next time you want to restore the instances to speed up the restore process.

For details on how to restore AWS and Google Cloud instances in a single session, see the following sections:


- [“Restoring multiple AWS instances” below](#)
- [“Restoring multiple Google Cloud instances” on page 126](#)

### Restoring multiple AWS instances


#### Limitation

You can use only the latest restore point to restore multiple instances.

#### Procedure

1. In the Instances panel, select the instances that you want to restore.
2. Click  **Bulk Restore**. The Bulk Restore Options dialog box opens.
3. Select **Restore Instances**, and then click **Next**.
4. From the Destination Source drop-down menu, select the source to which you want to restore the instances. You can choose from sources that belong to the currently selected protection set and that your user account can access.
5. From the Destination Region and Destination Zone drop-down menus, select the region and the zone to which you want to restore the instances.
6. *Optional*. Enter the destination instance postfix and the target disk postfix to add a postfix to the names of the destination instances and target disks.

7. *Optional.* In the Post-restore Script field, enter the path to the script or a command that R-Cloud should run on the restored instances after the restore.

 **Note** You can enter any command that the command-line interface of your instance supports.

8. Enable the **Overwrite Existing** switch to overwrite the existing instances. By default, this option is disabled and the restore of the instance fails if an instance with the same name exists in the destination zone.
9. *Only if you want to apply custom settings to an instance, or edit or download the restore specification.*


- a. Click **Advanced Settings**.
- b. *Only if you want to apply custom settings to an instance.* Do the following:
  - i. Under Instance Options, from the Instance drop-down menu, select the instance to which you want to apply the custom settings.
  - ii. *Only if you want to rename the instance.* Select **Rename Instance**, and then enter the new name for the instance.
  - iii. From the Destination Source drop-down menu, select the source to which you want to restore the instance.
  - iv. From the Destination Region drop-down menu, select the region to which you want to restore the instance.
  - v. From the Destination Zone drop-down menu, select the zone to which you want to restore the instance.
  - vi. From the Image drop-down menu, select the operating system image you want to use.  
To use a custom image, select **Use custom image** and enter the image AMI ID.
  - vii. *Only if you want to rename a restored disk.* Under Disk Options, do the following:
    - I. From the Disk drop-down menu, select the disk that you want to rename.
    - II. Select **Rename Disk**, and then enter a new name for the disk and click **Confirm**.
  - viii. Under Network Interfaces, review the list of networks that the original instance was configured in at the time of backup. The list shows the following for each such network:


- VPC ID
- Subnet ID


For each configured network interface, you can separately adjust its public and private IP address types. By default, the public IP address configuration of the original instance is kept.

### Modifying network settings


If you want to modify network settings, you can add an additional network interface, edit an existing network interface, or delete a network interface:

- Click **Add Network Interface** to add a network interface or click  **Edit** next to the network interface that you want to edit, and then follow these steps:
  - I. From the Subnet drop-down menu, select the subnet.
  - II. From the Security Groups drop-down menu, select the security groups.
  - III. In the Public Address Type field, select the public IP address for the network interface. You can select among the following options:


Option	Description
None	<p>The network interface does not use a public IP address.</p> <p>This option is preselected if the network interface of the original instance did not use a public IP address.</p>
Auto-assign	<p>The network interface uses an automatically allocated public IP address.</p> <p>This option is preselected if the network interface of the original instance used a public IP address.</p> <p> <b>Note</b> Auto-assign will not work if the Auto-assign public IPv4 address on a subnet option is set to No or if more than</p>

Option	Description
	<p>one network interface is specified.</p>
Elastic IP (Reserved)	The network interface uses an elastic public IP address that was reserved in Amazon EC2 in advance.
Elastic IP (New)	<p>The network interface uses a new elastic public IP address.</p> <p> <b>Note</b> Allocation of the IP address in Amazon EC2 is performed at the very beginning of the restore. If the allocation fails, the restore task is terminated without being logged.</p>

- IV. In the Private Address Type field, select the private IP address for the network interface. You can select between the following options:

Option	Description
Auto-assign	<p>The network interface uses an automatically allocated private IP address.</p> <p>This option is selected by default.</p>
Custom	<p>The network interface uses a private IP address that is defined by you.</p> <p> <b>Important</b> Use of this option might result in IP address conflicts.</p>

- V. Click **Add** or **Save**.


- Click  **Delete** next to the network interface that you want to delete. Keep in mind that you cannot restore the instance without a network interface.

- ix. *Only if you want to edit or download the restore specification.* Do the following:

- I. Click **Edit or download restore JSON**. The restore specification generated by R-Cloud for all selected instances is displayed, and

you can edit it as required.

- II. *Only if you want to download the restore specification.* Click **Download JSON**.
- III. *Only if you want to start the restore.* Click **Restore from .json**.

 **Note** If you want to edit the restore specification by using the REST API Explorer, you can use the URL that is displayed under Request URL.


- x. Click **Save Restore Settings** to apply the custom settings to the instance.
- c. Click **Continue to Summary**.
- d. *Only if you want to download the restore summary.* Click **Download Restore Summary as JSON**.
- e. Review the restore summary, and then click **Start Restore**.


## Restoring multiple Google Cloud instances

### Limitation

You can use only the latest restore point to restore multiple instances.

### Procedure

1. In the Instances panel, select the instances that you want to restore.
2. Click  **Bulk Restore**. The Bulk Restore Options dialog box opens.
3. Select **Restore Instances**, and then click **Next**.
4. From the Destination Source drop-down menu, select the source to which you want to restore the instances. You can choose from sources that belong to the currently selected protection set and that your user account can access.
5. From the Destination Region and Destination Zone drop-down menus, select the region and the zone to which you want to restore the instances.
6. *Optional.* Enter the destination instance postfix and the target disk postfix to add a postfix to the names of the destination instances and target disks.
7. *Optional.* In the Post-restore Script field, enter the path to the script or a command that R-Cloud should run on the restored instances after the restore.


 **Note** You can enter any command that the command-line interface of your instance supports.

8. Enable the **Overwrite Existing** switch to overwrite the existing instances. By default, this option is disabled and the restore of the instance fails if an instance with the same name exists in the destination zone.
9. *Only if you want to apply custom settings to an instance, or edit or download the restore specification.*
  - a. Click **Advanced Settings**.
  - b. *Only if you want to apply custom settings to an instance.* Do the following:
    - i. Under Instance Options, from the Instance drop-down menu, select the instance to which you want to apply the custom settings.
    - ii. *Only if you want to rename the instance.* Select **Rename Instance**, and then enter the new name for the instance.
    - iii. From the Destination Source drop-down menu, select the source to which you want to restore the instance.
    - iv. From the Destination Region drop-down menu, select the region to which you want to restore the instance.
    - v. From the Destination Zone drop-down menu, select the zone to which you want to restore the instance.
    - vi. *Only if you want to rename a restored disk.* Under Disk Options, do the following:
      - I. From the Disk drop-down menu, select the disk that you want to rename.
      - II. Select **Rename Disk**, and then enter a new name for the disk and click **Confirm**.
    - vii. Under Network Interfaces, review the list of networks that the original instance was configured in at the time of backup. The list shows the following for each such network:
      - Network type: Subnetwork for VPC networks and shared VPC networks, Legacy for legacy networks
      - Subnetwork name (for VPC networks and shared VPC networks) or network name (for legacy networks)
      - *Only in case of a shared VPC network.* Name of the host project of the network

For each configured network interface, you can separately adjust its public and private IP address types. By default, the public IP address configuration of the original instance is kept.

## Modifying network settings

If you want to modify network settings, you can add an additional network interface, edit an existing network interface, or delete a network interface:

- Click **Add Network Interface** to add a network interface or click  **Edit** next to the network interface that you want to edit, and then follow these steps:
  - I. From the Destination Network drop-down menu, select the destination network.
  - II. In the Public Address Type field, select the public IP address for the network interface. You can select among the following options:


Option	Description
None	<p>The network interface does not use a public IP address.</p> <p>This option is preselected if the network interface of the original instance did not use a public IP address.</p>
Ephemeral	<p>The network interface uses an automatically allocated public IP address.</p> <p>This option is preselected if the network interface of the original instance used a public IP address.</p>
Static (Reserved)	<p>The network interface uses a static public IP address that was reserved in Google Compute Engine in advance.</p>
Static (New)	<p>The network interface uses a static public IP address that is allocated at the time of the restore. If the allocation fails, the instance is assigned a temporary public IP address. Such fallback also sets the restore task status to Done with errors.</p>



- III. In the Private address type field, select the private IP address for the network interface. You can select between the following options:

Option	Description
Ephemeral (Automatic)	The network interface uses an automatically allocated private IP address. This option is selected by default for the preselected network interfaces.
Ephemeral (Custom)	The network interface uses a private IP address that is defined by you. <b>⚠ Important</b> Use of this option might result in IP address conflicts.
Static (Reserved)	<i>Not available for legacy networks.</i> The network interface uses a static private IP address that was reserved in Google Compute Engine or in advance.
Static (New)	<i>Not available for legacy networks.</i> The network interface uses a new static private IP address that is defined by you. <b>📄 Note</b> Allocation of the IP address in Google Compute Engine is performed at the very beginning of the restore. If the allocation fails, the restore task is terminated without being logged.

- IV. Click **Add** or **Save**.


- Click  **Delete** next to the network interface that you want to delete. Keep in mind that you cannot restore the instance without a network interface.

- viii. *Only if you want to edit or download the restore specification.* Do the following:

- I. Click **Edit or download restore JSON**. The restore specification generated by R-Cloud for all selected instances is displayed, and you can edit it as required.

II. *Only if you want to download the restore specification.* Click **Download JSON**.

III. *Only if you want to start the restore.* Click **Restore from .json**.

 **Note** If you want to edit the restore specification by using the REST API Explorer, you can use the URL that is displayed under Request URL.

ix. Click **Save Restore Settings** to apply the custom settings to the instance.

c. Click **Continue to Summary**.

d. *Only if you want to download the restore summary.* Click **Download Restore Summary as JSON**.

e. Review the restore summary, and then click **Start Restore**.


## Restoring multiple disks in a single session

You can restore disks belonging to multiple instances by using a single restore specification. After the restore specification is generated, you can use it immediately or further modify it according to your needs. You can also download the restore specification and use it the next time you want to restore multiple disks to speed up the restore process.


### Limitation

You can use only the latest restore point to restore disks belonging to multiple instances. Other restore points are not available.

### Procedure

1. In the Instances panel, select the instances whose disks you want to restore.
2. Click  **Bulk Restore**. The Bulk Restore Options dialog box opens.
3. Select **Restore disks**, and then click **Next**. The Restore Disks dialog box opens.
4. From the Destination source drop-down menu, select the source to which you want to restore the disks. You can choose from sources that belong to the currently selected protection set and that your user account can access.
5. From the Destination region and Destination zone drop-down menus, select the region and the zone to which you want to restore the disks.
6. *Optional.* Enter the target disk postfix to add a postfix to the names of the target disks.


7. Enable the **Overwrite existing** switch if you want to overwrite existing disks. By default, this option is disabled and the restore fails if a disk with the same name exists at the instance to which the disks are attached.
8. *Only if you want to rename a disk, or edit or download the restore specification.*
  - a. Click **Advanced Settings**.
  - b. *Only if you want to rename a restored disk.* Do the following:
    - i. From the Instance drop-down menu, select the instance to which the disk you want to rename is attached.
    - ii. Under Disk options, from the Disk drop-down menu, select the disk that you want to rename.
    - iii. Select **Rename Disk**, and then enter the new name for the disk and click **Confirm**.
    - iv. Click **Save Restore Settings** to rename the disk.
  - c. *Only if you want to edit the restore specification.* Do the following:
    - i. Click **Advanced**.
    - ii. Click **Edit or download restore JSON**. The restore specification generated by R-Cloud for all selected instances is displayed, and you can edit it as required.
    - iii. *Only if you want to download the restore specification.* Click **Download JSON**.
    - iv. *Only if you want to start the restore.* Click **Restore from .json**.
 

 **Note** If you want to edit the restore specification by using the REST API Explorer, you can use the URL that is displayed under Request URL.
  - d. Click **Continue to Summary**.
  - e. *Only if you want to download the restore summary.* Click **Download Restore Summary as JSON**.
  - f. Review the restore summary, and then click **Start Restore**.

## Restoring multiple instances or disks from a JSON file

If you previously downloaded a restore specification, you can use it to restore multiple instances or disks by uploading the JSON file to R-Cloud and restoring directly from it.

## Procedure

1. In the Instances panel, select the instances that you want to restore or the instances whose disks you want to restore.
2. Click  **Bulk Restore**. The Bulk Restore Options dialog box opens.
3. Select **Restore from JSON**, and then click **Next**.
4. Under Restore JSON, click **Browse**. Browse for and then select the JSON file that you want to use for the restore.
5. Click **Start restore**.

# Restoring individual files or folders

You can restore one or more individual files or folders to an instance or to a bucket.

Depending on where you want to restore individual files or folders, see one of the following sections:

- [“Restoring files or folders to an instance” on the next page](#)
- [“Restoring files or folders to a bucket” on page 138](#)

## Limitation

If you want to restore individual files or folders that are stored in the S3 Glacier Flexible Retrieval or S3 Glacier Deep Archive storage class, the data must be rehydrated first.

## Considerations

- Two views are available when choosing the files or folders to be restored:
  - The default view allows you to see the file structure as you would from within the operating system. This view is available for newly created backups if the instance was powered on and had credentials assigned during the backup.
  - The filesystem view is always available—even if the default view is not. In this view, the mount point related information from the operating system is not available and the view is suitable also for performing the restores from the file systems which are not mounted within the operating system.


- Some symbolic links in the filesystem view might point to a non-existing target because the file systems do not get mounted to the same location as they are mounted within the operating system.
- For details on how the restored individual files or folders are named, see [“Resources created by R-Cloud” on page 242](#).

### Consideration when restoring symbolic links

Depending on the file or folder and the location that you select for the restore, symbolic links and their targets are restored in the following ways:

Selected file or folder	Restore location	Restored resource
Symbolic link	Original location on the original instance	Symbolic link and its target
	<ul style="list-style-type: none"> <li>• Alternate location on the original instance</li> <li>• Custom location on a different instance</li> <li>• Bucket</li> </ul>	Symbolic link's target
	Any location on the original or a different instance	Symbolic link
Folder that includes a symbolic link	Bucket	Symbolic link's target

#### Accessing the Instances panel


To access the Instances panel, in the navigation pane, click  **Instances**.

## Restoring files or folders to an instance

You can restore one or more individual files or folders to the original or an alternate location on the original or a different instance. For details, see the following sections:

- [“Restoring files or folders to the original instance” on page 135](#)
- [“Restoring files or folders to a different instance” on page 136](#)

## Prerequisites

- The instance to which you plan to restore data must be up and running.
- The discovery status of the instance to which you plan to restore data must be .
- The files or folders that you plan to restore must reside on a supported file system. For details, see the *HYCU R-Cloud Compatibility Matrix*.
- A credential group must be assigned to the instance to which you plan to restore data, and the related credentials must belong to a user account with sufficient privileges:
  - *For Windows:* User from the Administrators group
  - *For Linux:* User with sudo privileges and the NOPASSWD option set

For instructions on how to assign access credentials, see [“Enabling access to data” on page 51](#).

- *For instances running Linux:* The cifs-utils package must be installed.
- To enable file transfers between the temporary instance and the instance to which you plan to restore data, the following prerequisites must be fulfilled:
  - Outbound port 3260 (Windows) or 445 (Linux) must be open on the instance to which you plan to restore data.
  - Inbound port 3260 (Windows) or 445 (Linux) must be open on the temporary instance. You can achieve this in the following ways:
    - Open port 3260 (Windows) or 445 (Linux) on the instance to which you want to restore data. The temporary instance will inherit inbound rules from the instance.
    - *For Google Cloud instances:* Do one of the following:
      - Create a firewall rule that opens the ports for the instances containing the `hycu-iscsi` or the `hycu-cifs` tags.
      - Assign the following permissions to the HMSA: `compute.firewalls.create`, `compute.firewalls.get`, `compute.firewalls.list`, and `compute.firewalls.delete`.
    - *For AWS instances:* Create an IAM policy that contains the following permissions: `ec2:CreateSecurityGroup`, `ec2:AuthorizeSecurityGroupEgress`, `ec2:RevokeSecurityGroupEgress`, `ec2>DeleteSecurityGroup`, and `ec2:ModifyInstanceAttribute`. The IAM policy can be limited to the instances containing the `hycu-worker=true` label and to the security groups containing the `hycu-iscsi=true` or the `hycu-cifs=true` labels.

## Restoring files or folders to the original instance


### Limitation


You cannot restore data to the original location by using the filesystem view if the storage configuration of the instance was changed after the backup in any of the following ways:


- If any disks were added or removed.
- If any partitions were added to or removed from the disk on which the data that you plan to restore is located.
- *For Linux instances:* If the name of a logical volume or the name of a volume group was changed.

### Procedure

1. In the Instances panel, click the instance that contains the files or folders that you want to restore. The Detail view appears at the bottom of the screen.

 **Note** The Detail view appears only if you click an instance. Selecting the check box before the name of the instance does not open the Detail view.

2. In the Detail view, select the preferred restore point, and then click  **Restore Files**.
3. Select **Restore to Instance**, and then click **Continue**.
4. In the Restore Files to Instance dialog box, do the following:
  - a. From the Restore From drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among which you can select:
    - **Automatic:** This option ensures the fastest restore.
    - **Backup (Snapshot)**
    - **Backup (Target)**
    - **Copy**
    - **Archive—(daily, weekly, monthly, yearly)**
  - b. Click **Next**.
5. In the Choose Files and Folders dialog box, from the list of available files and folders, select the ones that you want to restore.

 **Note** If both views are available, you can switch between the default view and the filesystem view. For details, see [“Considerations” on](#)

| [page 132](#).

If needed, click **<** or **>** to move between the pages, or enter a page number to go directly to that page. You can also adjust the number of items shown in a page.

**Tip** You can also search for a file or a folder by entering its name in the Search field and then pressing **Enter**.

6. Click **Next**.
7. Select whether you want to restore the files to the original or to an alternate location.

If you select an alternate location, specify the path in the following format:

- Linux:

```
/<Path>/<FolderName>
```

- Windows:

```
<DriveLetter>:\<Path>\<FolderName>
```

8. Specify which action should be performed during the restore operation if a file with the same name already exists in the selected location (overwrite the file, skip the file, rename the original file, or rename the restored file).  
For naming conventions, see [“Resources created by R-Cloud” on page 242](#).
9. Use the **Restore ACL** switch if you want to restore the original access control list. If enabled, R-Cloud preserves original ACLs. If disabled, R-Cloud applies inherited ACLs on the restored files (according to the file system ACL inheritance rules).
10. Click **Restore**.

## Restoring files or folders to a different instance

### Limitations



- You can restore data only to an instance that is running the same type of operating system as the original instance and belongs to the same protection set as the original instance.
- You can restore data to the original location by using the filesystem view only if the original instance and the instance to which you want to restore data have the same storage configuration (the number of disks, the number of partitions on each disk, the type of disks (MBR, GPT, or RAW), and, for Linux instances, the LVM layout).






- *For Google Cloud instances:* You cannot restore data to an instance that resides on a different cloud platform by using the Snapshot tier.


## Procedure

1. In the Instances panel, click the instance that contains the files or folders that you want to restore. The Detail view appears at the bottom of the screen.
 

 **Note** The Detail view appears only if you click an instance. Selecting the check box before the name of the instance does not open the Detail view.
2. In the Detail view, select the preferred restore point, and then click  **Restore Files**.
3. Select **Restore to Instance**, and then click **Continue**.
4. In the Restore Files to Instance dialog box, do the following:
  - a. From the Restore From drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among which you can select:
    - **Automatic:** This option ensures the fastest restore.
    - **Backup (Snapshot)**
    - **Backup (Target)**
    - **Copy**
    - **Archive—(daily, weekly, monthly, yearly)**
  - b. Select **Restore to a different instance**.
  - c. From the Instance drop-down menu, select the instance to which you want to restore data.
  - d. Click **Next**.
5. In the Choose Files and Folders dialog box, from the list of available files and folders, select the ones that you want to restore.

 **Note** If both views are available, you can switch between the default view and the filesystem view. For details, see [“Considerations” on page 132](#).

If needed, click  or  to move between the pages, or enter a page number to go directly to that page. You can also adjust the number of items shown in a page.

 **Tip** You can also search for a file or a folder by entering its name in the Search field and then pressing **Enter**.

6. Click **Next**.
7. Select whether you want to restore the files to the original or to an alternate location.

If you select an alternate location, specify the path in the following format:

- Linux:

```
/<Path>/<FolderName>
```

- Windows:

```
<DriveLetter>:\<Path>\<FolderName>
```

8. Specify which action should be performed during the restore operation if a file with the same name already exists in the selected location (overwrite the file, skip the file, rename the original file, or rename the restored file).

For naming conventions, see [“Resources created by R-Cloud” on page 242](#).

9. Use the **Restore ACL** switch if you want to restore the original access control list. If enabled, R-Cloud preserves original ACLs. If disabled, R-Cloud applies inherited ACLs on the restored files (according to the file system ACL inheritance rules).
10. Click **Restore**.

## Restoring files or folders to a bucket

### Prerequisite

At least one Amazon S3 or Google Cloud bucket must be available in the protection set that includes the source of the original instance.

### Limitations


- You cannot restore data to a bucket that has Object Lock (WORM) enabled.
- *For Google Cloud instances:* You cannot restore data to a bucket that resides on a different cloud platform by using the Snapshot tier.


### Consideration


If a file with the same name as the one you are restoring already exists in the location that you select for the restore, it will be overwritten.



## Procedure


1. In the Instances panel, click the instance that contains the files or folders that you want to restore. The Detail view appears at the bottom of the screen.

 **Note** The Detail view appears only if you click an instance. Selecting the check box before the name of the instance does not open the Detail view.

2. In the Detail view, select the desired restore point, and then click  **Restore Files**.
3. Select **Restore to Bucket**, and then click **Continue**.
4. In the Restore to Files to Bucket dialog box, do the following:
  - a. From the Restore From drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among which you can select:
    - **Automatic**: This option ensures the fastest restore.
    - **Backup (Snapshot)**
    - **Backup (Target)**
    - **Copy**
    - **Archive—(daily, weekly, monthly, yearly)**
  - b. From the Bucket drop-down menu, select the bucket to which you want to restore data.
  - c. Click **Next**.
5. In the Choose Files and Folders dialog box, from the list of available files and folders, select the ones that you want to restore.

 **Note** If both views are available, you can switch between the default view and the filesystem view. For details, see [“Considerations” on page 132](#).

If needed, click  or  to move between the pages, or enter a page number to go directly to that page. You can also adjust the number of items shown in a page.

 **Tip** You can also search for a file or a folder by entering its name in the Search field and then pressing **Enter**.

6. Click **Next**.
7. *Optional*. Specify the exact path on the bucket in the following format:

```
/<Path>/<FolderName>
```

If you do not specify the path, the files or folders will be restored to the root folder of the bucket.

8. Click **Restore**.

# Chapter 7

## Protecting buckets

R-Cloud enables you to protect your data in Amazon S3 and Google Cloud buckets with fast and reliable backup and restore operations. After you optionally configure bucket backup options and back up a bucket, you can choose to restore one or more individual files or folders inside the bucket.

### Prerequisites

*For Google Cloud:*

- The HYCU Managed Service Account (HMSA) must have the Compute Admin, Service Account User, and Storage Admin roles granted on the projects with the buckets that you plan to protect. For instructions on how to grant permissions to service accounts, see Google Cloud documentation.
- Cloud Resource Manager API, Compute Engine API, Cloud Identity and Access Management API, Cloud Billing API, and Cloud Storage API must be enabled on the Google Cloud projects that contain the buckets that you want to protect. For instructions on how to enable APIs, see Google Cloud documentation.

### Limitations

- Bucket data (backup data, copies of backup data, and data archives) can be stored only to targets, and not as a snapshot. For instructions on how to set up targets, see [“Setting up targets” on page 27](#).
- Protecting data in S3 compatible buckets is not supported.

### Considerations

- Keep in mind that the role you have assigned determines what kind of actions you can perform. For details on roles, see [“Managing roles” on page 225](#).
- *For Google Cloud:* R-Cloud uses an external IP address to access Google Cloud APIs if Private Google Access is disabled on subnets. If your data protection environment requires the use of an internal IP address, make sure Private

Google Access is enabled on subnets. For details, see Google Cloud documentation.

For details on how to protect bucket data efficiently, see the following sections:

- [“Configuring bucket backup options” below](#)
- [“Backing up buckets” on page 145](#)
- [“Restoring buckets” on page 146](#)

## Configuring bucket backup options

Before you start protecting data in buckets, you can adjust bucket protection to the needs of your data protection environment by configuring bucket backup options.

### Backup options

Backup option	Description
Pre/post scripts	Enables you to specify the pre-backup and post-backup scripts to perform necessary actions before and/or after the backup of a bucket is performed.
Temporary instance configuration	Enables you to specify the location and the subnet where you want R-Cloud to create a temporary instance during the backup. For Amazon S3 buckets, you can also specify the security group for the temporary instance. By default, the temporary instance is created in the original AWS account or Google Cloud project of the bucket.

### Prerequisites

- *Only if you plan to specify pre-backup and post-backup scripts.*
  - The `#!/usr/bin/env python3` header must be specified in the script.
  - For Google Cloud:
    - The HYCU Managed Service Account (HMSA) must have access to the bucket where the script is located.
    - *Only if using a service account for running the scripts.* The following line of code must be present in the script:

```
os.environ['GOOGLE_APPLICATION_CREDENTIALS'] =
'/tmp/hycu/serviceAccount.json'
```

- *Only if you plan to configure backup options for multiple buckets.* All buckets must have the same values set for each option that you plan to configure.

## Limitations

*Only if you plan to specify pre-backup and post-backup scripts.*

- Only Python scripts are supported.
- *For AWS:* The pre-backup and post-backup scripts must be located in the same account and in the same region as the bucket.
- *For Google Cloud:* Only the `googleapiclient` Python library can be used for making Google Cloud API calls.

## Considerations

*Only if you plan to configure the temporary instance.* If not specified otherwise, the temporary instance will be created:


- *For AWS:* In the same region as the bucket (for example, US-EAST-1).
- *For Google Cloud:* In the following region (based on the location type of the bucket):
  - *The region:* In the same region as the bucket (for example, US-CENTRAL1).
  - *The dual-region:*

Dual-region name	Temporary instance region
ASIA1	ASIA-NORTHEAST1
EUR4	EUROPE-NORTH1
NAM4	US-CENTRAL1


- *The multi-region:*

Multi-region name	Temporary instance region
ASIA	ASIA-EAST1
EU	EUROPE-WEST1
US	US-CENTRAL1

## Accessing the Buckets panel

To access the Buckets panel, in the navigation pane, click  **Buckets**.

### Procedure

1. In the Buckets panel, select the buckets for which you want to configure backup options.
2. Click  **Configuration**. The Bucket Configuration dialog box opens.
3. Depending on what you want to do, provide the required information:
  - *Only if you want to specify the pre-backup and post-backup scripts.* On the Pre/post Scripts tab, specify the scripts to perform necessary actions before and/or after the backup of the bucket is performed:
    - In the Pre-backup Script field, enter the path to the script that R-Cloud will run just before it performs the backup of the bucket.
    - In the Post-backup Script field, enter the path to the script that R-Cloud will run immediately after it performs the backup of the bucket.

**ⓘ Important** When entering the path to the script, make sure to enter it correctly, including lowercase and uppercase letters, as the path is case sensitive. You must specify the path in the following format:

- *For AWS:* s3://bucket-name/script.py parameter1 parameter2 ...
- *For Google Cloud:* gs://bucket-name/script.py parameter1 parameter2 ...

**Example** The following is an example of the first lines of a pre-backup script for a Google Cloud bucket:


```
#!/usr/bin/env python3
import os
import googleapiclient.discovery

os.environ['GOOGLE_APPLICATION_CREDENTIALS'] =
'/tmp/hycu/serviceAccount.json'

storage = googleapiclient.discovery.build('storage',
'v1')
```



- *Only if you want to configure the temporary instance.* On the Temporary Instance Configuration tab, provide the following information:
    - a. From the Region drop-down menu, select the preferred region.

 **Note** It is recommended that you select the same region as the one where the bucket resides. Otherwise, you will be charged for outbound data transfer. For details, see Amazon S3 or Google Cloud pricing.
    - b. From the Subnet drop-down menu, select the preferred subnet. By default, the temporary instance is created in the default subnet of the preferred region and zone.
    - c. *For Amazon S3 buckets:* Optionally, from the Security Group drop-down menu, select the preferred security group. By default, the temporary instance is created in the default security group of the preferred subnet.
4. Click **Save**.

## Backing up buckets

With R-Cloud, you can back up your Amazon S3 and Google Cloud bucket data securely and efficiently.

### Prerequisites

- *For Google Cloud:*
  - If you plan to back up buckets for which a Shared VPC subnet is specified in configuration, your user account or the HYCU Managed Service Account (HMSA) must be granted the following permissions in the Shared VPC host project: `compute.firewalls.list`, `compute.networks.list`, `compute.networks.get`, `compute.subnetworks.list`, `compute.subnetworks.use`, and `compute.subnetworks.get`.
  - *Only if the target defined in the policy that you plan to assign to buckets has a service account other than the HMSA specified.* The service account must be granted access to the projects with the buckets that you want to protect.


## Limitation

For AWS: Backing up bucket objects that are stored in the S3 Glacier Flexible Retrieval and S3 Glacier Deep Archive storage classes is not supported.


## Consideration

The information on the bucket size becomes available in the Detail view after you assign a policy to the bucket. Keep in mind that this size is always rounded up to the full unit, the minimum being 1 GiB.


### Accessing the Buckets panel


To access the Buckets panel, in the navigation pane, click  **Buckets**.

## Procedure

1. In the Buckets panel, select the buckets that you want to back up. You can update the bucket list by clicking  **Refresh**.

 **Tip** To narrow down the list of displayed buckets, you can use the filtering options as described in [“Filtering and sorting data” on page 181](#).

2. Click  **Set Policy**. The Set Policy dialog box opens.
3. From the list of available policies, select the preferred policy.
4. Click **Assign** to assign the policy to the selected buckets.

 **Important** *Only if one or more buckets that you want to back up reside on a different cloud platform than the target defined in the policy. Click **Assign Anyway** if you want to assign the policy to the bucket. This may result in data transfer fees. To cancel assigning the policy, click **Cancel**.*

After you assign a policy to a bucket, a backup task starts immediately. Subsequent backups are scheduled according to the values defined in the policy.

If required, you can also perform a manual backup of any bucket at any time. For details, see [“Performing manual backups” on page 187](#).

## Restoring buckets

R-Cloud enables you to restore one or more individual files or folders inside an Amazon S3 bucket or a Google Cloud bucket to the original or a different bucket.

## Prerequisite

*For Google Cloud:* If you plan to restore buckets for which a Shared VPC subnet is specified in configuration, your user account or the HYCU Managed Service Account (HMSA) must be granted the following permissions in the Shared VPC host project: `compute.firewalls.list`, `compute.networks.list`, `compute.networks.get`, `compute.subnetworks.list`, `compute.subnetworks.use`, and `compute.subnetworks.get`.

## Limitation

*Only if you plan to restore the original access control list.* The Restore ACL option is not available if you are restoring files:


- To a bucket residing on a different cloud platform.
- Inside Amazon S3 buckets.


## Consideration



For details on how the restored individual files or folders are named, see [“Resources created by R-Cloud” on page 242](#).

## Procedure

1. In the Buckets panel, click the bucket that contains the files or folders that you want to restore. The Detail view appears at the bottom of the screen.

 **Note** The Detail view appears only if you click a bucket. Selecting the check box before the name of the bucket does not open the Detail view.

2. In the Detail view, select the preferred restore point, and then click  **Restore Files**. The File Restore Options dialog box opens.

If needed, click  or  to move between the pages, or enter a page number to go directly to that page. You can also adjust the number of items shown in a page.


3. From the Restore from drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among which you can select:
  - **Automatic:** This option ensures the fastest and most cost-effective restore.
  - **Backup (Target)**
  - **Copy**
  - **Archive—(daily, weekly, monthly, yearly)**
4. Click **Next**. The Choose Files and Folders dialog box opens.



5. From the list of available files and folders, select the ones that you want to restore, and then click **Next**.



If needed, click **<** or **>** to move between the pages, or enter a page number to go directly to that page. You can also adjust the number of items shown in a page.

 **Tip** You can also search for a file or a folder by entering its name in the Search field and then pressing **Enter**.

6. Depending on where you want to restore data, select the preferred restore option, click **Next**, and then follow the instructions:

Restore option	Instructions
<b>Restore to original bucket</b>	<p>a. Select the location on the bucket where you want to restore the files or the folders, and then provide the required information:</p> <ul style="list-style-type: none"> <li>• <b>Original location</b> Specify which action should be performed during the restore operation if a file with the same name already exists in the selected location (overwrite the file, rename the original file, or rename the restored file).</li> <li>• <b>Alternate location</b> Specify the path to an alternate location on the bucket. The restored file overwrites the file with the same name that might exist in the alternate location.</li> </ul> <p>b. Use the <b>Restore ACL</b> switch if you want to restore the original access control list. If enabled, R-Cloud preserves original ACLs. If disabled, R-Cloud applies inherited ACLs on the restored files (according to the ACL permissions at the bucket or source level).</p> <p>c. <i>Only if you want to add custom metadata tags to the restored bucket objects.</i> Click <b>Advanced</b>, and then, in the Advanced section that opens, do the following:</p> <ol style="list-style-type: none"> <li>i. Click  <b>Manage</b>. The Custom Metadata dialog box opens.</li> </ol>

Restore option	Instructions
	<p>ii. Enter a key and a value, and then click <b>Add</b> for each custom metadata tag that you want to add.</p> <p>iii. Click <b>Save</b>.</p> <p> <b>Note</b> If you want to delete any of the added custom metadata tags, click <b>X Delete</b> next to it.</p>
<b>Restore to different bucket</b>	<p>a. From the Source drop-down menu, select the source that contains the bucket to which you want to restore data.</p> <p> <b>Note</b> You can select only among the sources inside the selected protection set.</p> <p>b. From the Bucket name drop-down menu, select the name of the bucket to which you want to restore data, and then click <b>Next</b>.</p> <p>c. Select the location on the bucket where you want to restore the files or folders, and provide the required information:</p> <ul style="list-style-type: none"> <li>• <b>Original location</b> Specify which action should be performed during the restore operation if a file with the same name already exists in the selected location (overwrite the file, rename the original file, or rename the restored file).</li> <li>• <b>Alternate location</b> Specify the path to an alternate location on the bucket.  The restored file overwrites the file with the same name that might exist in the alternate location.</li> </ul> <p>d. Use the <b>Restore ACL</b> switch if you want to restore the original access control list. If enabled, R-Cloud preserves original ACLs. If disabled, R-Cloud applies inherited ACLs on the restored files (according to the ACL permissions at the bucket or source level).</p>

Restore option	Instructions
	<p>e. <i>Only if you want to add custom metadata tags to the restored bucket objects.</i> Click <b>Advanced</b>, and then, in the Advanced section that opens, do the following:</p> <ol style="list-style-type: none"> <li>i. Click  <b>Manage</b>. The Custom Metadata dialog box opens.</li> <li>ii. Enter a key and a value, and then click <b>Add</b> for each custom metadata tag that you want to add.</li> <li>iii. Click <b>Save</b>.</li> </ol> <p> <b>Note</b> If you want to delete any of the added custom metadata tags, click <b>✕ Delete</b> next to it.</p>

7. Click **Restore**.

# Chapter 8

## Performing daily tasks

To ensure your data protection environment is in the optimal state in terms of security, reliability, and efficiency, R-Cloud provides various mechanisms to support your daily activities.

### Consideration

Keep in mind that the role you have assigned determines what kind of actions you can perform. For details on roles, see [“Managing roles”](#) on page 225.


I want to...	Instructions
Get an at-a-glance overview of the data protection environment topology and state, identify eventual bottlenecks, and inspect different areas of the data protection environment.	<a href="#">“Using the R-Cloud dashboard”</a> on the next page
See the visual representation of your data protection platform.	<a href="#">“Exploring R-Graph”</a> on page 215
View the information about my entities.	<a href="#">“Viewing information about entities”</a> on page 153
View policy information, edit a policy, or delete a policy.	<a href="#">“Managing policies”</a> on page 163
View target information, activate or deactivate a target, and edit or remove a target.	<a href="#">“Managing targets”</a> on page 165
Track tasks that are running in the data protection environment and get an insight into the status of a specific task.	<a href="#">“Checking task statuses”</a> on page 169
View all events that occurred in my data protection environment.	<a href="#">“Viewing events”</a> on page 170

I want to...	Instructions
Configure R-Cloud to send notifications when new events occur in my data protection environment.	<a href="#">“Configuring event notifications” on page 172</a>
Obtain R-Cloud reports on different aspects of the data protection environment.	<a href="#">“Using R-Cloud reports” on page 175</a>
Narrow down the list of displayed items by applying filters and sort the items in panels.	<a href="#">“Filtering and sorting data” on page 181</a>
Back up my data manually.	<a href="#">“Performing manual backups” on page 187</a>
Mark a restore point as expired.	<a href="#">“Expiring backups manually” on page 188</a>
Export data that I can view in a table in any of the panels to a JSON or CSV file.	<a href="#">“Exporting the contents of the panel” on page 190</a>
View subscription information.	<a href="#">“Viewing subscription information” on page 190</a>
Customize the R-Cloud web user interface to match your needs.	<a href="#">“Customizing your R-Cloud web user interface” on page 192</a>

## Using the R-Cloud dashboard

The R-Cloud dashboard enables you to monitor your data protection environment, observe the relevant data protection activity, and quickly identify the areas that need your attention. You can use this dashboard as a starting point for your everyday tasks because it enables you to easily access the area of interest by simply clicking the corresponding widget.


### Accessing the Dashboard panel

To access the Dashboard panel, in the navigation pane, click  **Dashboard**.

The following table describes what kind of information you can find within each widget.



Widget	Description
R-Graph	R-Graph is a visual representation of your data protection environment. For details on R-Graph, see <a href="#">“Exploring R-Graph” on page 215</a> .
Policies	Percentage of compliant policies. A policy is considered compliant if all entities to which a policy is assigned are compliant with the policy settings. For details on policies, see <a href="#">“Defining your backup strategy” on page 36</a> .
Targets	Number of targets in the protection set, and the information about how much space is available for storing the backup data. For details on targets, see <a href="#">“Setting up targets” on page 27</a> .
Backups	Number of backups performed per day in the protection set, and the backup task success rate for the last seven days in percentages. For details on backups, see <a href="#">“Defining your backup strategy” on page 36</a> .
Events	Total number of events in the protection set and the number of events according to their severity level (Success, Warning, Failed) in the last 48 hours. For details on events, see <a href="#">“Viewing events” on page 170</a> .

 **Tip** By clicking icons that denote different statuses within a widget, you are automatically taken to the corresponding panel with the data already filtered accordingly.

## Viewing information about entities

You can at any time view information about the entities that are part of your data protection environment. The entities include SaaS applications, Google Cloud applications, instances, and buckets. Depending on whether you want to see generic or detailed information about your entities, see one of the following sections:

- [“Viewing entity information” on the next page](#)
- [“Viewing entity details” on page 159](#)




## Viewing entity information

You can view the information about each entity in the SaaS, Applications, Instances, or Buckets panel.




Type of entity	See...
SaaS application	<a href="#">“Viewing SaaS application information” below</a>
Google Cloud application	<a href="#">“Viewing Google Cloud application information” on the next page</a>
Instance	<a href="#">“Viewing instance information” on page 156</a>
Bucket	<a href="#">“Viewing bucket information” on page 158</a>









## Viewing SaaS application information

SaaS application information	Description
Name	Name of the SaaS application.
Source	R-Cloud module to which the SaaS application belongs.
Policy	Policy that is assigned to the SaaS application.
Compliance	Shows whether the SaaS application is compliant with backup requirements: <ul style="list-style-type: none"> <li>✔ (Success): The time since the last successful backup is lower than the RPO defined in the policy and archiving was successfully finished, if specified in the policy.</li> <li>✘ (Failure): The time since the last successful backup is not lower than the RPO defined in the policy or archiving was not successfully finished.</li> <li>❓ (Undefined): The exclude policy is assigned to the SaaS application or the SaaS application does not have a policy assigned.</li> </ul>
Protection	Shows whether the SaaS application is protected: <ul style="list-style-type: none"> <li>✔ (Protected)</li> </ul>

SaaS application information	Description
	<ul style="list-style-type: none"> <li>•  (Unprotected)</li> <li>•  (Protected deleted): The SaaS application has been deleted from the source, but it has at least one valid restore point available in R-Cloud.</li> <li>•  (Undefined): The exclude policy is assigned to the SaaS application or the information about the SaaS application protection status is not available.</li> </ul>












## Viewing Google Cloud application information

GC application information	Description
Name	Name of the application.
Application type	Type of application—SAP HANA or Google Kubernetes Engine.
Source	Google Cloud project to which the application belongs.
Location	Location of the application.
Policy	Policy that is assigned to the application.
Compliance	<p>Shows whether the application is compliant with backup requirements:</p> <ul style="list-style-type: none"> <li>•  (Success): The time since the last successful backup is lower than the RPO defined in the policy and archiving was successfully finished, if specified in the policy.</li> <li>•  (Failure): The time since the last successful backup is not lower than the RPO defined in the policy or archiving was not successfully finished.</li> <li>•  (Undefined): The exclude policy is assigned to the application or the application does not have a policy assigned.</li> </ul>
Protection	Shows whether the application is protected:








GC application information	Description
	<ul style="list-style-type: none"> <li>•  (Protected)</li> <li>•  (Unprotected)</li> <li>•  (Protected deleted): The application has been deleted from the source, but it has at least one valid restore point available in R-Cloud.</li> <li>•  (Undefined): The exclude policy is assigned to the application or the information about the application protection status is not available.</li> </ul>
Discovery	<p>Shows the discovery status of the application:</p> <ul style="list-style-type: none"> <li>•  (Success)</li> <li>•  (Failure)</li> <li>•  (Warning): Connection to the application has been interrupted. By pausing on the icon, you can see the details of the warning.</li> <li>•  (Undefined): The protection status of the application is Protected deleted or the information about the application discovery status is not available.</li> </ul>
Credentials	<i>Applicable only to SAP HANA applications.</i> Credentials assigned to the application.

## Viewing instance information

Instance information	Description
Name	Name of the instance.
Source type	Type of source to which the instance belongs—AWS or Google Cloud.
Source	AWS account or the Google Cloud project to which the instance belongs.
Zone	Zone of the instance.
Policy	Policy that is assigned to the instance.


Instance information	Description
Compliance	<p>Shows whether the instance is compliant with backup requirements:</p> <ul style="list-style-type: none"> <li>•  (Success): The time since the last successful backup is lower than the RPO defined in the policy and archiving was successfully finished, if specified in the policy.</li> <li>•  (Failure): The time since the last successful backup is not lower than the RPO defined in the policy or archiving was not successfully finished.</li> <li>•  (Undefined): The exclude policy is assigned to the instance or the instance does not have a policy assigned.</li> </ul>
Protection	<p>Shows whether the instance is protected:</p> <ul style="list-style-type: none"> <li>•  (Protected)</li> <li>•  (Unprotected)</li> <li>•  (Protected deleted): The instance has been deleted from the source, but it has at least one valid restore point available in R-Cloud.</li> <li>•  (Undefined): The exclude policy is assigned to the instance or the information about the instance protection status is not available.</li> </ul>
Discovery	<p>Shows the discovery status of the instance:</p> <ul style="list-style-type: none"> <li>•  (Success): Connection to the instance has been established.</li> <li>•  (Failure): Connection to the instance could not be established.</li> <li>•  (Warning): Connection to the instance has been interrupted. By pausing on the icon, you can see the details of the warning.</li> <li>•  (Undefined): The protection status of the instance is Protected deleted or the information about the instance discovery status is not available.</li> </ul>
Credential group	Name of the credential group assigned to the instance.

## Viewing bucket information

Bucket information	Description
Name	Name of the bucket.
Source type	Type of source to which the bucket belongs—AWS or Google Cloud.
Source	AWS account or the Google Cloud project to which the bucket belongs.
Location	Location of the bucket.
Policy	Policy that is assigned to the bucket.
Compliance	<p>Shows whether the bucket is compliant with backup requirements:</p> <ul style="list-style-type: none"> <li>•  (Success): The time since the last successful backup is lower than the RPO defined in the policy and archiving was successfully finished, if specified in the policy.</li> <li>•  (Failure): The time since the last successful backup is not lower than the RPO defined in the policy or archiving was not successfully finished.</li> <li>•  (Undefined): The exclude policy is assigned to the bucket or the bucket does not have a policy assigned.</li> </ul>
Protection	<p>Shows whether the bucket is protected:</p> <ul style="list-style-type: none"> <li>•  (Protected)</li> <li>•  (Unprotected)</li> <li>•  (Protected deleted): The bucket has been deleted from the source, but it has at least one valid restore point available in R-Cloud.</li> <li>•  (Undefined): The exclude policy is assigned to the bucket or the information about the bucket protection status is not available.</li> </ul>





## Viewing entity details

You can view the details about each entity in the Detail view of the SaaS, Applications, Instances, or Buckets panel.

 **Note** The Detail view appears only if you click an entity. Selecting the check box before its name will not open the Detail view.

The following details are available:

Entity detail	Description
Summary	Shows detailed information about the selected entity.
Restore point	<p>Shows the following information for the restore point:</p> <ul style="list-style-type: none"> <li>• Creation date and time.</li> <li>• Available tiers from which you can restore data: <ul style="list-style-type: none"> <li>◦ <i>For SaaS applications, GKE applications, and instances:</i> <ul style="list-style-type: none"> <li>▪ <b>SNAP</b> or <b>S</b>: Snapshot. Displayed if a snapshot of the instance using persistent disks exists. Snapshots allow faster completion of restore tasks.</li> </ul> <p>For SaaS applications, the snapshot tier represents a backup to a staging target or remote storage.</p> <li>▪ <b>BCKP</b> or <b>B</b>: Backup data on a target. Displayed if backup data is stored on a target.</li> <li>▪ <b>COPY</b> or <b>C</b>: Copy of backup data. Displayed if a copy of a backup image (snapshot or backup data on a target) exists on another target.</li> <li>▪ <b>D ARCH</b> or <b>D A</b>: Data archive—daily. Displayed if a daily data archive exists on a target.</li> <li>▪ <b>W ARCH</b> or <b>W A</b>: Data archive—weekly. Displayed if a weekly data archive exists on a target.</li> <li>▪ <b>M ARCH</b> or <b>M A</b>: Data archive—monthly. Displayed if a monthly data archive exists on a target.</li> <li>▪ <b>Y ARCH</b> or <b>Y A</b>: Data archive—yearly. Displayed if a yearly data archive exists on a target.</li> <li>▪ <b>CTLG</b> or <b>C</b>: Catalog. Displayed if a restore of individual files or folders is available. (Available only</li> </li></ul> </li> </ul>

	<p>for instances.)</p> <p> <b>Note</b> A restore point may or may not include backup data of the entire instance. This depends on the disks included in the corresponding backup.</p> <p>Visual labels of the tiers may be specially marked to denote different statuses. For more information, see <a href="#">“Tier statuses” on page 162</a>.</p> <ul style="list-style-type: none"> <li>○ <i>For SAP HANA applications:</i> <ul style="list-style-type: none"> <li>▪ <b>FULL</b>: Full backup.</li> <li>▪ <b>INCR</b>: Incremental backup.</li> </ul> </li> <li>○ <i>For buckets:</i> <ul style="list-style-type: none"> <li>▪ <b>BCKP</b> or <b>B</b>: Backup data on a target.</li> <li>▪ <b>COPY</b> or <b>C</b>: Copy of backup data. Displayed if a copy of a backup data exists on another target.</li> <li>▪ <b>D ARCH</b> or <b>D A</b>: Data archive—daily. Displayed if a daily data archive exists on a target.</li> <li>▪ <b>W ARCH</b> or <b>W A</b>: Data archive—weekly. Displayed if a weekly data archive exists on a target.</li> <li>▪ <b>M ARCH</b> or <b>M A</b>: Data archive—monthly. Displayed if a monthly data archive exists on a target.</li> <li>▪ <b>Y ARCH</b> or <b>Y A</b>: Data archive—yearly. Displayed if a yearly data archive exists on a target.</li> </ul> </li> </ul>
Compliance	<p>Shows the compliance status of the backup (and the resulting restore point):</p> <ul style="list-style-type: none"> <li>•  (Success): The backup is compliant (the RPO setting in the policy assigned to the entity was met).</li> <li>•  (Failure): The backup is not compliant (the RPO setting in the policy assigned to the entity was not met).</li> <li>•  (Undefined): The backup compliance status is undefined (the backup is still running).</li> </ul> <p>By pausing on a compliance status icon, additional information about the backup is available. You can see backup frequency, the elapsed time since the last successful backup, and the expiration time for each available tier.</p>



Backup status	Shows the backup status of your entity. For more information, see <a href="#">“Viewing the backup status of entities”</a> below.
Restore status	Shows a progress bar indicating the progress of the restore for your entity.  <b>Tip</b> If you double-click a progress bar, you are directed to the Tasks panel where you can check details about the related task.

**Tip** To minimize the Detail view, click **Minimize** or press the Spacebar. To return the Detail view to its original size, click **Maximize** or press the Spacebar.

## Viewing the backup status of entities

The backup status of your entities determines whether it is possible to restore them.

Backup status	Restore a SaaS app, a GKE app, an instance, or a disk?	Restore files?	Restore an SAP HANA app?	Restore a bucket?
✓ (Done)	✓	✓ <sup>a</sup>	✓	✓
! (Done with warnings)	✓	✓ <sup>a</sup>	✓	✓
! (Done with errors)	✓ <sup>b</sup>	?	✓ <sup>d</sup>	✓ <sup>e</sup>
⊖ (Aborted)	×	×	×	×
⌚ (Pending)	×	×	×	×
? (Inaccessible)	×	×	×	×
✗ (Failed)	×	×	×	×
○ (Expired, Inaccessible on Source, or Deleted from Source)	×	×	×	×

<sup>a</sup>All disks were backed up successfully, but the disk catalog creation task might have failed. In this case, you will not be able to restore individual files or folders.

<sup>b</sup>This backup status may indicate one of the following:


- Not all entities were backed up successfully. Therefore, the entity can be restored only partially. If backing up a boot disk of an instance failed, you may not be able to start the instance after the restore.
- Creating a copy of backup data or a data archive failed. However, the entity can still be fully restored from the backup.
- The backup is not application-consistent.
- *Applicable only if you are using the pre-backup and post-backup scripts. The script or some actions specified by the script were not executed.*

<sup>c</sup>This backup status may indicate one of the following:

- Not all disks were backed up successfully and the disk catalog creation task might have failed. In this case, you will not be able to restore individual files or folders.
- Not all disks were backed up successfully, therefore only the files belonging to the disks that were successfully backed up can be restored.









<sup>d</sup>An application can be partially restored (only the databases that are displayed in the Restore dialog box).








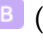
<sup>e</sup>*Applicable only if you are using the pre-backup and post-backup scripts. The script or some actions specified by the script were not executed.*

 **Note** By pausing on a backup status icon, additional information about the restore point is shown. You can see the backup duration and ID.

## Tier statuses

Tier labels may be visually marked to represent backup statuses of individual tiers. These statuses define whether it is possible to restore an entity. The following is an example of possible marks:


Tier status	Restore an entity?
 or  (Done)	✓
 or  (Done with warnings)	✓ For details on what data can be restored if one of these backup statuses is shown, see <a href="#">“Viewing the backup status of entities” on the previous page.</a>
 or  (Done with errors)	✗
 or  (Aborted)	✗

Tier status	Restore an entity?
 or  (Inaccessible on source)	×
 or  (Deleted from the source)	×
 or  (Failed)	×
 or  (Expired)	×

## Managing policies




You can view policy information, edit policy properties, or delete a policy if you do not want to use it for protecting data anymore.

### Accessing the Policies panel

To access the Policies panel, in the navigation pane, click  **Policies**. Alternatively, in the Dashboard panel, click the **Policies** widget title.




## Viewing policy information

You can view information about each policy in the list of policies in the Policies panel.

Property name	Description
Name	Policy name.
Compliance	Compliance status of the policy: <ul style="list-style-type: none"> <li>The  icon: The policy is compliant.</li> <li>The  icon: The policy is non-compliant.</li> <li>The  icon: Policy compliance is undefined. The policy is not assigned to any entity, or this is the exclude policy.</li> </ul>
SaaS Count	Number of the SaaS applications that have the policy assigned to them.

Property name	Description
Instance Count	Number of the instances that have the policy assigned to them.
Application Count	Number of the applications that have the policy assigned to them.
Bucket Count	Number of the buckets that have the policy assigned to them.
Description	Description of the policy.

To open the Detail view where you can find more details about the policy, click the preferred policy.


 **Tip** To minimize the Detail view, click  **Minimize** or press the Spacebar. To return the Detail view to its original size, click  **Maximize** or press the Spacebar.

## Creating a policy

See [“Creating custom policies” on page 37](#).

## Editing a policy

Procedure


1. In the Policies panel, select the policy that you want to edit, and then click  **Edit**. The Edit Policy dialog box appears.
2. Edit the selected policy as required. For details about policy properties, see [“Creating custom policies” on page 37](#).
3. Click **Save**.

## Deleting a policy

Limitation

You cannot delete the exclude policy.


## Procedure

1. In the Policies panel, select the policy that you want to delete, and then click  **Delete**.
2. Click **Delete** to confirm that you want to delete the selected policy.

# Managing targets

You can view target information, edit a target, deactivate or activate a target, or remove a target if you do not want to use it for storing backup data anymore.


## Accessing the Targets panel




To access the Targets panel, in the navigation pane, click  **Targets**.


Alternatively, in the Dashboard panel, click the **Targets** widget title.

## Viewing target information




You can view information about each target in the list of targets in the Targets panel. This allows you to have an overview of the general status of the targets. The following information is available for each target:

Property name	Description
Name	<p>Target name (globally unique).</p> <p>A target that has Object Lock (WORM) enabled is represented by the  icon in the list of targets.</p> <p>For information on how automatically created targets are named, see <a href="#">“Resources created by R-Cloud” on page 242</a>.</p>
Target type	Type of target for storing data protected by R-Cloud. For a list of supported targets, see the <i>HYCU R-Cloud Compatibility Matrix</i> .
Location	Geographical region in which the target resides.
Storage class	Storage classes define the storage availability and pricing model. The default object storage class is displayed for the target.

Property name	Description
	<p>The available options are:</p> <ul style="list-style-type: none"> <li>• <i>For Amazon S3 targets:</i> S3 Standard, S3 Intelligent-Tiering, S3 Standard-IA, S3 One Zone-IA, S3 Glacier Instant Retrieval, S3 Glacier Flexible Retrieval, and S3 Glacier Deep Archive</li> <li>• <i>For Azure targets:</i> Hot and Cool</li> <li>• <i>For Google Cloud targets:</i> Standard, Nearline, Coldline, and Archive</li> <li>• <i>For S3 compatible targets:</i> S3 Standard</li> </ul>
State	<p>Status of the target:</p> <ul style="list-style-type: none"> <li>• <b>Active:</b> You can use the target for backing up data, creating data archives, and restoring data.</li> <li>• <b>Inactive:</b> The target has been deactivated within R-Cloud. Until it is activated, you can use it only for restoring data.</li> <li>• <b>Inaccessible on source:</b> R-Cloud cannot access the target.</li> <li>• <b>Deleted from source:</b> The target no longer exists in cloud.</li> </ul> <p>For instructions on how to change the status of active or inactive targets, see <a href="#">“Deactivating and activating targets” on the next page</a>.</p>
Size limit	<p>Maximum amount of the target storage space (expressed in MiB, GiB, or TiB) that is allowed to be used by backup data created by R-Cloud. The amount represents a soft limit, therefore actual usage may exceed it.</p>
Health	<p>Health status of the target:</p> <ul style="list-style-type: none"> <li>• The  icon: Indicates one of the following: <ul style="list-style-type: none"> <li>◦ The target health has not been determined yet.</li> <li>◦ The target is inactive.</li> </ul> </li> <li>• The  icon: The target is in a healthy state. Utilization of storage space for backup data in the target is less than 90 percent of the configured size limit.</li> <li>• The  icon: Utilization of storage space for backup data in the target is over 90 percent and under 100 percent of the</li> </ul>


Property name	Description
	<p>configured size limit, or the target is publicly accessible in cloud.</p> <ul style="list-style-type: none"> <li>• The  icon: Indicates one of the following: <ul style="list-style-type: none"> <li>◦ Target storage space occupied by backup data exceeds the configured size limit.</li> <li>◦ The target is not accessible due to an I/O error, insufficient permissions, or some other reason.</li> <li>◦ Active lifecycle rules are configured for the target.</li> </ul> </li> </ul>
Utilization	Ratio (expressed in percentage) between the target storage space occupied by backup data and the configured size limit.
Tags	<p>The tag shows if the target:</p> <ul style="list-style-type: none"> <li>• Was created automatically by R-Cloud (Automatic).</li> <li>• Is a staging target for a SaaS application (Staging).</li> </ul>

To open the Detail view where you can find more details about the target, click the preferred target.

 **Tip** To minimize the Detail view, click  **Minimize** or press the Spacebar. To return the Detail view to its original size, click  **Maximize** or press the Spacebar.

## Editing targets

### Procedure

1. In the Targets panel, select the target that you want to edit, and then click  **Edit**. The Edit Target dialog box appears.
2. Edit the selected target as required.
3. Click **Save**.

## Deactivating and activating targets

After you deactivate a target, you can use it only for restoring data. Targets that were created automatically by R-Cloud cannot be deactivated.

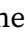
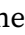
### Prerequisite

*For target deactivation:* The target must not be specified in the Target option of any policy or data archive.

### Considerations

- After deactivating a target, the target cannot be selected for the Target option of a policy until it is activated again.
- Targets that are specified as staging targets for storing SaaS application data cannot be deactivated.

### Procedure

1. In the Targets panel, select the target that you want to deactivate or activate.
2. Change the status of the selected target: Click  **Deactivate** or  **Activate**.
3. *Only if you are deactivating a target.* Click **Deactivate** to confirm that you want to deactivate the selected target.

## Removing targets

You can remove a target from R-Cloud if it does not contain any protected data. After removing a target, no backup or restore actions including this target are possible anymore.


### Prerequisite

The target must not be specified in the Target option of any policy or data archive.

### Considerations

- You cannot remove targets that were created automatically by R-Cloud unless they have been deleted from cloud.
- Targets that are specified as staging targets for storing SaaS application data cannot be removed from R-Cloud.

### Procedure

1. In the Targets panel, select the target that you want to remove, and then click  **Remove**.
2. Click **Remove** to confirm that you want to remove the selected target.



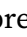





# Checking task statuses

In the Tasks panel, you can do the following:

- Check the overall status of the tasks in your data protection environment.
- Check the status of tasks that are currently running.
- Check the status of completed and stopped tasks.
- Check more details about a specific task.

The information is presented in the Detail view that appears at the bottom of the screen after you select the task.


 **Tip** To minimize the Detail view, click  **Minimize** or press the Spacebar. To return the Detail view to its original size, click  **Maximize** or press the Spacebar.

- Generate a report about a specific task.  
To generate the report, select a task, and then click  **View Task Report**.  
To copy the report to the clipboard, in the View Task Report dialog box that opens, click  **Copy to Clipboard**.
- Cancel any currently running task by selecting it, and then clicking  **Abort Task**.

## Consideration

R-Cloud periodically deletes tasks (as well as all associated task reports and events) from the database. Tasks related to backups, copies of backups, and archives are deleted 90 days after the corresponding restore points are removed from R-Cloud. All other tasks are deleted 90 days after they are created.

### Accessing the Tasks panel

To access the Tasks panel, in the navigation pane, click  **Tasks**.

Alternatively, in the Dashboard panel, click the **Tasks** widget title.


Task information	Description
Description	Summary of the task (for example, running a backup, performing a restore, restoring individual files or folders).
Status	Current status of a task (for example, Ready, a progress bar indicating the Running status, Done, Done with errors, Failed, or Aborted).

Task information	Description
Subtask	The list of subordinate tasks.
Started	The task's start date and time.
Finished	The task's finish date and time.



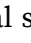
## Viewing events

In the Events panel, you can do the following:

- View all events that occurred in your data protection environment.
- Check more details about a specific event in the Detail view that appears at the bottom of the screen after you select the event.

 **Tip** If you click the related task link in the Detail view, you are directed to the Tasks panel where you can view more details about the related task.


- List the events that match the specified filter.
- Configure R-Cloud to send notifications when new events occur in your data protection environment. For details, see [“Configuring event notifications” on page 172](#).

 **Tip** To minimize the Detail view, click  **Minimize** or press the Spacebar. To return the Detail view to its original size, click  **Maximize** or press the Spacebar.

### Consideration




R-Cloud periodically deletes events from the database. Events related to backups, copies of backups, and archives are deleted 90 days after the corresponding restore points are removed from R-Cloud. All other events are deleted 90 days after they are created.

#### Accessing the Events panel

To access the Events panel, in the navigation pane, click  **Events**. Alternatively, in the Dashboard panel, click the **Events** widget title.

The following information is available for each event:

Severity	Severity level of the event:
----------	------------------------------



	<ul style="list-style-type: none"> <li>•  (Info): Events representing regular service operation.</li> <li>•  (Warning): Potentially harmful situations that do not represent an immediate threat to service operation.</li> <li>•  (Error): Errors that immediately affect service operation.</li> </ul>
Message	Description of the event.
Category	<p>R-Cloud functional area to which the event belongs:</p> <ul style="list-style-type: none"> <li>• Administration: Protection environment changes, such as updated configurations, added/removed sources, or added/removed targets.</li> <li>• Archive: Creation or deletion of archives.</li> <li>• Backup: Events that take place during backup and notifications about skipped backup tasks.</li> <li>• Backup_Window: Events that take place when a backup misses a time period defined for the backup window.</li> <li>• Configuration: Events related to setting backup options for entities.</li> <li>• Credentials: Events related to instance credentials management.</li> <li>• Export: Events related to the export of data from the panels of the web user interface.</li> <li>• Migration: Events related to the SpinUp functionality. For details on the SpinUp functionality, see HYCU R-Cloud Hybrid Cloud Edition documentation.</li> <li>• Notification: Possible failures or system malfunctions.</li> <li>• Policies: Creation, updates, or removal of policies.</li> <li>• Reporting: Events related to report management and generation.</li> <li>• Restore: Events that take place during restore.</li> <li>• IAM: Added or removed users, updates of user roles, and status changes.</li> <li>• System: Events not related to any other category. Events of this type usually take place independently of your interaction with R-Cloud.</li> <li>• Targets: Events related to target management.</li> </ul>
Timestamp	Event creation date and time.

# Configuring event notifications

You can configure R-Cloud to send notifications when new events occur in your data protection environment. This allows you to monitor and manage your data protection environment more efficiently, and to immediately respond to the events if required. You can set up emails or webhooks as a notification channel.

**ⓘ Important** Make sure to configure event notifications for each protection set separately.

## Accessing the Notifications dialog box


To access the Notifications dialog box, click  **Events** in the navigation pane, and then click  **Notifications** in the toolbar.

Depending on which notification channel you want to use, see one of the following sections:

- [“Creating email notifications” below](#)
- [“Creating webhook notifications” on the next page](#)





## Creating email notifications

### Procedure

1. In the Notifications dialog box, click the **Email** tab, and then click  **New**.
2. In the Subject field, enter a subject for the email notification.
3. From the Category drop-down menu, select one or more categories. To include all categories, click **Select All**. For a description of categories, see [“Viewing events” on page 170](#).
4. From the Status drop-down menu, select one or more statuses. To include all statuses, click **Select All**. For a description of statuses, see [“Viewing events” on page 170](#).
5. In the Email address field, enter the recipient's email address. If you are entering more than one email address, make sure to press the Spacebar after entering each one.
6. Click **Save**.


Your changes take effect immediately and email notifications are sent to any email address that you specified in the notification settings.

After you create a notification, you can do the following:

- Edit the email notification settings by clicking  **Edit**.
- Temporarily deactivate the email notification and stop sending emails by clicking  **Deactivate**. To resume sending emails, click  **Activate**.
- Delete the email notification that you no longer need by clicking  **Delete**.


## Creating webhook notifications

### Procedure

1. In the Notifications dialog box, click the **Webhooks** tab, and then click  **New**.
2. Enter a name for the webhook notification and, optionally, its description.
3. From the Category drop-down menu, select one or more event categories. To include all categories, click **Select All**. For a description of categories, see [“Viewing events” on page 170](#).
4. From the Status drop-down menu, select one or more statuses. To include all statuses, click **Select All**. For a description of statuses, see [“Viewing events” on page 170](#).
5. In the Post URL field, enter the URL of the endpoint the webhook notifications should be sent to in one of the following formats:

```
https://<Host>
https://<Host>/<Path>
```

6. *Only if the receiving endpoint requires sender's identification.* From the Authentication type drop-down menu, select one of the following authentication types:
  - **Authentication by secret**, and then enter the secret to connect to your webhook endpoint.
  - **Basic authentication**, and then enter the user name and password associated with your webhook endpoint.
7. Click **Next**.
8. *Optional.* Customize the body of the request that is sent by R-Cloud. You can click the appropriate fields in the HYCU fields list to easily insert event variables into the body.





 **Important** Make sure the format you define in the body is supported by the platform to which webhook notifications will be sent.

For details on the format of the data that R-Cloud sends to the specified URL, see [“Configuring event notifications” on page 172](#).

9. Click **Save**.

Your changes take effect immediately and webhook notifications are sent to any URL that you specified in the notification settings.

After you create a notification, you can do the following:

- Edit the webhook notification settings by clicking  **Edit**.
- Temporarily deactivate the webhook notification and stop sending the notifications by clicking  **Deactivate**. To resume sending the notifications, click  **Activate**.
- Delete the obsolete webhook notification by clicking  **Delete**.

## Webhook data format


The webhook data format is defined by:

- HTTP request header sent by R-Cloud
- HTTP request body sent by R-Cloud
- HTTP response code sent by the webhook endpoint and received by R-Cloud

### HTTP request headers

The request headers are sent in the following format:


```
content-type = application/json
x-hycu-signature = base64(hmac(body, secret, 'sha256'))
```

 **Note** The `x-hycu-signature` request header is sent only if the webhook secret is specified.

### HTTP request body

The request body is sent in the following format:

```
{
  "severity": "<severity-value>",
  "created": "<created-value>",
  "details": "<details-value>",
  "category": "<category-value>",
  "message": "<message-value>",
  "user": "<user-value>",
  "taskId": "<taskId-value>"
}
```


 **Note** Null values are ignored.

HTTP response code

Your webhook URL should return a response with HTTP status code 204.


## Using R-Cloud reports

R-Cloud reports provide you with a visual presentation of data protection environment resources within the currently selected protection set. This comprehensive and precise presentation allows you to have an optimum view for analyzing data so that you can make the best decisions when it comes to protecting your data. Report data can be presented as a table or as a chart.

 **Important** Reports reflect the state of your data protection environment with an up to 60-minute latency period.


After you get familiar with the reports as described in [“Getting started with reporting” below](#), you can continue as follows:

- View reports. For details, see [“Viewing reports” on page 178](#).
- Generate reports. For details, see [“Generating reports” on page 179](#).
- Schedule reports. For details, see [“Scheduling reports” on page 179](#).

 **Note** When scheduling the reports, you can also choose to send them by email.

- Export and import reports. For details, see [“Exporting and importing reports” on page 180](#).

Accessing the Reports panel


To access the Reports panel, in the navigation pane, click  **Reports**.

## Getting started with reporting

You can take advantage of predefined reports or create additional reports to better understand your data protection environment, identify potential problems, and improve performance.




For a list of predefined reports, see [“Predefined reports” on the next page](#). For instructions on how to create reports, see [“Creating reports” on page 177](#).

## Predefined reports

Predefined reports, represented by the  icon, provide you with information on the key aspects of your data protection environment, such as the size of disks and the total size of protected instance data. These reports cannot be edited or deleted.

Name	Description
backup-tasks-for-last-24-hours	List of backup tasks for the last 24 hours.
potential-consumption <sup>a</sup>	Collection of potential consumption within the protection set.
protected-data-on-targets-per-vm	Amount of protected data on targets for each protected instance.
protected-vm-disk-capacity-per-policy	Amount of protected instance disk capacity for each policy.
total-vm-disk-capacity-trend	Total amount of instance disk capacity through time.
transferred-data-per-vm-for-previous-month	Amount of transferred data for each protected instance (per backup tier) for the previous month.
unprotected-vms	List of unprotected instances.
vm-compliance-status	List of instances, their compliance statuses, assigned policies, and the corresponding policy tiers.
vm-protected-data-on-targets-per-policy	Amount of protected instance data on targets for each policy.
vm-protected-data-on-targets-per-storage-class	Amount of protected instance data on targets for each storage class.
vm-total-protected-data-on-targets-trend	Total amount of protected instance data on targets through time.

<sup>a</sup> The Google Cloud application data is not included in the collection of potential consumption.

 **Tip** To minimize the Detail view, click  **Minimize** or press the Spacebar. To return the Detail view to its original size, click  **Maximize** or






press the Spacebar.

## Creating reports

If none of the predefined reports meets your reporting requirements, you can create a new report and tailor it to your needs.

Depending on whether you want to create a new report from scratch or edit an existing report and save it as a new report, do the following:

I want to...	Procedure
Create a new report from scratch.	<ol style="list-style-type: none"> <li>1. Click  <b>New</b>. The New Report dialog box opens.</li> <li>2. Enter a report name and, optionally, its description.</li> <li>3. Select the type of report (a table or a chart).</li> <li>4. Select the aggregation value that you want to use to perform a calculation on a set of collected data.</li> <li>5. Specify the time range for the report. The Time Range drop-down menu allows you to: <ul style="list-style-type: none"> <li>• Select one of the predefined time ranges.</li> <li>• Define a custom from-to time range. Click <b>Custom</b> to define the custom time range using a date and time picker.</li> </ul> </li> <li>6. Distribute the report tags for the collected data that you want to include in your report between x-axis and y-axis to determine how the collected data will be presented in the report. <div style="border-left: 2px solid #4a7ebb; padding-left: 10px; margin: 10px 0;"> <p> <b>Important</b> When distributing the report tags, keep in mind that some report tags are not compatible with each other and are, therefore, grayed out after you add a specific report tag to the x-axis or y-axis tags.</p> </div> </li> <li>7. Click <b>Save</b>.</li> </ol>
Edit an existing report and save it as a new report.	<ol style="list-style-type: none"> <li>1. From the list of reports, select the one that you want to edit and save as a new report, and then click  <b>Edit</b>. The Preview Report dialog box opens.</li> </ol>



I want to...	Procedure
	<ol style="list-style-type: none"> <li>2. Enter a new name for the report, and then make the required modifications.</li> <li>3. Click <b>Save as</b> to save the edited report as a new report or <b>Save</b> to save the changes to the existing report.</li> </ol>


## Viewing reports

You can view the reports on the current state of your data protection environment or the saved report versions that were generated either manually or automatically.

### Limitation

You cannot preview the potential-consumption predefined report.

I want to...	Procedure
View a report on the current state of my data protection environment.	From the list of reports, select the preferred report, and click  <b>Preview</b> .
View a saved report version.	<ol style="list-style-type: none"> <li>1. From the list of reports, select the preferred report.</li> <li>2. In the Detail view that appears at the bottom of the screen, select the preferred report version, and then click  <b>View</b>.</li> </ol> <p>For instructions on how to generate report versions manually or automatically, see <a href="#">“Generating reports” on the next page</a> or <a href="#">“Scheduling reports” on the next page</a>.</p>

In the dialog box that opens, besides viewing the report data, you can also download and export the report in the PDF, PNG, or CSV format. To do so, click  **Download**, and then select one of the available formats.

## Generating reports


When you generate a report, you save a copy of the current version of the selected report (a report version) for future reference.


### Consideration



In a protection set with a large number of sources, generating the potential-consumption predefined report may take a while.

### Procedure

1. From the list of reports, select the one that you want to generate.


 **Note** If none of the available reports meets your reporting requirements, you can create a new report. For details, see [“Creating reports” on page 177](#).

2. In the Detail view that appears at the bottom of the screen, click  **Generate**. The Generate Report Version dialog box opens.
3. *Optional.* Enter a description for the report version.
4. Click **Generate**.

 **Tip** You can save a version of the selected report also by clicking  **Preview** followed by **Generate**.

The generated report version is added to the list of report versions in the Detail view that appears at the bottom of the screen when you select a corresponding report.

You can later do the following:

- View the saved report versions. For details, see [“Viewing reports” on the previous page](#).
- Delete the saved report versions that you do not need anymore. To do so, select the preferred report version, and then click  **Delete**.


## Scheduling reports


You can use scheduling to generate report versions automatically at a particular time each day, week, or month. You can view these report versions in the web browser or schedule them by email.

## Limitation



You cannot use scheduling for the potential-consumption predefined report.

## Procedure



1. From the list of reports, select the one that you want to be generated on a regular basis, and then click  **Scheduler**. The Report Scheduler dialog box opens.

 **Note** If none of the available reports meets your reporting requirements, you can create a new report. For details, see “[Creating reports](#)” on page 177.

2. In the Schedule date field, specify the date and the time of day when you want the report generation to begin.
3. From the Interval drop-down menu, select how often you want the report versions to be generated (daily, weekly, or monthly).
4. Use the **Send** switch if you want to schedule the automatic delivery of the reports to email recipients, and then do the following:
  - a. From the Report format drop-down menu, select a file format for your report (PDF, PNG, or CSV).
  - b. In the Email address field, enter one or more email recipients that should receive the reports. If you are entering more than one email address, make sure to press the Spacebar after entering each one.
5. Click **Schedule**.

 **Tip** The reports that are generated automatically are marked by the  icon in the Scheduled column of the Reports panel.

You can later do the following:

- Edit scheduling options of any of the scheduled reports. To do so, select the report, click  **Scheduler**, make the required modification, and then click **Schedule**.
- Unschedule any of the reports if you do not want them to be generated automatically anymore. To do so, select the report, click  **Scheduler**, and then click **Unschedule**.


## Exporting and importing reports

R-Cloud enables you to share user-created reports among different R-Cloud subscriptions by exporting the reports to a JSON file and then importing the

reports from the JSON file.

## Exporting reports


Procedure


From the list of all reports, select the one that you want to export, and then click  **Export**.

The selected report will be exported to a JSON file and saved to the download location on your system.

## Importing reports

Procedure

1. Click  **Import**. The Import Report dialog box opens.
2. Browse your file system for a report that you want to import.
3. Enter a name for the report and, optionally, its description.

 **Note** If the JSON file name and description are already defined in the file itself, the Name and Description fields will be populated automatically. You can, however, use another name and description.

4. Click **Import**.

A new report will be added to the list of the reports.

## Filtering and sorting data

R-Cloud enables you to filter data in the panels so you can easily find what you need. After you apply any of the filters, only data that matches the filter criteria is displayed and you can easily find what you need. For example, filtering the data in the Instances panel helps you to focus only on the instances that you are interested in.

In addition, to make it easier to work with the tables in the panels that have a large number of columns, you can also sort the data in ascending or descending order.

Depending on whether you want to filter or sort your data, see one of the following sections:


- [“Filtering data in panels” below](#)
- [“Sorting data in panels” on page 187](#)

## Filtering data in panels

### Tips

- You can filter data also by using the Search field on the left side of the panel. Typing text in this field automatically filters and displays only the matching items.
- *Only if filtering entity data.* To find out more information about each of the entity statuses mentioned in this section, see [“Viewing information about entities” on page 153](#).

### Procedure

1. In the selected panel, click  **Filters**. The Filters side pane opens.
2. In the side panel that opens, select your filter criteria.
3. Click **Apply Filters**.

For details about the available filtering options, see one of the following sections:

- [“Filtering options in the SaaS panel” below](#)
- [“Filtering options in the Applications panel” on the next page](#)
- [“Filtering options in the Instances panel” on the next page](#)
- [“Filtering options in the Buckets panel” on page 184](#)
- [“Filtering options in the Policies panel” on page 184](#)
- [“Filtering options in the Targets panel” on page 184](#)
- [“Filtering options in the Tasks panel” on page 185](#)
- [“Filtering options in the Events panel” on page 186](#)
- [“Filtering options in the IAM panel” on page 186](#)

### Filtering options in the SaaS panel

The following filtering options are available:

Filtering option	Filter SaaS applications by one or more...
Source	Sources to which the SaaS applications belong.
Policy	Policies assigned to the SaaS applications.

Filtering option	Filter SaaS applications by one or more...
Compliance	Compliance statuses of the SaaS applications (Success, Failure, and/or Undefined).
Protection	Protection statuses of the SaaS applications (Protected, Unprotected, Protected deleted, and/or Undefined).

## Filtering options in the Applications panel

The following filtering options are available:

Filtering option	Filter Google Cloud applications by one or more...
Source	Sources to which the applications belong.
Type	Application types (SAP HANA and/or GKE ).
Policy	Policies assigned to the applications.
Compliance	Compliance statuses of the applications (Success, Failure, and/or Undefined).
Protection	Protection statuses of the applications (Protected, Unprotected, Protected deleted, and/or Undefined).
Discovery	Discovery statuses of the applications (Success, Failure, Warning and/or Undefined).

## Filtering options in the Instances panel

The following filtering options are available:

Filtering option	Filter instances by one or more...
Source	Sources to which the instances belong.
Source type	Instance source types (Google Cloud and/or AWS).
Policy	Policies assigned to the instances.
Credential group	Credential groups assigned to the instances.
Zone	Zones to which the instances belong.
Compliance	Compliance statuses of the instances (Success, Failure, and/or Undefined).
Protection	Protection statuses of the instances (Protected,

Filtering option	Filter instances by one or more...
	Unprotected, Protected deleted, and/or Undefined).
Discovery	Discovery statuses of the instances (Success, Failure, Warning and/or Undefined).

## Filtering options in the Buckets panel

The following filtering options are available:

Filtering option	Filter buckets by one or more...
Source	Sources to which the buckets belong.
Source type	Bucket source types (Google Cloud and/or AWS).
Policy	Policies assigned to the buckets.
Location	Locations to which the buckets belong.
Compliance	Compliance statuses of the buckets (Success, Failure, and/or Undefined).
Protection	Protection statuses of the buckets (Protected, Unprotected, Protected deleted, and/or Undefined).

## Filtering options in the Policies panel

The following filtering option is available:

Filtering option	Filter policies by one or more...
Compliance	<p>Compliance statuses of the policies:</p> <ul style="list-style-type: none"> <li>• Success: All entities to which the policies are assigned are compliant with the policy settings.</li> <li>• Failure: Not all entities to which the policies are assigned are compliant with the policy settings.</li> <li>• Undefined: The exclude policy is assigned to the entities or the entities do not have a policy assigned.</li> </ul>

## Filtering options in the Targets panel

The following filtering options are available:



Filtering option	Filter targets by one or more...
Storage class	Target storage classes: <ul style="list-style-type: none"> <li>• <i>For Amazon S3 targets:</i> <ul style="list-style-type: none"> <li>◦ S3 Standard</li> <li>◦ S3 Standard-IA</li> <li>◦ S3 One Zone-IA</li> <li>◦ S3 Intelligent-Tiering</li> <li>◦ S3 Glacier Flexible Retrieval</li> <li>◦ S3 Glacier Instant Retrieval</li> <li>◦ S3 Glacier Deep Archive</li> </ul> </li> <li>• <i>For Azure targets:</i> <ul style="list-style-type: none"> <li>◦ Hot</li> <li>◦ Cool</li> </ul> </li> <li>• <i>For Google Cloud targets:</i> <ul style="list-style-type: none"> <li>◦ Standard</li> <li>◦ Nearline</li> <li>◦ Coldline</li> <li>◦ Archive</li> </ul> </li> <li>• <i>For Wasabi S3 compatible targets:</i> <ul style="list-style-type: none"> <li>◦ S3 Standard</li> </ul> </li> <li>• <i>For OVHcloud S3 compatible targets:</i> <ul style="list-style-type: none"> <li>◦ S3 Standard</li> </ul> </li> </ul>
Health	Target health statuses (OK, Warning, Error, and/or Undefined).
Tags	Target tags (Automatic and/or Staging).

## Filtering options in the Tasks panel

The following filtering options are available:

Filtering option	Filter tasks by one or more...
Source	Sources to which the tasks belong.
Username	User names to include only the tasks started by the selected accounts.
Type	Task types.

Filtering option	Filter tasks by one or more...
Status	Task statuses (Ready, Running, Aborting, Aborted, Done, Failed, Done with errors, Done with warnings, and/or Skipped).
Time range	Time ranges: You can select one of the predefined time ranges (Last 1 hour, Last 24 hours, and/or Last week), or use the calendar to select a start date and hour and an end date and hour of the time range for tasks to be displayed.

### Filtering options in the Events panel

The following filtering options are available:

Filtering option	Filter events by one or more...
Source	Sources to which the events belong.
Category	Event categories.
Username	User names to include only the events started by the selected accounts.
Severity	Event severity (Success, Warning, and/or Failed).
Time range	Time ranges: You can select one of the predefined time ranges (Last 1 hour, Last 24 hours, and/or Last week), or use the calendar to select a start date and hour and an end date and hour of the time range for events to be displayed.



### Filtering options in the IAM panel

The following filtering options are available:

Filtering option	Filter accounts by one or more...
Type	Account types (User and/or Service account).
Status	Account statuses (Active and/or Deactivated).

## Sorting data in panels

### Procedure

1. In the selected panel, click the column header of the column that you want to sort. The column is sorted in ascending order, which is indicated by the  icon.
2. Click the column header again to sort the data in descending order, which is indicated by the  icon.

## Performing manual backups

R-Cloud backs up your data automatically after you assign a policy to the selected entities. However, you can also back up your data manually at any time, for example, for testing purposes or if an automatic backup fails.


### Prerequisite



A policy other than the exclude policy must be assigned to the entity.

### Consideration

When the assigned policy uses a backup window, manual backups may prevent the scheduled backup from starting within the defined time frame. This may result in data not being protected until the next backup window or the next manual backup.

### Procedure

1. In the SaaS, Applications, Instances, or Buckets panel, select which entities you want to back up.
2. Click  **Backup** to perform the backup of the selected entities.
3. Click **Yes** to confirm that you want to start the manual backup.

 **Tip** In the navigation pane, click  **Tasks** to check the overall progress of the backup.

## Expiring backups manually

R-Cloud expires backups automatically according to the retention period that is set for the backup data in the policy. However, if there is a restore point that you do not want to use for restoring data anymore, you can at any time expire it manually. You can do this also for restore points whose backup status is Failed or Aborted if you want to free storage space.

A restore point represents data that was backed up at a specified point in time. Your restore point can contain one or more tiers—Backup, Copy, Archive—that can be marked as expired also individually. Keep in mind that the Catalog tier cannot be marked as expired.

Depending on whether the selected restore point belongs to a SaaS application, a Google Cloud application, an instance, or a bucket, it can contain one or more tiers that you can mark as expired:

- For *SaaS applications, Google Kubernetes Engine applications, and instances*: Backup (Target), Backup (Snapshot), Copy, and/or Archive

**ⓘ Important** Only the Backup tier is available for GKE applications not using persistent volumes.

- For *SAP HANA applications*: Full or Incremental

**ⓘ Important** Only Full can be marked as expired if at least one successful full backup has been created after it.

- For *buckets*: Backup, Copy, and/or Archive

You can mark as expired one of the following:

- Entire restore point





Make sure that all tiers are marked for expiration.

- One or more tiers:

Make sure that only the tiers that you want to expire are marked for expiration.

**ⓘ Important** Marking a restore point or its tiers as expired cannot be undone. If you are marking an application restore point as expired, keep in mind that all previous backups are also marked for expiration.

Depending on whether you want to expire backups for a SaaS application, a Google Cloud application, an instance, or a bucket, access one of the following panels:



- **Accessing the SaaS panel**  
To access the SaaS panel, in the navigation pane, click  **SaaS**.
- **Accessing the Applications panel**  
To access the Applications panel, in the navigation pane, click  **Applications**.
- **Accessing the Instances panel**  
To access the Instances panel, in the navigation pane, click  **Instances**.
- **Accessing the Buckets panel**  
To access the Buckets panel, in the navigation pane, click  **Buckets**.

### Limitation

You cannot manually expire tiers on targets with Object Lock (WORM) enabled.

### Procedure

1. In the relevant panel, click the entity for which you want to expire a backup. The Detail view appears at the bottom of the screen.
 

 **Note** The Detail view appears only if you click a backup entity. Selecting the check box before its name does not open the Detail view.
2. In the Detail view, select the restore point that you want to mark as expired.
3. Click  **Expire**.
4. *Only if marking an entity restore point as expired and its backup status is not Failed or Aborted.* Select the tiers that you want to mark as expired:
  - **Backup (Snapshot):** *Available only for SaaS applications, GKE applications using persistent volumes, and instances.*
  - **Backup (Target)**
  - **Copy**
  - **Archive—(daily, weekly, monthly, yearly)**

The tiers that are available for expiration are based on the options that you set in your policy. By selecting all the tiers, you mark the entire restore point as expired.

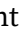
5. Click **Yes** to confirm that you want the selected tiers to be marked as expired.

The next retention maintenance task in R-Cloud removes the corresponding data from the storage locations.


## Exporting the contents of the panel

Data that you can view in a table in any of the panels can be exported to a file in JSON or CSV format.

### Consideration

If you want to export only specific data, click  **Filters**, select your filter criteria based on what kind of data you want to export to a file, and then click **Apply Filters**. You can also use the Search field on the left side of the main panel to filter the data.

### Procedure


1. Navigate to the panel whose data you want to export.
2. Click  **Export**, and then, from the drop-down menu, select one of the following options:

Option	Description
<b>Export to JSON (Current)</b>	Exports the current table page to a JSON file.
<b>Export to JSON (All)</b>	Exports all table data to a JSON file.
<b>Export to CSV (Current)</b>	Exports the current table page to a CSV file.
<b>Export to CSV (All)</b>	Exports all table data to a CSV file.

## Viewing subscription information

This section describes the R-Cloud subscription information that is provided in the web user interface.

Accessing the Subscription Information dialog box

To access the Subscription Information dialog box, click  **<EmailAddress>** in the toolbar, and then select **Subscription Information**.

The following information is displayed in the Subscription Information dialog box for the R-Cloud subscription:

<b>Subscriber</b>	
First name	Information about the person who subscribed to R-Cloud.
Last name	
Company	
<b>Notification email recipients</b>	<p>A list of recipients to whom notifications related to the selected R-Cloud subscription will be sent.</p> <p>If this field is empty, all the important notifications related to the R-Cloud subscription, such as support and upgrade information, are by default sent to all users that are using the service. It is recommended that you verify these email addresses and, if required, update the list of email addresses to which the notifications are sent.</p>
<b>Subscription Details</b>	
Subscription ID	The identifier that is automatically generated by R-Cloud and assigned to your HYCU subscription during registration. The subscription ID is used by R-Cloud when addressing the reported issues.
Subscription plan	The plan that your R-Cloud subscription is using. Subscriptions that are not based on a quote are using the Basic plan (also called the Pay-as-you-go plan). For more information, see <a href="#">“Backup and data retention pricing” on page 16</a> .
Subscribed on	The date of subscribing to R-Cloud.
Version	Current R-Cloud version.
<b>HYCU Account</b>	
HYCU Account ID	The identifier that is automatically generated by R-Cloud and assigned to your HYCU account during registration. The HYCU Account ID is used to sign in to R-Cloud.

Login URL	The login URL for the HYCU account.
Alias	An alias for your HYCU account that you can use to sign in to R-Cloud.

## Customizing your R-Cloud web user interface

The R-Cloud web user interface is designed to be customized to match your needs. When customizing your R-Cloud web user interface, you can do the following:

- Adjust the table density to determine how close or far apart the text in the tables should be.
- Choose to show or hide the separator line between the rows.
- Switch your R-Cloud web user interface to light or dark mode. R-Cloud by default uses the color mode exposed by your browser.

### Accessing the Customization dialog box

To access the Customization dialog box, click  **<EmailAddress>** in the toolbar, and then select **Customization**.

### Procedure

1. In the Customization dialog box, do the following:
  - Under Table density, select **Default density** or **High density** depending on how close or far apart you want the text in the tables to be.
  - Enable the **Row dividers** switch if you want to show the separator line between the rows.
  - Do one of the following:
    - *If you want to use the dark mode:* Enable the **Dark mode** switch.
    - *If you want to use the light mode:* Disable the **Dark mode** switch.
2. Click **Close**.

The changes take place immediately without the need to sign out and sign in again to the R-Cloud web user interface. The preferred customization is remembered for the next time you sign in to R-Cloud.



# Chapter 9

## Customizing R-Cloud

After you subscribe to R-Cloud, you can perform various tasks to customize R-Cloud for your data protection environment.

### Consideration

Keep in mind that the role you have assigned determines what kind of actions you can perform. For details on roles, see [“Managing roles” on page 225](#).

If you have the Administrator role assigned, the scope of tasks you can perform depends on the user interface context you select. You can switch between the following two contexts:


- Subscription

In the subscription context, only the IAM panel and the dashboard are active. Use this context to perform administration tasks related to your subscription, such as adding identity providers, adding or removing users, or changing roles.

- Protection set

In the protection set context, you select the scope of data protection by selecting a specific protection set.

### Switching the user interface context

1. On the toolbar, click  next to the name of the selected protection set or subscription.
2. From the drop-down menu, select the context.

The R-Cloud web user interface switches the context. The context that you select is remembered for the next time you sign in.

### Tasks

Task	Instructions
Manage R-Cloud protection sets.	<a href="#">“Managing protection sets” on the next page</a>

Task	Instructions
Add cloud accounts.	<a href="#">“Adding cloud accounts” on page 199</a>
Add, edit, or remove sources.	<a href="#">“Managing sources” on page 206</a>
Configure service discovery and explore your data protection environment.	<a href="#">“Discovering SaaS services” on page 213</a>
Manage identity providers, add or remove users, and add or remove roles.	<a href="#">“Managing identity and access” on page 221</a>
Hide instances from R-Cloud.	<a href="#">“Excluding instances from synchronization by tagging the instance in AWS or Google Cloud” on page 228</a>
Stop protecting individual sources.	<a href="#">“Stopping protection for individual sources” on page 227</a>

## Managing protection sets

By default, a predefined protection set is created automatically (named default-protection-set) and if your business needs require no additional protection sets, all AWS accounts, Google Cloud projects, and/or R-Cloud modules that you add to R-Cloud as sources are added to this default protection set. However, you can at any time create additional protection sets and distribute your sources among them. By doing so, you define different scopes of data protection that best suit your needs and gain an at-a-glance view of the sources in each protection set.

### Prerequisite


You must have the Administrator role assigned at the subscription level.

You can perform the following tasks related to protection sets:

Task	Instructions
Create a protection set and include preferred sources in it.	<a href="#">“Creating protection sets” on the next page</a>
Edit an existing protection set.	<a href="#">“Editing protection sets” on page 196</a>

Task	Instructions
Add a Google Cloud project to a protection set by using a label.	“Adding Google Cloud projects to a protection set by using a label” on page 197
Remove a Google Cloud project from a protection set by using a label.	“Removing Google Cloud projects from a protection set by using a label” on page 197
Delete a protection set that you no longer need.	“Deleting protection sets” on page 198

### Accessing the Protection Sets dialog box

To access the Protection Sets dialog box, click  **Administration**, and then select **Protection Sets**.


## Creating protection sets


You can create additional protection sets that allow you to have different data protection setup for different groups of sources.

### Considerations

- If you move a source to a different protection set, consider the following:
  - Policies will be automatically unassigned from the entities in the source.
  - If you move an AWS account or a Google Cloud project, the credential groups that were manually assigned to the instances in the account or the project will be automatically unassigned from those instances.
- An AWS account cannot be moved to a different protection set if its default AWS IAM role is assigned to an existing target. To be able to move the AWS account to a different protection set, you must either delete the target or make sure that the target uses a different cloud account.

### Procedure

1. In the Protection Sets dialog box, click  **New**.
2. Enter a name for your protection set and, optionally, its description.
3. From the list of available sources, select one or more sources that you want to include in the protection set.

 **Tip** You can search for a source by entering its name in the search field and then pressing **Enter**. By selecting the Source check box, you select all sources at once.

#### 4. Click **Save**.

The protection set is created and added to the list of protection sets.

## Editing protection sets

You can change the name of a protection set, add sources to the protection set, or remove sources from the protection set.

When you remove a source from the protection set other than the default one, the source is automatically moved to the default protection set. If you want to completely remove the source from R-Cloud and stop protecting its resources, you must remove the source from the default protection set.


As an alternative to adding or removing sources by using the R-Cloud web user interface, you can also add or remove Google Cloud projects from protection sets by using a label. For details, see the following sections:

- [“Adding Google Cloud projects to a protection set by using a label” on the next page](#)
- [“Removing Google Cloud projects from a protection set by using a label” on the next page](#)

### Considerations

- If you move a source to a different protection set, consider the following:
  - Policies will be automatically unassigned from the entities in the source.
  - If you move an AWS account or a Google Cloud project, the credential groups that were manually assigned to the instances in the account or the project will be automatically unassigned from those instances.
- An AWS account cannot be moved to a different protection set if its default AWS IAM role is assigned to an existing target. To be able to move the AWS account to a different protection set, you must either delete the target or make sure that the target uses a different cloud account.

### Procedure

1. In the Protection Sets dialog box, from the list of protection sets, select the one that you want to edit, and then click  **Edit**.

2. Edit the name of the protection set and its description.
3. *Only if you want to add sources to the protection set.* From the list of sources, select one or more sources that you want to add to the protection set. The sources that already belong to the protection set are preselected.
4. *Only if you want to remove sources from the protection set.* From the list of sources, deselect one or more sources that you want to remove from the protection set. The sources that belong to the protection set are preselected.
5. Click **Save**.
6. *Only if you want to add or remove sources from the protection set.* Click **Yes** to confirm that you want to add or remove the sources from the protection set.

## Adding Google Cloud projects to a protection set by using a label

As an alternative to adding a project to a protection set by using the R-Cloud web user interface, you can also add a project to a protection set by attaching the `hycu-protection-set` label to the project in Google Cloud.

### Prerequisite

The protection set to which you want to add the project must be created in R-Cloud.

### Procedure

In Google Cloud, attach the label to the project as the following key/value pair:

Key	Value
<code>hycu-protection-set</code>	<code>&lt;ProtectionSetName&gt;</code> In this case, <code>&lt;ProtectionSetName&gt;</code> is the name of the protection set to which you want to add the project.

For detailed instructions on how to create and manage labels, see Google Cloud documentation.

## Removing Google Cloud projects from a protection set by using a label

As an alternative to removing a project from a protection set by using the R-Cloud web user interface, you can also remove a project from a protection set by attaching the `hycu-protection-set` label to the project in Google Cloud.

### Consideration

If after excluding a project from a protection set and R-Cloud by using the `hycu-project-exclude` label, you need to add the same project to R-Cloud again, contact [HYCU Support](#).

### Procedure

In Google Cloud, add the label to the project as the following key/value pair:

Key	Value
<code>hycu-project-exclude</code>	<code>true</code>

After you add the label to the project, it is no longer included in the protection set and R-Cloud no longer retrieves its information from Google Cloud.

For detailed instructions on how to create and manage labels, see [Google Cloud documentation](#).

## Deleting protection sets

You can at any time delete protection sets that you no longer need.


### Prerequisites

- The protection set that you want to delete must be empty with no included sources.
- The current data protection scope must be set to a protection set other than the protection set that you want to delete.

### Consideration

The default protection set created by R-Cloud cannot be deleted.

### Procedure

1. In the Protection Sets dialog box, from the list of protection sets, select the one that you want to delete from R-Cloud, and then click  **Delete**.
2. Click **Delete** to confirm that you want to delete the selected protection set.

# Adding cloud accounts

If you want to use a specific cloud account for performing operations related to data protection, you must add it to R-Cloud.

## Prerequisite

You must have the Administrator role assigned.

## Consideration


Depending on whether you are a subscription administrator or a protection set administrator, consider the following:

- *If you are a subscription administrator:* You can view and manage all cloud accounts in all protection sets in your data protection environment.
- *If you are a protection set administrator:* You can view and manage only the cloud accounts in the currently selected protection set.

Depending on what type of cloud account you want to add to R-Cloud, see one of the following:

Cloud account	Use this cloud account to...	Instructions
AWS IAM role	Perform all operations on an Amazon S3 target.	<a href="#">“Adding AWS IAM roles” on the next page</a>
Azure service principal	Perform all operations on an Azure target.	<a href="#">“Adding Azure service principals” on page 201</a>
Google Cloud service account	Perform all operations on a Google Cloud target or to authenticate with an R-Cloud module.	<a href="#">“Adding Google Cloud service accounts” on page 202</a>
S3 compatible account	Perform all operations on an S3 compatible target.	<a href="#">“Adding S3 compatible accounts” on page 203</a>

### Accessing the Cloud Accounts dialog box

To access the Cloud Accounts dialog box, click  **Administration**, and then select **Cloud Accounts**.

## Adding AWS IAM roles

To allow a specific AWS IAM role to perform all operations on an Amazon S3 target, you must add the role to R-Cloud as a cloud account (as an alternative to creating an AWS IAM role as part of adding the AWS account where the target resides to R-Cloud), and then specify it when setting up the target.

For details on how to specify an AWS IAM role when setting up an Amazon S3 target, see [“Setting up an Amazon S3 target” on page 29](#).

### Prerequisites

- An AWS IAM role must be created in AWS. The role must have the policies with the permissions for the S3 services attached. If you plan to set up a directory bucket as an Amazon S3 target, the permissions for the S3 Express service must also be included.

Service	Permissions
S3	ListBucket ListBucketVersions GetBucketLocation GetBucketObjectLockConfiguration GetBucketPublicAccessBlock GetBucketTagging GetBucketVersioning GetEncryptionConfiguration GetLifecycleConfiguration GetObject GetObjectTagging DeleteObject DeleteObjectVersion PutBucketTagging PutObjectTagging PutObject ListAllMyBuckets  <i>For targets that have Object Lock (WORM) enabled, the following additional permissions are required:</i> PutObjectRetention PutObjectLegalHold




S3 Express	CreateSession ListAllMyDirectoryBuckets
------------	--

For details on policies and permissions in IAM, see AWS documentation.



- Your AWS IAM role must have a trust relationship established with R-Cloud that includes the following:
  - The AWS principal: `arn:aws:iam::<HYCUAWSAccountID>:root`. To get your HYCU AWS account ID, contact [HYCU Support](#).
  - The `sts:AssumeRole` action.

For details on how to establish a trust relationship, see AWS documentation.

### Procedure

1. In the Cloud Accounts dialog box, click  **New**.
2. Select **Add AWS IAM Role**, and then click **Next**.
3. In the Name field, enter a name for your IAM role.
4. From the Protection Set drop-down menu, select the protection set to which you want to add your IAM role.
5. In the S3 ARN field, enter the Amazon Resource Name (ARN) of your IAM role.
6. In the External ID field, enter the external ID of your IAM role trust relationship.
7. Click **Save**.

The IAM role is added to the list of cloud accounts in R-Cloud.

You can at any time edit any of the IAM roles (click  **Edit** and make the required modifications) or delete the ones that you do not need anymore (click  **Delete**). Keep in mind that deleting the IAM role from R-Cloud does not remove it from AWS.

## Adding Azure service principals

To allow a specific service principal to perform all operations on an Azure target, you must first add the service principal to R-Cloud, and then specify it when setting up the target.

For details on how to specify an Azure service principal when setting up an Azure target, see [“Setting up an Azure target” on page 31](#).

## Prerequisite

A service principal must be created in Azure and it must have the following roles assigned:

- Storage Blob Data Contributor
- Storage Blob Data Owner
- A custom role that contains the following permissions:
  - Microsoft.Storage/storageAccounts/blobServices/containers/blobs/tags/read
  - Microsoft.Storage/storageAccounts/blobServices/containers/blobs/tags/write
  - Microsoft.Storage/storageAccounts/blobServices/read
  - Microsoft.Storage/storageAccounts/managementPolicies/read
  - Microsoft.Storage/storageAccounts/read
  - Microsoft.Storage/storageAccounts/write

## Procedure

1. In the Cloud Accounts dialog box, click **+** **New**.
2. Select **Add Azure Service Principal**, and then click **Next**.
3. In the Name field, enter a name for your service principal.
4. From the Protection Set drop-down menu, select the protection set to which you want to add your service principal.
5. In the Tenant ID field, enter your tenant ID.
6. In the Application ID field, enter the ID of the service principal.
7. In the Client Secret field, enter the client secret value.
8. Click **Save**.

The service principal is added to the list of cloud accounts in R-Cloud.

You can at any time edit any of the service principals (click **✎** **Edit** and make the required modifications) or delete the ones that you do not need anymore (click **🗑** **Delete**). Keep in mind that deleting the service principal from R-Cloud does not remove it from Azure.

## Adding Google Cloud service accounts


To allow a specific service account to perform all operations on a Google Cloud target or to authenticate with an R-Cloud module, you must first add the service account to R-Cloud, and then specify it when setting up the target or adding an R-Cloud module to R-Cloud.

For details on how to specify a service account when setting up a Google Cloud target, see [“Setting up a Google Cloud target” on page 33](#). For details on how to specify a service account when adding an R-Cloud module to R-Cloud as a source, see [“Adding R-Cloud modules” on page 210](#).

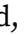

### Prerequisite

A service account must be configured in Google Cloud and you must have access to a valid JSON file that stores the service account information, including its private key.

### Procedure

1. In the Cloud Accounts dialog box, click  **New**.
2. Select **Add Google Cloud Service Account**, and then click **Next**.
3. In the Name field, enter a name for the service account that you want to add.
4. From the Protection Set drop-down menu, select the protection set to which you want to add your service account.
5. Click **Browse**.
6. Select the JSON file with the service account information, and then click **Open**.
7. Review the service account information, and then click **Upload**.

The service account is added to the list of cloud accounts in R-Cloud.

You can at any time edit any of the existing service accounts (click  **Edit** and make the required modifications) or delete the ones that you do not need anymore (click  **Delete**). Keep in mind that deleting the service account from R-Cloud does not remove it from Google Cloud.

## Adding S3 compatible accounts


To allow a specific account to perform all operations on an S3 compatible target, you must first add the account to R-Cloud, and then specify it when setting up the target.


For details on how to specify an S3 compatible account when setting up an S3 compatible target, see [“Setting up an S3 compatible target” on page 34](#).

## Prerequisite

An account must be configured in an S3 compatible provider environment, and it must have the `s3:ListAllMyBuckets` permission granted.



## Procedure

1. In the Cloud Accounts dialog box, click  **New**.
2. Select **Add S3 Compatible Account**, and then click **Next**.
3. In the Name field, enter a name for your S3 compatible account.
4. From the Protection Set drop-down menu, select the protection set to which you want to add your S3 compatible account.
5. From the S3 Provider drop-down menu, select your S3 compatible service provider.
6. In the Access Key ID field, enter the access key ID of the S3 compatible account.
7. In the Secret Access Key field, enter the secret access key of the S3 compatible account.

 **Note** The access key ID and the secret access key are used to authenticate S3 REST API service calls.

8. Click **Save**.

The S3 compatible account is added to the list of cloud accounts in R-Cloud.


You can at any time edit any of the S3 compatible accounts (click  **Edit** and make the required modifications) or delete the ones that you do not need anymore (click  **Delete**). Keep in mind that deleting the S3 compatible account from R-Cloud does not remove it from the S3 compatible provider environment.

## Navigating the HYCU Marketplace



The HYCU Marketplace enables you to quickly and easily find and configure R-Cloud solutions to help you meet your SaaS application business needs. It includes products for:

- SaaS applications that can be protected with R-Cloud using R-Cloud modules
- Solutions that can be protected with R-Cloud (for example, Amazon S3 and Google Kubernetes Engine)

- Solutions that can be protected with other HYCU products (for example, Microsoft 365 and Azure Gov Cloud)


 **Note** Some of the HYCU Marketplace entries may direct you to other HYCU data protection solutions. For details on these solutions, see the related HYCU documentation.

### Accessing the Marketplace panel

To access the Marketplace panel, in the navigation pane, click  **Marketplace**. Alternatively, in the toolbar, click .

### Searching the HYCU Marketplace

To find R-Cloud modules for your SaaS applications, in the Search field type the name or part of the name and the list is filtered as you type.

To further narrow your search, you can use the platform or category filters to filter the HYCU Marketplace by platform (for example, AWS or Google Cloud) or by category (for example, DevOps or Sales & Marketing). To clear your selection, click  next to it.

If R-Cloud does not provide data protection for the SaaS application that you want to protect, you can request application support by clicking **Request Application Support**.

### Navigating the product details page

To open the details page, select the product, and then click **Show more**.

For each product, you can perform the following actions:

I want to ...	Actions
Get more information about the product.	Each product details page contains a more detailed description. You can find links to additional resources such as the <i>R-Cloud Module Guide</i> , product sheets, or tutorials. Click the link to open a resource.
Add an R-Cloud module, an AWS account, or a Google project as a source to R-Cloud.	<ol style="list-style-type: none"> <li>1. Click <b>Configure</b> to open the Sources &gt; New dialog box.</li> <li>2. Enter the required information to finish the configuration. For details on how to add an R-Cloud module, an AWS account, or a Google Cloud project as a source to R-Cloud, see <a href="#">“Managing sources” on the next page</a>.</li> </ol>

To return to the Marketplace panel, click **Back to Marketplace**.

## Managing sources

R-Cloud provides data protection for the following sources:

- AWS accounts
- Google Cloud projects
- R-Cloud modules

### Prerequisite

You must have the Administrator role assigned.

### Consideration

Depending on whether you are a subscription administrator or a protection set administrator, consider the following:


- *If you are a subscription administrator:* You can view and manage all sources in all protection sets in your data protection environment.
- *If you are a protection set administrator:* You can view and manage only the sources in the currently selected protection set.

You can add, edit, or remove the sources directly from R-Cloud. When adding sources, R-Cloud may request you to grant the required permissions or roles to R-Cloud.

For details on how to manage sources, see the following topics:

- [“Managing AWS accounts” below](#)
- [“Managing Google Cloud projects” on page 208](#)
- [“Managing R-Cloud modules” on page 210](#)

### Accessing the Sources dialog box

To access the Sources dialog box from the toolbar, click  **Administration**, and then select **Sources**.

To access the Sources dialog box from the R-Graph, in an empty R-Graph with no sources configured, click **Set Up Source**.

## Managing AWS accounts

You can perform the following tasks related to AWS accounts:

Task	Instructions
Add an AWS account.	<a href="#">“Adding AWS accounts” below</a>
Edit an existing AWS account.	<a href="#">“Editing AWS accounts” on the next page</a>
Remove an AWS account that you no longer need.	<a href="#">“Removing AWS accounts” on the next page</a>

## Adding AWS accounts

### Prerequisite

*Only if you plan to add your AWS account to a protection set other than the default one.* The protection set must be created. For instructions, see [“Creating protection sets” on page 195](#).

### Consideration

You can also add sources directly from the Marketplace panel by clicking **Configure** in the product details page, bypassing the initial Sources dialog box.


### Procedure

1. *Not applicable if you are redirected from the Marketplace panel.* In the Sources dialog box, click the **AWS** tab, and then click **New**.
2. From the Protection Set drop-down menu, select the protection set to which you want to add the AWS account. By default, the AWS account is added to the currently selected protection set.
3. Enter the account ID, and, optionally, a display name for the AWS account, and then click **Add**.
4. Click **Create IAM Role**. The AWS Management Console opens.

**ⓘ Important** You must be signed in to AWS Management Console with the account that you are adding to R-Cloud. If you are already signed in to AWS Management Console with a different account when you create the IAM roles, the creation fails.


5. In the AWS Management Console, on the Quick create stack page, confirm the capabilities required by R-Cloud by clicking **I acknowledge that AWS CloudFormation might create IAM resources with custom names**, and then click **Create stack**.
6. Return to the R-Cloud web user interface, and then click **Save**.

The AWS account is added to the list of sources.

 **Note** If you do not complete the IAM role creation step in the AWS Management Console or if you enter an incorrect account ID, the source adding procedure is suspended after a timeout and the status of the source is Preparing. If this happens, remove the source, and then add it to R-Cloud again. For instructions, see [“Removing AWS accounts”](#) below.

## Editing AWS accounts

Procedure

1. In the Sources dialog box, from the list of account IDs, select the AWS account that you want to edit, and then click  **Edit**.
2. Edit the selected AWS account as required.
3. Click **Save**.


## Removing AWS accounts

You can at any time remove AWS accounts that you no longer need.

Consideration

Removing the AWS account from R-Cloud does not delete any IAM resources that were created in the AWS account.

Procedure

1. In the Sources dialog box, from the list of account IDs, select the AWS account that you want to remove from R-Cloud, and then click  **Delete**.
2. Click **Delete** to confirm that you want to remove the selected AWS account.

## Managing Google Cloud projects

You can perform the following tasks related to Google Cloud projects:

Task	Instructions
Add a Google Cloud project and enable the HYCU Managed Service Account for it.	<a href="#">“Adding Google Cloud projects”</a> on the next page
Remove a Google Cloud project that you no longer need.	<a href="#">“Removing Google Cloud projects”</a> on page 210



## Enabling the HYCU Managed Service Account

The HYCU Managed Service Account (HMSA) is a special type of account that is designed specifically for R-Cloud to run data protection operations. It provides business continuity of your data protection environment by enforcing a single service account that cannot be deleted accidentally, and at the same time it also delivers enhanced security by uniquely identifying the service and using key rotation to limit risks associated with potential service account key leaks.

You enable the HMSA for a project by following the HYCU Managed Service Account configuration wizard when adding a Google Cloud project as a source.

## Adding Google Cloud projects



### Prerequisite

*Only if you plan to add your Google Cloud project to a protection set other than the default one.* The protection set must be created. For instructions, see [“Creating protection sets” on page 195](#).

### Consideration

You can also add sources directly from the Marketplace panel by clicking **Configure** in the product details page, bypassing the initial Sources dialog box.


### Procedure

1. *Not applicable if you are redirected from the Marketplace panel.* In the Sources dialog box, click the **Google Cloud** tab, and then click  **New**.
2. From the Protection Set drop-down menu, select the protection set to which you want to add the Google Cloud project. By default, the Google Cloud project is added to the currently selected protection set.
3. Enter the project ID, and then click **Add**. The HMSA email is displayed.
4. Click  **Copy to Clipboard** to copy the HMSA email to the clipboard. You need the email address to assign permissions to the HMSA.
5. Click **Grant Access** to open the HYCU Managed Service Account configuration wizard.

The HYCU Managed Service Account configuration wizard guides you through all the required steps of enabling the HMSA for the project.

6. After you successfully complete all the steps, return to the R-Cloud web user interface, and then click **Save**.

The Google Cloud project is added to the list of sources.

 **Note** If you do not complete the steps in the HYCU Managed Service Account configuration wizard or if you enter an incorrect project ID, the source adding procedure is suspended after a timeout and the status of the source is Preparing. If this happens, remove the source, and then add it to R-Cloud again. For instructions, see [“Removing Google Cloud projects”](#) below.


## Removing Google Cloud projects

You can at any time remove Google Cloud projects that you no longer need.

### Consideration

Removing the Google Cloud project from R-Cloud does not delete any IAM resources that were created in the Google Cloud project.

### Procedure

1. In the Sources dialog box, on the Google Cloud tab, select the Google Cloud project that you want to remove from R-Cloud, and then click  **Delete**.
2. Click **Delete** to confirm that you want to remove the selected project.

## Managing R-Cloud modules

You can perform the following tasks related to R-Cloud modules:

Task	Instructions
Add an R-Cloud module as a source.	<a href="#">“Adding R-Cloud modules”</a> below
Edit an existing R-Cloud module.	<a href="#">“Editing R-Cloud modules”</a> on page 212
Remove an R-Cloud module that you no longer need.	<a href="#">“Removing R-Cloud modules”</a> on page 212

## Adding R-Cloud modules

To be able to protect SaaS application data, you must add an R-Cloud module to R-Cloud as a source.

If the R-Cloud module supports storing data on a staging target, as part of adding an R-Cloud module, you also add an Amazon S3 bucket, a Google Cloud bucket, or an S3 compatible bucket to R-Cloud as a staging target. The staging target is used either to temporarily store SaaS application data before it is

moved to the target that you define in the R-Cloud policy, or to store SaaS application data as a snapshot. For information on whether your R-Cloud module supports staging targets, see the [R-Cloud Module Guides](#).

### Prerequisites

- *Only if you plan to add your R-Cloud module to a protection set other than the default one.* The protection set must be created. For instructions, see [“Creating protection sets” on page 195](#).
- *Only if your R-Cloud module supports storing data on a staging target.* The staging target that you plan to add to R-Cloud must be created in Amazon S3 or Google Cloud Storage.

### Limitations

*Only if your R-Cloud module supports storing data on a staging target.* When adding a staging target to R-Cloud, the following limitations apply:

- Targets that are specified in any of the R-Cloud policies cannot be used as staging targets.
- Targets with Object Lock (WORM) enabled cannot be used as staging targets.
- Automatically created staging targets are created only in Google Cloud Storage.
- The staging target that you add to R-Cloud when adding an R-Cloud module, and the target that is defined in the policy that is assigned to the related SaaS application must reside on the same cloud platform.

### Considerations

- *Only if your R-Cloud module supports storing data on a staging target.* When adding a staging target to R-Cloud, consider the following:
  - The staging target must be dedicated exclusively to SaaS application backups.
  - Data belonging to different R-Cloud modules cannot be stored on the same staging target (one staging target per R-Cloud module).
  - If you use an automatically created staging target, the HMSA must be configured to perform all operations on the target specified in the policy that is assigned to the related SaaS application. Alternatively, the same cloud account must be configured to perform all operations on both targets (the staging target and the target specified in the policy that is assigned to the related SaaS application).

- You can also add sources directly from the Marketplace panel by clicking **Configure** in the product details page, bypassing the initial Sources dialog box.

### Procedure

1. *Not applicable if you are redirected from the Marketplace panel.* In the Sources dialog box, click the **SaaS** tab, and then click **+ New**.
2. From the R-Cloud Module drop-down menu, select the R-Cloud module that you want to add to R-Cloud.
3. In the Display Name field, enter a display name for the R-Cloud module.
4. From the Protection Set drop-down menu, select the protection set to which you want to add the R-Cloud module. By default, the R-Cloud module is added to the currently selected protection set.
5. *Only if your R-Cloud module supports storing data on a staging target.* From the Staging Target drop-down menu, select one of the following for storing data:
  - *Only if your R-Cloud module supports automatically created targets.*  
**Automatically selected**  
 If you select this option, R-Cloud automatically creates a staging target and uses it to temporarily store the data.
  - Any available staging target of your choice
6. Provide the required authentication information, such as the organization name, the user name, API tokens, the preferred service account, and so on.
7. Click **Save**.

## Editing R-Cloud modules


### Procedure

1. In the Sources dialog box, from the list of R-Cloud modules, select the one that you want to edit, and then click **✎ Edit**.
2. Edit the selected R-Cloud module as required.
3. Click **Save**.


## Removing R-Cloud modules

You can at any time remove R-Cloud modules that you no longer need.

## Prerequisites

- All the policies must be unassigned from all the SaaS applications in the R-Cloud module. To unassign the policies from the SaaS applications, in the SaaS panel, select the applications, and then click  **Set Policy**. Click **Unassign**, and then click **Yes** to confirm that you want to unassign the policies from the selected SaaS applications.
- No restore points must be present for any of the SaaS applications in the R-Cloud module. If any of the SaaS applications in the R-Cloud module still have valid restore points, you must expire them manually and wait for the next retention maintenance task to finish before removing the R-Cloud module. For details on how to expire restore points, see [“Expiring backups manually” on page 188](#).
- No tasks with the Ready status or a progress bar indicating the Running status must be present for the R-Cloud module.

## Procedure

1. In the Sources dialog box, from the list of R-Cloud module, select the one that you want to remove from R-Cloud, and then click  **Delete**.
2. Click **Delete** to confirm that you want to remove the selected R-Cloud module.

# Discovering SaaS services


As part of SaaS native data protection, R-Cloud allows you to discover all SaaS services to which you are subscribed and map them by using R-Graph. This helps you to understand the scope of protected and unprotected SaaS data, and to use R-Cloud to establish the data protection environment according to your business needs.

SaaS service discovery starts automatically after you allow R-Cloud to access the required SaaS service information by configuring an identity provider for SaaS service discovery in R-Cloud. After your SaaS services are discovered, you can explore R-Graph, a visual representation of your SaaS data protection environment, which enables you to quickly gain insight into the status of your SaaS application data protection. For details, see [“Exploring R-Graph” on page 215](#).


## Prerequisites

- *For Microsoft Entra ID:* Your service principal must have the `Microsoft Graph/Application.Read.All` application resource permissions.
- *For Okta:* You must use a custom role with permissions to view applications and their details, and a custom resource set that is constrained to all applications.

### Accessing the Discover Services dialog box



To access the Discover Services dialog box, click  **Administration**, and then select **Discover Services**.

## Procedure

1. In the Discovery Services dialog box, click  **New**.
2. From the Identity provider drop-down menu, select your identity provider, and then follow the instructions:

Identity provider	Instructions
<b>Microsoft Entra ID</b>	<ol style="list-style-type: none"> <li>In the Display name field, enter a display name for Microsoft Entra ID.</li> <li>In the Client ID field, enter the application ID that is generated by Microsoft Entra ID.</li> <li>In the Client secret field, enter the application secret that is associated with the client ID and generated by Microsoft Entra ID.</li> <li>In the Tenant ID, enter the identifier for the tenant that is based on the tenant name and can be found in the Microsoft Entra ID properties.</li> </ol>
<b>Okta</b>	<ol style="list-style-type: none"> <li>In the Display name field, enter a display name for Okta.</li> <li>In the Okta URL field, enter your Okta domain that you can find in the global header located in the upper-right corner of the Admin Dashboard.</li> <li>In the API token field, enter the token that is used to authenticate requests to Okta APIs.</li> </ol>

3. Click **Save**. The identity provider is added to the list of identity providers.

You can later edit any of the existing identity providers (click  **Edit** and make the required modifications) or delete the ones that you do not need anymore (click  **Delete**).

## Exploring R-Graph

R-Graph is a visual representation of your data protection environment, displaying the topology, data protection and compliance statuses of different data sources—cloud workloads, applications and databases, and SaaS applications.

### Prerequisite


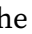


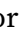
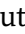

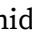
An identity provider that R-Cloud uses to perform SaaS service discovery must be added to R-Cloud. For details, see [“Discovering SaaS services” on page 213](#).

### Consideration

If no identity provider is configured, only the protection sets and sources that you already configured in R-Cloud are shown. If no sources are configured, an empty R-Graph is shown.

### Navigating R-Graph

Use the following actions to navigate R-Graph:

- Switch between layouts: In the top right corner of the graph, click  or  to switch between the tree layout and the force-directed layout of the graph.
- Expand R-Graph over the entire display pane: Click the  arrow in the bottom right corner of the graph.
- Reset the graph to its initial position and scale to fit the screen: Click  in the bottom right corner of the graph.
- Zoom in and zoom out: Scroll the mouse wheel to zoom in or out, click  or  in the bottom right corner of the graph to zoom in or out, or double-click an empty area of the graph to zoom in.
- Display hidden service nodes: By default, hidden service nodes are not displayed. Click  in the bottom right corner of the graph. The display status icon changes to  and hidden nodes are displayed shaded. To hide the hidden nodes, click the display status icon again.

- **Move around:** When zoomed in, click in the area between nodes and drag the graph to display the part of the graph in which you are interested.

Perform additional actions on nodes:

- **Display additional information about the node:** Pause the pointer on the node.
- **Expand and collapse a node:** In the top right corner of the node, click + to collapse the elements of the node and – to expand the elements of the node.
- **Open HYCU Marketplace to configure an R-Cloud module:** In the Subscription context, pause the pointer on the module availability indicator to expand it, and then click the arrow on the right side.
- **Create a protection set and add a source to it:** Double-click the generic protection set placeholder and click **Set up Source**.
- **Show entities of a particular source:** Right-click the source node and select **Filter by Source** or double-click the source node to open the Detail view of a panel and filter entities by this source.
- **Switch to a protection set:** In subscription context, right-click the protection set node and select **Go to Protection Set** or double-click the protection set node.
- **Hide a node:** Right-click the source node and select **Hide Node**. You can hide only the service nodes that have no elements. To make a hidden node visible again, right-click it and select **Unhide Node**.
- **Request an R-Cloud module:** Right-click the service node and select **Request Module**.

## R-Graph elements and structure

R-Graph can be viewed in two layouts: the tree layout or the force-directed layout. Both layouts use the same elements to represent your environment. Each node in the graph represents an element of the data protection environment—protection sets, services, and sources. Connections between nodes represent relationships between these elements—sources are grouped under services and services are grouped under protection sets. Intuitive icons help you to quickly glance the status of data protection for each node and overlay icons show their compliance status.

The following layers of objects represent the data protection environment:



- Root element. HYCU subscription (visible only in the Subscription context).
- Protection sets. The default protection set, protection sets that you create, and generic protection set placeholders that group newly discovered services.
- SaaS services. Various types of services or platforms that can have one or more sources. Examples of SaaS service types are: Amazon EC2, Salesforce, Atlassian Jira, and Google Cloud Storage.
- Sources. AWS accounts, Google Cloud projects, or R-Cloud modules.

The following figure shows an example of a zoomed-in part of R-Graph in the force-directed layout with three layers of objects (protections sets ①, services ②, and sources ③), their protection status icons ④, an indicator that there is an R-Cloud module available in HYCU Marketplace ⑤, and a tooltip ⑥ with protection details:

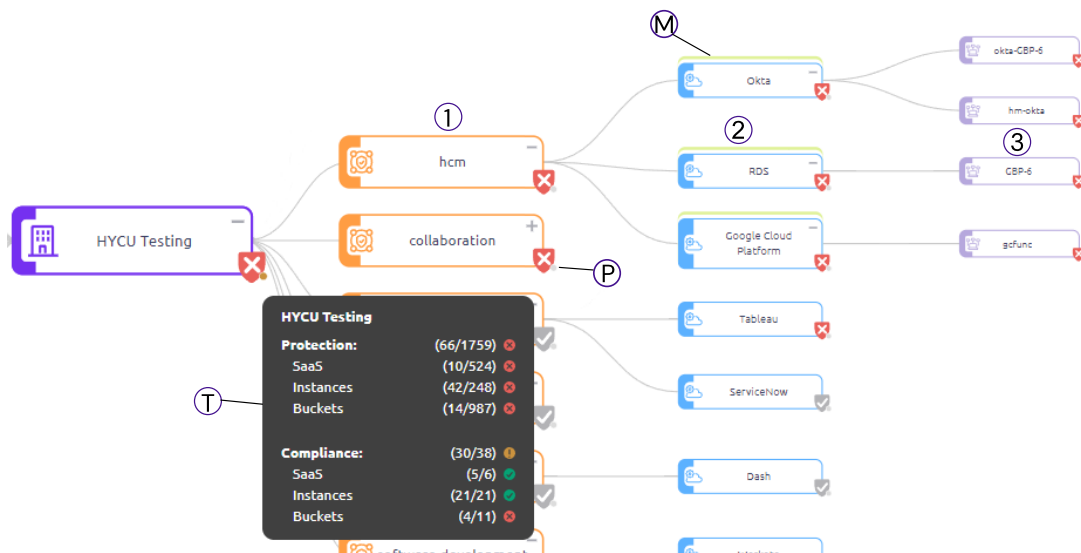


Figure 9-1: R-Graph elements

During the initial scan, R-Cloud matches the discovered SaaS services with existing protection sets and maps them accordingly. If no source for a particular SaaS service is added to any protection set, R-Cloud shows the service as part of a generic protection set that would be automatically created when you add the first source, using a name based on the industry or domain under which the service is classified. If you do not want to create such a protection set, you can add the source to an existing protection set or create a different protection set.

**Note** If a SaaS service is discovered, but R-Cloud cannot determine its type, the service is placed in a protection set placeholder (node) named saas-

| general.

## Node protection status








The node protection status is indicated by status icons, compliance overlay icons, and tooltips with more details.

The node protection status depends on whether the discovered services have sources added or not:

- If sources are added, the indicators reflect the node protection and compliance. For details, see [“Services with added sources”](#) below.
- If no sources are added, the protection indicators reflect the native data protection capabilities of the service. For details, see [“Services discovered through identity providers”](#) on the next page.





### Services with added sources

#### Protection indicators

Status icon	Description
	80% or more included entities have the protection status  .
	Between 60% and 80% of included entities have the protection status  .
	Less than 60% of included entities have the protection status  .
	There are no entities in the node.

#### Compliance indicators

Node compliance is indicated with overlay icons:

Overlay icon	Description
 Green	80% or more included entities have the compliance status Green.
 Yellow	Between 60% and 80% of included entities have the compliance status Green.
 Red	Less than 60% of included entities have the compliance status Green.
 Gray	There are no entities in the node.

### Detailed information about a node

Pause the pointer on the node to display additional information. Each node has the following properties:

Property	Description
Protection	Shows the protection status of the node and individual groups of entities in the node.
Compliance	Shows the compliance status of the node and individual groups of entities in the node.
RPO	<i>Available only for SaaS applications.</i> Shows the RPO as defined by the SaaS application.
Retention	<i>Available only for SaaS applications.</i> Shows the retention period as defined by the SaaS application.

### Protection and compliance calculation and inheritance

The protection and compliance status of a node is calculated based on the status of child entities:

- An entity is protected if it has at least one valid restore point available and the entity has a policy assigned.
- An entity is compliant if the RPO set in the assigned policy is met.
- Only entities with an assigned policy are included in compliance calculation.
- Entities with the Exclude policy assigned are excluded from the protection calculation.

The status of a source, service, or protection set node is based on the status of all entities that are included in the node.



### Services discovered through identity providers

The protection and compliance of services discovered through identity providers are shown as follows:

- A service is not compliant until it is added as a source and is protected by R-Cloud.
- An entity is protected if the RPO of the SaaS application is defined, otherwise it is marked as unprotected.
- When a service with available information about native data protection

capabilities is discovered, the protection status is based on its native protection capabilities.

### Protection indicators

Status icon	Description
	Information about native data protection capabilities is available and one or more of the required capabilities is not available. The actual data protection status of the service is not known.
	Information about native data protection capabilities is available and all required capabilities are available. The actual data protection status of the service is not known.


### Compliance indicators


The compliance indicator is always  Gray until a source is added for the service.

### Detailed information about native data protection capabilities

Pause the pointer on the node to display information about native data protection capabilities:

Capability	Description
Backup automation	Shows if the service offers backup automation. This means that you can: <ul style="list-style-type: none"> <li>Schedule automated backups in a user interface provided by the service.</li> <li>Customize the frequency of the backups to meet your RPO and backup retention needs.</li> </ul>
Off-site storage automation	Shows if the service enables you to automatically export data to an off-site storage location. This means that you can: <ul style="list-style-type: none"> <li>Automatically export data to an off-site storage location outside of the service.</li> <li>Select the target (for example, Google Cloud Storage or S3 compatible storage) for copies of backup data from a user interface provided by the service.</li> </ul>

Self-service restore	Shows if the service provides a user interface to restore deleted or corrupted data from backups.   <b>Note</b> Recycle bins are not included as data is deleted after a certain amount of time and data deleted from a recycle bin is unrecoverable.
----------------------	--

 **Note** R-Graph only shows if native data protection capabilities are available. The actual data protection status of the service is not known.

## Managing identity and access

You can use the Identity and access management (IAM) panel to manage identity providers, users, and user roles in R-Cloud.

The scope of tasks you can perform depends on your assigned roles and the selected user interface context:

- **Subscription:**

Task	Instructions
Add, edit, or remove identity providers from R-Cloud.	“Managing identity providers” on the next page
Add, deactivate, or remove users.	“Managing users” on page 223
Add or remove user roles.	“Managing roles” on page 225
Send password reset requests	“Requesting a password reset” on page 227

- **Protection set:**

Task	Instructions
Add users.	“Managing users” on page 223
Assign or unassign user roles.	“Managing roles” on page 225

Send password reset requests	<a href="#">“Requesting a password reset” on page 227</a>
------------------------------	---

### Accessing the IAM panel

To access the IAM panel, in the navigation pane, click **IAM**.

## Managing identity providers

You can integrate R-Cloud with identity providers that support the OpenID Connect authentication protocol, such as Google, Microsoft, and Okta, to give users the possibility to securely sign in to R-Cloud by using these identity providers, without the need to maintain dedicated credentials for R-Cloud.

### Prerequisites

*Only when adding identity providers that support the OpenID Connect authentication protocol.* R-Cloud must be registered as a web application within the identity provider that you plan to add to R-Cloud. When registering R-Cloud, make sure the following is done:

- *Only if you are using Microsoft as an identity provider.* In Azure, R-Cloud must be given access permissions to the following Azure API: Microsoft Graph with delegated permissions for User . Read.
- *Only if you are using Okta as an identity provider.* In Okta, you must select **Authorization Code** under Client acting on behalf of a user as the grant type.


For instructions on how to register an application, see the respective identity provider documentation.

### Accessing the Identity Providers dialog box

To access the Identity Providers dialog box, in the Subscription context, in the IAM panel, click  **Identity Providers**.


## Adding an identity provider to R-Cloud

### Procedure



1. In the Identity Providers dialog box, click  **New**.
2. Enter a name for the identity provider. The name that you specify can contain only lowercase letters and hyphens, must begin and end with a lowercase letter, and cannot be longer than 63 characters.

3. From the Type drop-down menu, select one of the following types of identity providers, and then follow the instructions:

Identity provider type	Instructions
<b>Google</b>	a. In the Client ID field, enter the application ID that is generated by the identity provider. b. In the Client secret field, enter the application secret that is associated with the client ID and generated by the identity provider.
<b>Microsoft</b>	a. In the Client ID field, enter the application ID that is generated by the identity provider.
<b>Okta</b>	b. In the Client secret field, enter the application secret that is associated with the client ID and generated by the identity provider.
<b>OIDC</b>	
<b>Cognito</b>	c. In the Issuer field, enter the URL of the issuer of the identity provider.

4. Click  **Copy to Clipboard** to copy the redirect URL that you need to input when you create the application integration with R-Cloud.
5. Click **Save**.
6. Configure your identity provider and enter the redirect URL that you copied. For details on the required format, see the respective identity provider documentation.

You can later do the following:

- Edit information about any of the existing identity providers by clicking  **Edit** and making the required modifications.
- Delete any of the existing identity providers by clicking  **Delete**.

## Managing users


The R-Cloud user management system provides security mechanisms to help prevent unauthorized users from accessing protected data. Only users that are given specific rights have access to the data protection environment. These users can be authenticated either by HYCU or any of the supported identity


providers. For details on identity providers, see [“Managing identity providers” on page 222](#).

### Consideration


The scope of tasks you can perform depends on the selected UI context. In the Protection set context, you can only add users but cannot deactivate or remove them.

## Adding a user

1. In the IAM panel, click  **New User**. The New User dialog box opens.
2. Enter the email address of the user that you want to add.
3. *Optional, if the user will sign in using an identity provider.* Select **Generate password** to automatically generate a password. The user must change the generated password during the first sign-in.

 **Important** If the user has no identity provider configured and you do not generate a password, the user will not be able to sign in to R-Cloud.

4. *Only if you are adding a user in the Subscription context.* Select one of the following options:
  - **Assign to subscription**  
Assign the user to the subscription.
  - **Assign to protection set**  
From the list of protection sets, select the one to which you assign the user.

 **Tip** You can search for a protection set by entering its name in the Protection set search field and then pressing **Enter**. By selecting the Name check box, you select all protection sets at once.

5. From the Role drop-down menu, select the role for the user.  
You can select more than one role if needed. For more information about user roles, see [“R-Cloud roles” on page 226](#).
6. Click **Save**.

## Deactivating a user

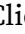
### Consideration

When you deactivate a user, the user can no longer perform any actions. However, the inactive account is preserved in cloud, including all the data that



the user has backed up.

#### Procedure

1. In the IAM panel, from the list of available users, select the user that you want to deactivate.
2. Click  **Deactivate**. The Deactivate User dialog box opens.
3. Click **Deactivate** to confirm the deactivation of the user.


## Deleting a user


#### Considerations

- Deleting a user from R-Cloud does not remove it from cloud.
- You cannot delete yourself from R-Cloud.
- Any upcoming data protection tasks related to the user that you delete will be automatically assigned to you.

#### Procedure

1. In the IAM panel, from the list of available users, select the one that you want to delete.

 **Tip** You can also search for a user by entering their name in the Search field.

2. Click  **Remove**. The Remove Account dialog box opens.
3. Click **Remove** to confirm that you want the selected user to be deleted from R-Cloud.

## Managing roles

A role determines the scope of actions that can be performed in the R-Cloud data protection environment by a specific user or service account. This means that access to data and information within the data protection environment is limited based on the assigned role. As an administrator, you can manage these roles and define what actions can be performed by each user or service account.

#### Considerations

- Each user that signs in to R-Cloud or each configured service account has by default the Administrator role assigned.

- At least one user with the Administrator role assigned must exist in the data protection environment for each subscription, at the subscription level.
- User roles are inherited from the subscription level to all protection sets under one subscription. User roles set in a protection set are local to that protection set.

## R-Cloud roles

A user or a service account can be assigned one or more of the following roles:

Role	Allowed actions
Administrator	Perform all actions in the data protection environment.
Backup Operator	Define backup strategies, back up SaaS applications, applications, instances, and buckets, and acquire the same information as Viewer.
Restore Operator	Restore SaaS applications, applications, instances, and buckets, and acquire the same information as Viewer.
Viewer	Acquire information about SaaS applications, applications, instances, buckets, policies, targets, tasks, events, reports, service accounts, and protection sets in the data protection environment.


## Assigning or unassigning roles

### Consideration

If you plan to remove your own Administrator role, keep in mind the following:

- At least one user with the Administrator role assigned must exist in the data protection environment for each subscription.
- You will not be able to change your role back to Administrator yourself.


### Procedure

1. In the IAM panel, from the list of available users, select the user for whom you want to change the roles and then click  **Edit**.
2. In the Edit Role dialog box, from the drop-down list, select the roles that you want to assign or unassign. You can select or deselect roles individually or you can click **Select all** to select all roles at once.
3. Click **Save** to save the selected roles.

## Requesting a password reset

If a user signs in to R-Cloud by using the HYCU credentials and their password should be changed due to company policy requirements or safety reasons, send the user a password reset request.

### Procedure

1. In the IAM panel, from the list of available users, select the user that should reset their password, and then click  **Edit**.
2. Click **Request password reset**.
3. Click **Request password reset** to confirm that you want to request a password reset for this user.

The user will receive an email containing the password verification code that allows them to reset the password the next time they sign in to R-Cloud.

## Stopping protection for individual sources

This section provides instructions that you must follow to stop protecting individual sources in R-Cloud.

 **Note** If you want to stop using R-Cloud completely, see [“Unsubscribing from R-Cloud” on page 235](#).

### Procedure

1. In R-Cloud, unassign policies from all protected entities in the source. For instructions, see [“Stopping service charges” on page 235](#).
2. In R-Cloud, manually mark restore points of all entities in the source as expired. For instructions, see [“Expiring backups manually” on page 188](#).
3. Remove the source from any protection set. For instructions, see [“Editing protection sets” on page 196](#).

When a source is no longer protected, irrelevant notifications are prevented, and the unneeded associated charges are avoided.

# Excluding instances from synchronization by tagging the instance in AWS or Google Cloud

This section provides information on how to make selected instances invisible to R-Cloud. The needs of your environment may require that some instances are not protected by R-Cloud. For example, your Google Cloud projects may include managed instance groups and employ an autoscaler. To leave some instances unprotected, you can exclude them from synchronization so that they are not visible to R-Cloud. The invisible instances cannot be assigned policies in any way.

## Procedure

1. Depending on your cloud platform, do the following:

- *For AWS:*
  - a. In the AWS Management Console, choose the AWS account to which the instances that you want to leave unprotected belong.
  - b. Within the AWS account, choose an instance and add it the `hycu-instance-sync` tag in Amazon EC2. Use the following key/value pair:



Key	Value
<code>hycu-instance-sync</code>	<code>false</code>

Custom tags can be added from the Amazon EC2 console. For instructions, see AWS documentation.

- *For Google Cloud:*
  - a. In the Google Cloud Console, choose the Google Cloud project to which the instances that you want to leave unprotected belong.
  - b. Within the project, choose an instance and add it the `hycu-instance-sync` custom metadata tag in Google Compute Engine. Use the following key/value pair:

Key	Value
<code>hycu-instance-sync</code>	<code>false</code>

Custom metadata tags can be added from the Google Cloud Console, the `gcloud` command line, or by using the Google Cloud API. For instructions, see Google Cloud documentation.

2. Repeat step 1b for each additional instance that you want to make invisible to R-Cloud.
3. Sign in to the R-Cloud web user interface.
4. Select the protection set that includes the same AWS account or Google Cloud project as you selected in step 1 of the procedure. For instructions on selecting protection sets in R-Cloud, see [“Selecting an R-Cloud protection set” on page 26](#).
5. In the navigation pane, click  **Instances**.
6. Click  **Refresh** or wait until the next instance synchronization cycle.

In the Instances panel, the names of the instances that you excluded from synchronization are not present.

# Chapter 10

## Troubleshooting

If you encounter a problem while using R-Cloud, use the following approach to troubleshoot it:

1. Check if your problem is described in [“Known problems and solutions” on the next page](#) and apply the recommended solution.
2. If you cannot find the problem in the list of the known problems, try to solve it on your own. When doing so, you first need to identify the cause of the problem, collect and analyze all available information about it, and then solve the problem. Answering the following questions may help you to solve your problem:
  - a. Did you fulfill all the prerequisites and are you aware of all the limitations that come with R-Cloud?
  - b. Do you receive any errors?

You can view all events that occurred in your environment in the Events panel. In addition, you can track tasks that are running in your data protection environment and get an insight into the specific task status. For this purpose, use the Tasks panel. For detailed information on events and tasks, see [“Viewing events” on page 170](#) and [“Checking task statuses” on page 169](#).
  - c. Is your problem related to any third-party hardware or software?

In this case, contact the respective vendor for support.
3. If the problem still persists, contact [HYCU Support](#). It is recommended that you collect and send the following information to HYCU Support:
  - Description of your data protection environment
  - Description of your problem
  - Results of any testing you have done (if available)

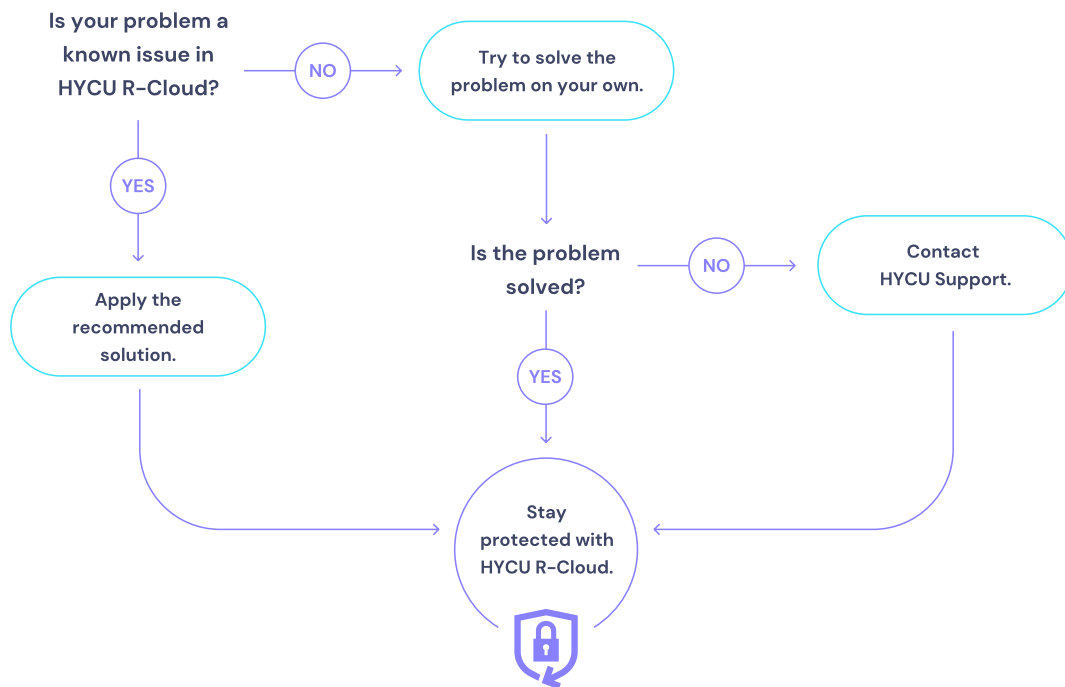


Figure 10-1: Major steps of the troubleshooting process

## Known problems and solutions

This section lists all known problems that you may encounter while using R-Cloud, along with their solutions.

### Restore of individual files ends with errors or fails

#### Problem

When a restore of individual files completes, the status of the corresponding task is Done with errors or Failed. Closer inspection reveals that some or all of your selected resources have not been restored.

#### Cause

The original disk no longer exists, or the credential group that is assigned to the original instance in R-Cloud includes a user account with insufficient privileges.

### Solution

Restore your files to an alternate location on the original instance, to a custom location on a different instance, or to an available bucket, or update the configuration of the credential group that is assigned to the original instance in R-Cloud.

## Inability to change the protection set or to sign in

### Problem

Although you have access to Google Cloud projects that are included in multiple protection sets in R-Cloud, only the currently selected protection set is available in the Protection set UI context. After your web user interface session ends, you are unable to sign in again.

### Solution

Contact [HYCU Support](#).

## Problem with sorting data in the Events panel

### Problem

In the R-Cloud web user interface, sorting data in the Events panel by the Message parameter does not function properly.

### Solution

There is no solution available for this problem. If data disappears while being sorted, sign out of the web user interface and sign in again to repopulate the table with data.

## Inability to set up manually created Google Cloud targets

### Problem

When you try to set up a manually created target, R-Cloud reports that the target is inaccessible.



### Solution

In the Google Cloud Storage service, grant your Google Account the Storage Admin role on the Google Cloud project of the target.

For information on the required roles for the general use of the service, see [“Signing in to R-Cloud” on page 22](#).

## Assigning a policy to a Google Cloud instance fails

### Problem

After adding the `hycu-policy` custom metadata tag to an instance in Google Compute Engine, no policy is assigned to the instance in R-Cloud.

### Cause

The symptom may indicate one of the following:

- The instance belongs to a project that is not included in any protection set.
- The policy that is specified for the metadata tag value does not exist.

### Solution

Find the corresponding entry in the event log to identify the root cause of the problem:

1. In the R-Cloud web user interface, go to the Events panel and search for the following error message:

```
Failed to assign a policy
```

2. Click the message entry, check the Message details section for the root cause of the problem, and act accordingly.

## Snapshot creation fails for instances in a specific Google Cloud project

### Problem

When a backup task for any instance in a specific Google Cloud project is started, the snapshot creation task fails and reports an error.

## Solution

In Google Compute Engine, grant your Google Account the Compute Admin role on the Google Cloud project.

For information on the required roles for general use of the service, see [“Signing in to R-Cloud” on page 22](#).

## Task progress indicator remains at 0% during the backup of a Google Cloud instance

### Problem

You experience one of the following symptoms:

- When you start a backup task, its child task for creating disk catalog never makes any progress.
- After you start a backup or restore task, the task gets started, but it never makes any progress.

### Solution

Check if the Google Cloud project that the instance belongs to has the Cloud Pub/Sub API enabled. If it does not, enable the API for the project through the Google Cloud Console.

# Chapter 11

## Unsubscribing from R-Cloud



If for whatever reason you decide that you no longer want to use R-Cloud for protecting your data, you can easily unsubscribe from the service.









Unsubscribing from R-Cloud includes the following tasks:

Task	Instructions
1. Stop being charged for using R-Cloud.	<a href="#">“Stopping service charges” below</a>
2. Prevent R-Cloud to access your account.	<a href="#">“Preventing account access” on page 237</a>
3. <i>Optional.</i> Remove the HYCU Managed Service Account permissions.	<a href="#">“Removing the HYCU Managed Service Account permissions” on page 239</a>
4. Cancel your R-Cloud subscription in the cloud.	<a href="#">“Canceling your R-Cloud subscription” on page 239</a>

## Stopping service charges

To avoid unnecessary charges for the backup and recovery service, perform the following tasks:

Task	Instructions
1. Stop charges for backup and recovery.	<p>In R-Cloud, unassign policies from all protected entities:</p> <ul style="list-style-type: none"><li>• To unassign policies from SaaS applications:<ol style="list-style-type: none"><li>1. In the navigation pane, click  <b>SaaS</b>.</li><li>2. Select all SaaS applications with assigned policies, and then click  <b>Set Policy</b>.</li><li>3. Click <b>Unassign</b>, and then click <b>Yes</b> to confirm that</li></ol></li></ul>

	<p>you want to unassign the policies from the selected SaaS applications.</p> <ul style="list-style-type: none"> <li>• To unassign policies from applications: <ol style="list-style-type: none"> <li>1. In the navigation pane, click  <b>Applications</b>.</li> <li>2. Select all applications with assigned policies, and then click  <b>Set Policy</b>.</li> <li>3. Click <b>Unassign</b>, and then click <b>Yes</b> to confirm that you want to unassign the policies from the selected applications.</li> </ol> </li> <li>• To unassign policies from instances: <ol style="list-style-type: none"> <li>1. In the navigation pane, click  <b>Instances</b>.</li> <li>2. Select all instances with assigned policies, and then click  <b>Set Policy</b>.</li> <li>3. Click <b>Unassign</b>, and then click <b>Yes</b> to confirm that you want to unassign the policies from the selected instances.</li> </ol> </li> <li>• To unassign policies from buckets: <ol style="list-style-type: none"> <li>1. In the navigation pane, click  <b>Buckets</b>.</li> <li>2. Select all buckets with assigned policies, and then click  <b>Set Policy</b>.</li> <li>3. Click <b>Unassign</b>, and then click <b>Yes</b> to confirm that you want to unassign the policies from the selected buckets.</li> </ol> </li> </ul> <p> <b>Important</b> If multiple protection sets are available in your data protection environment, make sure to follow these steps for each protection set separately.</p>
<p>2. Stop charges for backup data storage.</p>	<ol style="list-style-type: none"> <li>1. Manually mark restore points of all entities as expired. For instructions, see <a href="#">“Expiring backups manually” on page 188</a>.</li> </ol> <p> <b>Important</b> If multiple protection sets are available in your data protection environment, make sure to do this for each protection set separately.</p>

	<p>2. <i>Only if SAP HANA application data was backed up by using the Backint agent.</i> Disable log backups and remove all existing log backups from the Google Cloud buckets from the following location:</p> <pre>&lt;SAPHANAAppName&gt;/usr/sap/ &lt;SAPHANAAppName&gt;/SYS/global/ hdb/backint/&lt;DatabaseName&gt;</pre> <p>For details on how to disable log backups, see SAP HANA documentation.</p> <p>3. Remove all backup data created by R-Cloud from cloud (delete all automatically or manually created targets that contain only backup data, and delete all backup data that is stored on automatically or manually created targets that contain also other kind of data). For the target naming conventions, see <a href="#">“Resources created by R-Cloud” on page 242</a>. For instructions on how to delete targets and remove backup data from targets, see the respective cloud documentation.</p> <p>4. Remove all snapshots created by R-Cloud from cloud. For the snapshot naming conventions, see <a href="#">“Resources created by R-Cloud” on page 242</a>. For instructions on how to remove snapshots, see the respective cloud documentation.</p>
--	--

## Preventing account access

As part of unsubscribing from R-Cloud, you must prevent R-Cloud to access your account.


Cloud platform	Instructions
AWS	<a href="#">“Preventing access to an AWS account” on the next page</a>
Google Cloud	<a href="#">“Preventing access to a Google Cloud account” on the next page</a>

## Preventing access to an AWS account

When you added an account as a source to R-Cloud, you assigned R-Cloud IAM roles to your AWS account. After you stop using the solution, you must remove the roles.

### Procedure

1. Open a web browser, go to the [Sign in page](#) of the AWS Management Console and sign in.
2. Open the AWS CloudFormation console and in the navigation pane, choose **Stacks**.
3. In the list of stacks, select `CreateHycuRole` and delete it. When prompted, confirm the deletion.

 **Note** If multiple sources are available in your data protection environment, make sure to follow these steps for each source.

For details on removing AWS stacks, see AWS documentation.

## Preventing access to a Google Cloud account

When you subscribed to R-Cloud, you granted it access to your Google Account. After you stop using the solution, you must remove the access permission.

### Procedure

1. Open a web browser, go to the [Sign in & security](#) page of the Google website, and then click **Sign in**.
2. Sign in with your Google Account.
3. Click **Security**.
4. Locate the Third party apps with account access section, and then click **Manage third party access**.
5. Under Third party apps with account access, click **HYCU R-Cloud**, and then click **REMOVE ACCESS**.
6. Click **OK** to confirm that you want to remove the access permission.

For information on access permissions, see Google Cloud documentation.

## Removing the HYCU Managed Service Account permissions

After you cancel your R-Cloud subscription, your HYCU Managed Service Account (HMSA) is kept together with other data for 14 days before it is permanently deleted. However, if for any reason you want to remove the HMSA permissions immediately, you can do it by using one of the following methods:

Method	Instructions
Manual	In Google Cloud, remove the HMSA permissions. For instructions on how to remove service account permissions, see Google Cloud documentation.
Automatic	Click the following link to open Google Cloud Shell, and then follow the instructions in the tutorial: <a href="#">Open Google Cloud Shell</a>

**ⓘ Important** If you remove the HMSA permissions by using either of these methods, keep in mind that to add the HMSA back to R-Cloud, you will have to grant the HMSA the following roles in Google Cloud on each project that you plan to protect:

- Compute Admin, Service Account User, and Storage Admin
- *Required only if protecting GKE applications.* Kubernetes Engine Admin

For instructions on how to grant permissions to service accounts, see Google Cloud documentation.

## Canceling your R-Cloud subscription

As part of unsubscribing from R-Cloud, you must cancel your R-Cloud subscription.

Cloud platform	Instructions
AWS	<a href="#">“Canceling the R-Cloud subscription in the AWS Marketplace” on the next page</a>
Google Cloud	<a href="#">“Canceling the R-Cloud subscription in the Google Cloud Marketplace” on the next page</a>

## Canceling the R-Cloud subscription in the AWS Marketplace

### Prerequisite

Your user account has the `AWSMarketplaceManageSubscriptions` predefined role assigned.

### Procedure

1. Open a web browser and go to the [HYCU | AWS Market](#) webpage.
2. Search for HYCU R-Cloud to find your subscription.
3. On the Manage Subscription page, cancel the subscription. For details on how to cancel an AWS subscription, see AWS documentation.

After you cancel your R-Cloud subscription, your data is kept for 14 days before it is permanently deleted. If during this period you change your mind and you want to continue using R-Cloud, resubscribe from the same account.

## Canceling the R-Cloud subscription in the Google Cloud Marketplace

### Prerequisites

- You are signed in to Google with a Google Account that is granted the Billing Account Administrator (`roles/billing.admin`) role on the billing account that is used for the R-Cloud subscription.
- Your currently selected project in the Google Cloud Console is linked to the billing account that is used for the R-Cloud subscription.

### Procedure

1. Open a web browser and go to the [HYCU | Marketplace - Google Cloud](#) webpage.
2. Navigate to the Pricing section, and then click **Manage orders**.
3. Select the billing account that is used for the R-Cloud subscription.
4. On the list of your orders, select **Turn off auto-renewal** for your R-Cloud subscription.



After you cancel your R-Cloud subscription, your data is kept for 14 days before it is permanently deleted. If during this period you change your mind and you want to continue using R-Cloud, resubscribe to R-Cloud from the project that belongs to the billing account that was used for the R-Cloud subscription.

# Appendix A

## Resources created by R-Cloud

During data protection tasks, R-Cloud creates temporary and persistent HYCU resources in your AWS accounts or Google Cloud projects. Temporary HYCU resources exist only for the duration of a task, and persistent HYCU resources are preserved after tasks are completed.

**⚠ Caution** With the exception of the restored files and unless specifically instructed to do so, never rename or delete any HYCU resources.

Names or location path templates of persistent HYCU resources created during backup tasks

- Snapshot:
  - *For AWS:*  
HYCU-*<Instance>*-snapshot
  - *For Google Cloud:*  
hycu-snap-*<TaskUUID>*-*<Disk>*
- Automatically created target:  
hycu-*<CloudStorageRegion>*-*<UUID>*
- Target folder with a backup, a backup copy, or a data archive:  
hycu/backups/*<Source>*/*<Region/Zone>*/*<Instance>*/disks/*<Disk>*/*<StorageClass>*
- Target folder with a disk catalog:  
hycu/backups/*<Source>*/*<Region/Zone>**<Instance>*  
/tasks/*<TaskUUID>*/*<Disk>*

Names or location path templates of persistent HYCU resources created during restore tasks

- Renamed original file (at the original location on an instance):  
*<OriginalFileName>*.hycu.orig[.*<OriginalFileExtension>*]
- Renamed restored file (at the original location on an instance):

- *For AWS:*  
`<OriginalFileName>  
[.<OriginalFileExtension>].<TimeStamp>.restored`
- *For Google Cloud:*  
`<OriginalFileName>.hycu.restored[.<OriginalFileExtension>]`
- Target folder with restored files or folders:
  - *For AWS:*  
`hycu/restores/<Source>/<Instance>/<TaskUUID>/<Path>`
  - *For Google Cloud:*  
`hycu/restores/<Source>/<Zone>/<Instance>/  
<TaskUUID>/<Disk>/<Volume>/<Path>`
- Restored file:  
`<FileName>.<FileExtension>.<TimeStamp>.restored`
- *For Google Cloud:* External IP address resource automatically allocated by R-Cloud during cloning:  
`hycu-static-external-<UUID>`
- *For Google Cloud:* Internal IP address resource automatically allocated by R-Cloud during cloning:  
`hycu-static-internal-<UUID>`
- Cloned disk, attached to an instance:
  - *For AWS:*  
`<OriginalDiskName>`
  - *For Google Cloud:*  
`hycu-disk-<TaskUUID>-<UUID>-<Disk>`
- Cloned or moved disk, unattached:
  - *For AWS:*  
`hycu-export-<Disk>`
  - *For Google Cloud:*  
`hycu-disk-<TaskUUID>-<UUID>-<Disk>`

### Name templates of temporary HYCU resources created during backup and restore tasks

- Temporary disk:
  - *For AWS:*  
`hycu-temporary-<SnapshotID>`

- *For Google Cloud:*

`hycu-disk-tmp-<TaskUUID>-<OriginalDiskName>`

## Appendix B

# Bulk restore specifications

Based on the bulk restore options you specify when restoring multiple instances or disks belonging to multiple instances, R-Cloud generates a bulk restore specification.

## Elements of a bulk restore specification

The basic elements of a bulk restore specification include the type of the bulk restore specification (`bulkRestoreType`), a flag whether to overwrite existing items (`overwriteExisting`), and the items to restore.

### Syntax

```
{
  "bulkRestoreType": "VMS" | "DISKS",
  "overwriteExisting": false | true,
  "items": [
    {
      "source":
      {
        "path": "<path>",
        "disks": [<disk>,...]
      },
      "destination":
      {
        "path": "<path>",
        "disks": [<disk>, ...],
        "networkInterfaces": [<networkInterface>, ...],
        "metadata": {},
        "labels": {},
        "tags": []
      }
    }
  ]
}
```

```

    },
  },
  ...
],
}

```

### Basic elements

- `bulkRestoreType`: "VMS" | "DISKS"  
The bulk restore type, VMS for instances and DISKS for disks.
- `overwriteExisting`: false | true  
If set to true, instances in the destination region and zone with the same name as the source instances or disks attached to instances, with the same name as the source disks are overwritten during restore.  
Default: false
- `items` []  
An array of items to restore, each item element contains a source and a destination.

### Items to restore

Each item consists of a source and a destination record:

- `source`  
The source record contains the path and an array listing the disks.
- `destination`  
The destination record contains the path, an array listing the disks, an array listing the network interfaces, and tags and labels.

### Source and destination elements

- `path`  
The path in the format  
`projects/<Project>/zones/<TargetZone>/instance/<InstanceName>`.
- `disks`  
An array containing the disks to be restored. Disks can either contain the disk name for the source disks or a record with the following elements in the case of destination disks:

- `sourceType`  
The source type of the original disk: AWS for AWS, or GC for Google Cloud.
- `diskName`  
The name of the original disk.
- `newDiskName`  
The name of the restored disk, including the specified postfix. If no postfix is specified, the new name equals the original disk name.
- `newDeviceName`  
The name of the restored disk device, including the specified postfix. If no postfix is specified, the new name equals the original device name.
- `diskType`  
One of the available disk types for the restored disk.  
For AWS, this is I02 or I01 for provisioned IOPS SSD disks, GP3 or GP2 for general purpose SSD disks, SC1 for cold HDD disks, and ST1 for throughput optimized HDD disks.  
For Google Cloud, this is BALANCED for balanced persistent disks, EXTREME for extreme persistent disks, HDD for standard persistent disks, or SSD for SSD persistent disks.  
By default, the original disk type is used.
- `region`  
Specifies the region. You can use it to define a different destination region for the disk.
- `replicaZones`  
An array listing the replica zones.  
For regional disks, by default only the same replica zone as in the path is added. You need to add the second zone.
- `networkInterfaces`  
An array containing network interfaces. Each interface is a record with the following elements:
  - `sourceType`  
The source type of the network interface: AWS for AWS, or GC for Google Cloud.
  - `path`  
The path to the network device.

- `externalIpType`  
The external IP type for the network interface.  
For AWS, the `AUTO_ASSIGN` option is always selected.  
For Google Cloud, you can select among the following options: `NONE`, `Ephemeral`, `STATIC_RESERVED`, `STATIC_NEW`.
- `externalIp`  
The external IP value, if supported by the external IP type.
- `internalIpType`  
The internal IP address type for the network interface.  
For AWS, the `AUTO_ASSIGN` option is always selected.  
For Google Cloud, you can select among the following options: `Ephemeral_Automatic`, `Ephemeral_Custom`, `STATIC_RESERVED`, `STATIC_NEW`.
- `internalIp`  
The internal IP value, if supported by the internal IP type.

The interfaces are selected in the following order:

1. A legacy network with same name.
  2. Shared subnetworks that are accessible in destination projects.
  3. A subnetwork with the name "default".
  4. The first subnetwork in the specified region (sorted by name alphabetically).
- Labels, metadata, and tags
    - `metadata`  
Custom metadata tags, consisting of a key and a value pair.
    - `labels`  
Labels for the restored disk, consisting of a key and a value pair.
    - `tags`  
An array of tags (strings).
  - *For AWS*: Operating system image AMI ID
    - `imageId`  
AMI ID of the custom operating system image.



## Appendix C

# Least-privilege permissions used by R-Cloud

To access your data protection environment and perform different tasks such as discovering entities, backing up data, and restoring data, R-Cloud does the following:


- *For AWS:* Creates an IAM role for your AWS account with a predefined set of permissions.
- *For Google Cloud:* Uses the permissions that you granted to the Google Account, the Google Service Account, or the HMSA in Google Cloud.

However, if you need to create a custom role with the least-privilege permissions needed to access your data protection environment, you can use the R-Cloud role template that contains a predefined set of these permissions. Depending on your cloud platform, see one of the following sections:

Cloud platform	Instructions
AWS	<a href="#">“Using a role template for AWS” below</a>
Google Cloud	<a href="#">“Using a role template for Google Cloud” on page 254</a>


## Using a role template for AWS

### Prerequisite

You must have the HYCU account ID of your subscription. To get the HYCU account ID, click  **<EmailAddress>** in the toolbar, and then click **Subscription Information** to open the Subscription Information dialog box. The account ID is listed under the HYCU Account section.

## Considerations

- Make sure that the account for which you are creating the role is not already added as a source in R-Cloud, otherwise the creation of the least-privileges role will fail and the role with default permissions will stay in place. If you already added the account as a source, delete its role or the AWS CloudFormation stack with which you created the original role before you start the process or use a different account.
- By default, the IAM role created by R-Cloud allows performing data protection actions on all resources in your AWS account. In addition to creating a custom role with the least-privilege permissions as described in this section, you can also limit R-Cloud to have access only to the resources that are relevant for data protection by specifying them in the AWS Management Console. You should do this each time you add new resources that should be accessible to R-Cloud.

 **Note** To make sure that the targets that are automatically created by R-Cloud will be accessible, keep the default resource accessibility settings for the automatically created targets. For details on how to identify the automatically created targets, see “Resources created by R-Cloud” on page 242.

For details on how Amazon S3 works with IAM and how to specify the resources that should be accessible to R-Cloud, see AWS documentation.


## Procedure

To add the role template to your AWS account, perform the following:

1. Open the following URL in your browser:

[https://us-east-2.console.aws.amazon.com/cloudformation/home?region=us-east-2#/stacks/quickcreate?templateUrl=https%3A%2F%2Fhycu-resources.s3.amazonaws.com%2Fcloudformation%2F08082022-HycuRoleTemplate-AWSLeastPermissions.json&stackName=HycuStack&param\\_ExternalId=<HycuAccountId>](https://us-east-2.console.aws.amazon.com/cloudformation/home?region=us-east-2#/stacks/quickcreate?templateUrl=https%3A%2F%2Fhycu-resources.s3.amazonaws.com%2Fcloudformation%2F08082022-HycuRoleTemplate-AWSLeastPermissions.json&stackName=HycuStack&param_ExternalId=<HycuAccountId>)

In this URL, *<HycuAccountId>* at the end of the URL is the account ID of your subscription.

 **Important** You must be signed in to the AWS Management Console with the account for which you are creating roles. If you are already signed in to the AWS Management Console with a different account when

- you create the IAM roles, the creation fails.
- In the AWS Management Console, on the Quick create stack page, confirm the capabilities required by R-Cloud by clicking **I acknowledge that AWS CloudFormation might create IAM resources**, and then click **Create stack**.

## AWS permissions required by R-Cloud

The following is a list of AWS permissions required by R-Cloud:

Service	Permissions
S3	ListAllMyBuckets ListBucket ListBucketVersions GetBucketLocation GetBucketLogging GetBucketObjectLockConfiguration GetBucketPublicAccessBlock GetBucketTagging GetBucketVersioning GetEncryptionConfiguration GetLifecycleConfiguration GetObject GetObjectTagging DeleteJobTagging DeleteObject DeleteObjectVersion DeleteObjectTagging DeleteObjectVersionTagging DeleteStorageLensConfigurationTagging PutBucketTagging PutJobTagging PutObjectTagging PutObjectVersionTagging PutStorageLensConfigurationTagging ReplicateTags CreateBucket PutObject  <i>For targets that have Object Lock (WORM)</i>

	<i>enabled, the following additional permissions are required:</i> PutObjectRetention PutObjectLegalHold
S3 Express	CreateSession ListAllMyDirectoryBuckets
STS	AssumeRole
SQS	GetQueueUrl ListQueues ReceiveMessage CreateQueue DeleteMessage DeleteQueue SendMessage
IAM	GetAccountSummary PassRole
EC2	DescribeAddresses DescribeAvailabilityZones DescribeInstances DescribeInstanceStatus DescribeInstanceTypes DescribeRegions DescribeSecurityGroups DescribeSnapshots DescribeSubnets DescribeVolumes GetConsoleOutput CreateTags AllocateAddress AssociateAddress AttachVolume CopyFpgaImage CopyImage CopySnapshot CreateNetworkInterface

	<p>CreateSnapshot  CreateSnapshots  CreateVolume  DeleteSnapshot  DeleteVolume  DeregisterImage  DetachVolume  ImportImage  ImportInstance  ImportKeyPair  ImportSnapshot  ImportVolume  RegisterImage  RunInstances  StartInstances  StopInstances  TerminateInstances</p>
Elastic Block	<p>Store CompleteSnapshot  StartSnapshot  GetSnapshotBlock  ListChangedBlocks  ListSnapshotBlocks  PutSnapshotBlock</p>
SNS	<p>ListSubscriptions  ListSubscriptionsByTopic  ListTopics  GetSubscriptionAttributes  GetTopicAttributes  ListTagsForResource  TagResource  UntagResource  ConfirmSubscription  CreateTopic  DeleteTopic  Publish  SetSubscriptionAttributes  SetTopicAttributes  Subscribe</p>

	Unsubscribe
S3 Object Lambda	ListBucket ListBucketMultipartUploads ListBucketVersions ListMultipartUploadParts GetObject GetObjectRetention PutObject PutObjectLegalHold PutObjectRetention RestoreObject WriteGetObjectResponse

## Using a role template for Google Cloud

### Prerequisite

Your account must have the `iam.roles.create` permission. If you are a project or organization owner, you have this permission by default. If you are not an owner, you must have either the Organization Role Administrator or the IAM Role Administrator role assigned.

### Procedure

1. Download the R-Cloud service role template that contains the role definitions. The template is available at the following location:  
[https://storage.googleapis.com/hycu-public/custom-role/hycu\\_service\\_role.yaml](https://storage.googleapis.com/hycu-public/custom-role/hycu_service_role.yaml)
2. Create a role and grant it the permissions required by R-Cloud. To do so, run the following command:

```
gcloud iam roles create <RoleID> --project=<ProjectID> --
file=<RoleDefinitionFilePath>
```

In this command, `<RoleID>` is the name of the role (for example `hycuRole`), `<ProjectID>` is the name of your project, and `<RoleDefinitionFilePath>`

is the path to the location of the downloaded template that contains the custom role definition.

For details on creating and managing custom roles, see Google Cloud documentation.

## Google Cloud permissions required by R-Cloud

The following is a list of Google Cloud permissions required by R-Cloud:

Service	Permissions
Mandatory for all services	resourcemanager.projects.get
Google Compute Engine	compute.acceleratorTypes.get compute.addresses.create compute.addresses.createInternal compute.addresses.get compute.addresses.list compute.disks.create compute.disks.createSnapshot compute.disks.delete compute.disks.get compute.disks.list compute.disks.setLabels compute.disks.use compute.disks.useReadOnly compute.firewalls.get compute.firewalls.list compute.firewalls.update compute.globalOperations.get compute.images.getFromFamily compute.images.getIamPolicy compute.images.setIamPolicy compute.images.useReadOnly compute.instances.attachDisk

---

<code>compute.instances.create</code>
<code>compute.instances.delete</code>
<code>compute.instances.deleteAccessConfig</code>
<code>compute.instances.detachDisk</code>
<code>compute.instances.get</code>
<code>compute.instances.getSerialPortOutput</code>
<code>compute.instances.list</code>
<code>compute.instances.setLabels</code>
<code>compute.instances.setMachineType</code>
<code>compute.instances.setMetadata</code>
<code>compute.instances.setServiceAccount</code>
<code>compute.instances.setTags</code>
<code>compute.instances.start</code>
<code>compute.instances.stop</code>
<code>compute.instances.update</code>
<code>compute.licenses.get</code>
<code>compute.machineImages.useReadOnly</code>
<code>compute.machineTypes.get</code>
<code>compute.machineTypes.list</code>
<code>compute.networks.get</code>
<code>compute.networks.list</code>
<code>compute.networks.updatePolicy</code>
<code>compute.networks.use</code>
<code>compute.networks.useExternalIp</code>
<code>compute.projects.get</code>
<code>compute.regionOperations.get</code>
<code>compute.regions.get</code>
<code>compute.regions.list</code>
<code>compute.snapshots.create</code>
<code>compute.snapshots.delete</code>
<code>compute.snapshots.get</code>
<code>compute.snapshots.list</code>
<code>compute.snapshots.setLabels</code>
<code>compute.snapshots.useReadOnly</code>
<code>compute.subnetworks.get</code>
<code>compute.subnetworks.list</code>
<code>compute.subnetworks.use</code>
<code>compute.subnetworks.useExternalIp</code>

---



	<p>compute.zoneOperations.get  compute.zones.get  compute.zones.list</p>
Google Kubernetes Engine	<p>container.clusterRoleBindings.list  container.clusterRoles.list  container.configMaps.list  container.controllerRevisions.list  container.cronJobs.list  container.customResourceDefinitions.list  container.daemonSets.list  container.deployments.list  container.endpoints.list  container.jobs.list  container.limitRanges.list  container.networkPolicies.list  container.podTemplates.list  container.replicationControllers.list  container.resourceQuotas.list  container.roleBindings.list  container.roles.list  container.secrets.list  container.statefulSets.list  container.thirdPartyObjects.list</p>
Google Cloud Storage	<p>storage.buckets.create  storage.buckets.createTagBinding  storage.buckets.delete  storage.buckets.get  storage.buckets.getIamPolicy  storage.buckets.list  storage.buckets.listTagBindings  storage.buckets.setIamPolicy  storage.buckets.update  storage.objects.create  storage.objects.delete  storage.objects.get  storage.objects.getIamPolicy  storage.objects.list  storage.objects.setIamPolicy</p>

	storage.objects.update
--	------------------------

## Appendix D

# Deploying a HYCU backup controller

If you are employing the SpinUp functionality in HYCU R-Cloud Hybrid Cloud Edition (HYCU), you can use the R-Cloud web user interface to deploy a HYCU backup controller to AWS or Google Cloud. This enables you to restore your data in the event of a disaster in your HYCU R-Cloud Hybrid Cloud Edition data protection environment.

For details on the supported HYCU R-Cloud Hybrid Cloud Edition infrastructures and how to employ the SpinUp functionality, see HYCU documentation.

Depending on the cloud platform to which you want to deploy the HYCU backup controller, see one of the following sections:

Cloud platform	Instructions
AWS	<a href="#">“Deploying a HYCU backup controller to AWS”</a> below
Google Cloud	<a href="#">“Deploying a HYCU backup controller to Google Cloud”</a> on page 263

## Deploying a HYCU backup controller to AWS

To deploy a HYCU backup controller to AWS, follow the procedure described in this section. After you deploy the HYCU backup controller, you must also configure a port in AWS to be able to access the HYCU web user interface. For details, see [“Accessing the HYCU web user interface”](#) on page 263.

## Prerequisites

- You must own the HYCU and HYCU R-Cloud licenses. For details on how to obtain these licenses, see HYCU documentation.
- You must have the Administrator role assigned.

## Considerations

- The recommended requirements for the HYCU backup controller are 4 vCPU cores and 8 GiB of memory.
- Each HYCU backup controller is by default deployed with the system disk size of 10 GiB and the data disk size of 32 GiB.


### Accessing the HYCU Controller Deployment dialog box

To access the HYCU Controller Deployment dialog box, click 

**Administration**, and then select **HYCU Controller Deployment**.


## Procedure

1. In the HYCU Controller Deployment dialog box, select AWS.
2. From the Source drop-down menu, select the account to which you want to deploy the HYCU backup controller.
3. From the Region drop-down menu, select the geographic region for the HYCU backup controller.

 **Important** Make sure that at least one virtual network is configured in the selected region.

4. From the Zone drop-down menu, select the zone for the HYCU backup controller.
5. Click **Next**.
6. In the Instance name field, enter a name for the HYCU backup controller.
7. In the vCPU cores field, enter the number of virtual CPUs to be assigned to the HYCU backup controller multiplied by the number of cores per virtual CPU. The value that you specify must be a whole number and cannot be higher than 448.
8. In the Memory field, enter the amount of memory (in GiB) to be assigned to the HYCU backup controller. The value that you specify must be a whole number and cannot be higher than 24576.
9. From the Available versions drop-down menu, select the preferred version for the HYCU backup controller. By default, the latest version is selected.

10. From the Instance type drop-down menu, select the instance type.


 **Note** The list of available instance types is based on the number of virtual CPU cores and the amount of memory that you specified. If no instance type corresponds to the specified values, the list is empty and you need to adjust the values in the vCPU and Memory fields.


11. Under Network interfaces, you can view the network interface that will be added to the HYCU backup controller. By default, this is the first network interface from the region and zone that you selected for the HYCU backup controller.


If you want to modify network settings, you can add an additional network interface, edit an existing network interface, or delete a network interface:

#### Modifying network settings


To modify a network interface:


- Click **Add Network Interface** to add a network interface or click  **Edit** next to the network interface that you want to edit, and then follow these steps:
  - a. From the Subnet drop-down menu, select the subnet.
  - b. From the Security groups drop-down menu, select one or more security groups.
  - c. In the Public address type field, select the public IP address for the network interface. You can select among the following options:

Option	Description
None	<p>The network interface does not use a public IP address.</p> <p>This option is preselected if the network interface of the original instance did not use a public IP address.</p>
Auto-assign	<p>The network interface uses an automatically allocated public IP address.</p> <p>This option is preselected if the network interface of the original instance used a public IP address.</p> <p> <b>Note</b> Auto-assign will not work if the Auto-assign public IPv4 address on a subnet option is</p>

	<p>set to No or if more than one network interface is specified.</p>
Elastic IP (Reserved)	The network interface uses an elastic public IP address that was reserved in Amazon EC2 in advance.
Elastic IP (New)	<p>The network interface uses a new elastic public IP address.</p> <p> <b>Note</b> Allocation of the IP address in Amazon EC2 is performed at the very beginning of the deployment. If the allocation fails, the deployment task is terminated without being logged.</p>

- d. In the Private address type field, select the private IP address for the network interface. You can select between the following options:

Option	Description
Auto-assign	<p>The network interface uses an automatically allocated private IP address.</p> <p>This option is selected by default.</p>
Custom	<p>The network interface uses a private IP address that is defined by you.</p> <p> <b>Important</b> Use of this option might result in IP address conflicts.</p>

- e. Click **Add**.
- Click  **Delete** next to the network interface that you want to delete. Keep in mind that you cannot deploy the HYCU backup controller without a network interface.

12. Click **Deploy**.

## Accessing the HYCU web user interface

After you deploy the HYCU backup controller, you must configure a port in AWS to be able to access the HYCU web user interface.

### Procedure

Add rules to the security groups to allow inbound traffic. Specify the following settings:

- Source port ranges: 0.0.0.0/0 (to allow any source port)
- Destination port ranges: 8443

For instructions, see AWS documentation.

You can access the HYCU web user interface by entering the following URL:

```
https://<HYCUBackupControllerPublicIPAddress>:8443
```

On the logon page, enter your logon name and password. You can use the default user name and password for initial access:

User name: **admin**

Password: **admin**

**ⓘ Important** For security purposes, it is highly recommended that you change the default password.

## Deploying a HYCU backup controller to Google Cloud

To deploy a HYCU backup controller to Google Cloud, follow the procedure described in this section. After you deploy the HYCU backup controller, you must also configure a port in Google Cloud to be able to access the HYCU web user interface. For details, see [“Accessing the HYCU web user interface” on page 267](#).

### Prerequisites

- You must own the HYCU and HYCU R-Cloud licenses. For details on how to obtain these licenses, see HYCU documentation.

- You must have the Administrator role assigned.
- The Compute Engine default service account must be enabled for the project to which you plan to deploy the HYCU backup controller.

### Considerations

- The recommended requirements for the HYCU backup controller are 4 vCPU cores and 8 GiB of memory.
- Each HYCU backup controller is by default deployed with the system disk size of 10 GiB and the data disk size of 32 GiB.


#### Accessing the HYCU Controller Deployment dialog box

To access the HYCU Controller Deployment dialog box, click 

**Administration**, and then select **HYCU Controller Deployment**.


### Procedure

1. In the HYCU Controller Deployment dialog box, select Google Cloud, and then click **Next**.
2. From the Source drop-down menu, select the project to which you want to deploy the HYCU backup controller.
3. From the Region drop-down menu, select the geographic region for the HYCU backup controller.
 

 **Important** Make sure that at least one virtual network is configured in the selected region.
4. From the Zone drop-down menu, select the zone for the HYCU backup controller.
5. Click **Next**.
6. In the Instance name field, enter a name for the HYCU backup controller.
7. In the vCPU cores field, enter the number of virtual CPUs to be assigned to the HYCU backup controller multiplied by the number of cores per virtual CPU. The value that you specify must be a whole number and cannot be higher than 1024.
8. In the Memory field, enter the amount of memory (in GiB) to be assigned to the HYCU backup controller. The value that you specify must be a whole number and cannot be higher than 4096.
9. From the Available versions drop-down menu, select the preferred version for the HYCU backup controller. By default, the latest version is selected.



10. From the Instance type drop-down menu, select the instance type.


 **Note** The list of available instance types is based on the number of virtual CPU cores and the amount of memory that you specified. If no instance type corresponds to the specified values, the list is empty and you need to adjust the values in the vCPU and Memory fields.

11. Under Network interfaces, you can view the network interface that will be added to the HYCU backup controller. By default, this is the first network interface from the region and zone that you selected for the HYCU backup controller.


If you want to modify network settings, you can add an additional network interface, edit an existing network interface, or delete a network interface:

#### Modifying network settings

To modify a network interface:

- Click **Add network interface** to add a network interface or click  **Edit** next to the network interface that you want to edit, and then follow these steps:

- a. *Only if you are adding a network interface.* From the Destination Networks drop-down menu, select the destination network.

 **Note** The list of available destination networks includes only the ones within the region you selected for the HYCU backup controller.

- b. In the Public address type field, select the public IP address for the network interface. You can select among the following options:


Option	Description
None	The network interface does not use a public IP address.  This option is preselected if the network interface of the original instance did not use a public IP address.
Ephemeral	The network interface uses an automatically allocated public IP address.  This option is preselected if the network interface of the original instance used a public IP address.

Static (Reserved)	The network interface uses a static public IP address that was reserved in Google Compute Engine in advance.
Static (New)	The network interface uses a static public IP address that is allocated at the time of the deployment. If the allocation fails, the instance is assigned a temporary public IP address. Such a fallback also sets the deployment task status to Done with errors.

- c. In the Private address type field, select the private IP address for the network interface. You can select between the following options:

Option	Description
Ephemeral (Automatic)	The network interface uses an automatically allocated private IP address.  This option is selected by default for the preselected network interfaces.
Ephemeral (Custom)	The network interface uses a private IP address that is defined by you.  <b>ⓘ Important</b> Use of this option might result in IP address conflicts.
Static (Reserved)	<i>Not available for legacy networks.</i> The network interface uses a static private IP address that was reserved in Google Compute Engine in advance.
Static (New)	<i>Not available for legacy networks.</i> The network interface uses a new static private IP address that is defined by you.  <b>📄 Note</b> Allocation of the IP address in Google Compute Engine is performed at the very beginning of the deployment. If the allocation fails, the deployment task is terminated without being logged.

- d. Click **Add**.

- Click  **Delete** next to the network interface that you want to delete. Keep in mind that you cannot deploy the HYCU backup controller without a network interface.

12. Click **Deploy**.

## Accessing the HYCU web user interface

After you deploy the HYCU backup controller, you must configure a port in Google Cloud to be able to access the HYCU web user interface.

### Procedure

Create an inbound security rule to allow traffic. Specify the following settings:

- Source port ranges: 0.0.0.0/0 (to allow any source port)
- Destination port ranges: 8443

For instructions, see Google Cloud documentation.

You can access the HYCU web user interface by entering the following URL:

```
https://<HYCUBackupControllerPublicIPAddress>:8443
```

On the logon page, enter your logon name and password. You can use the default user name and password for initial access:

User name: **admin**

Password: **admin**

**ⓘ Important** For security purposes, it is highly recommended that you change the default password.

# Provide feedback

For any suggestions and comments regarding this product or its documentation, send us an e-mail to:

[info@hycu.com](mailto:info@hycu.com)

We will be glad to hear from you!

