USER GUIDE

HYCU Data Protection as a Service for Google Cloud

December 2022

HYCU®

Legal notices

Copyright notice

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.

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

Contents

Chapter 1

About HYCU for Google Cloud

HYCU Data Protection as a Service for Google Cloud (HYCU for Google Cloud) is a fully managed backup and recovery service for Google Cloud 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.

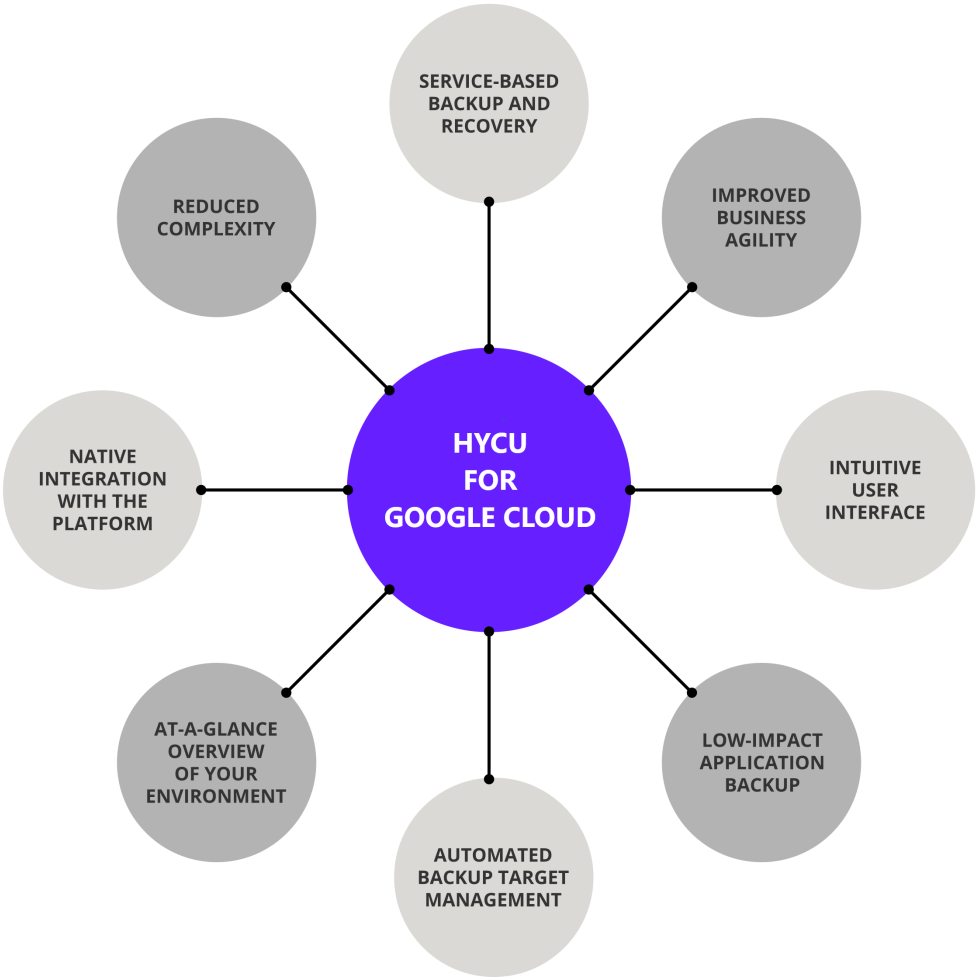

Figure 1–1: Introduction to HYCU for Google Cloud

Key features and benefits

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

- **Protection against data loss**

  Delivers native data protection for Google Compute Engine instances, applications running on Google Compute Engine instances and Google Kubernetes Engine clusters, and Google Cloud Storage buckets, 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 HYCU for Google Cloud, with no deployment and configuration concerns.

- **Application discovery and protection**

  In-built application discovery provides new-found visibility into Google Compute Engine instances and Google Kubernetes Engine 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 Google Cloud projects related to the same service subscription 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.

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

  The HYCU for Google 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 the Google billing system**

  Cost of data protection is billed by Google through existing billing accounts, without requiring you to provide additional billing information.

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

  HYCU Protégé ensures data resilience by using the SpinUp functionality to migrate protected data between the on-premises and Google Cloud infrastructures. In the event of a disaster in your on-premises environment, HYCU Protégé provides disaster recovery of data to cloud. For details on the supported on-premises infrastructures and how to employ HYCU Protégé, see HYCU for Enterprise Clouds documentation.

Data protection environment overview

Before you start protecting data with HYCU for Google Cloud, make yourself familiar with the following terms related to the data protection environment:

Term	Description
HYCU for Google Cloud web user interface	An interface for protecting Google Compute Engine instances, applications running on Google Compute Engine instances or Google Kubernetes Engine clusters, and Google Cloud Storage buckets.
Projects	Containers that include resources managed as a group—instances, buckets, and temporary objects.
Protection sets	Groups that join together Google Cloud projects for which you have access permissions within a subscription.
Instances	Instances to which you can assign policies and for which you therefore provide data protection. Data is always protected at a granular level, allowing you to restore either the entire instances, individual disks, or individual files.
Applications	Applications running on instances or Google Kubernetes Engine clusters to which you can assign a policy and for which you therefore provide data protection.

Term	Description
Buckets	Containers in Google Cloud Storage holding your data to which you can assign a policy and for which you therefore provide data protection. Buckets can also be added to HYCU for Google Cloud as targets for storing backup data.
Targets	Buckets that HYCU for Google Cloud uses for storing backup data. Backup data can also be stored as snapshots.

HYCU for Google Cloud data protection

With the HYCU for Google 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 billing account that you select for your HYCU for Google Cloud subscription defines the scope of data protection—the set of Google Cloud projects that can be protected by HYCU for Google Cloud within the same subscription.

HYCU for Google Cloud enables you to protect instances, applications, 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 HYCU for Google Cloud

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

Task	Instructions
Getting familiar with HYCU for Google Cloud pricing concepts	"Service pricing" below
Subscribing to HYCU for Google Cloud	"Subscribing to the service" on page 16
Signing in to the HYCU for Google Cloud web user interface	"Signing in to HYCU for Google Cloud" on page 17

Service pricing

Because HYCU for Google Cloud utilizes Google Cloud 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:

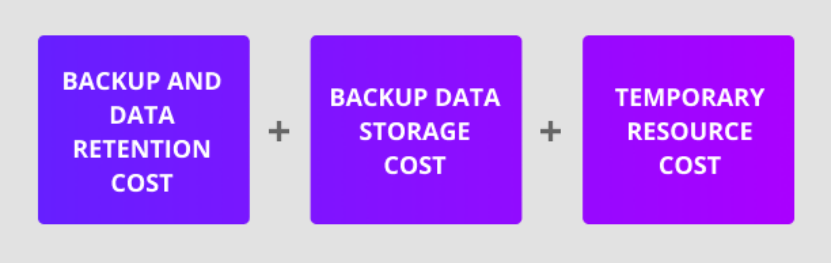

Figure 2–1: Data protection cost

Cost	Details
Backup and data retention	Cost of backing up data and data retention. For details, see "Backup and data retention pricing" on the next page.
Backup data storage	Cost of storing backup data. The following factors are

Cost	Details
	considered: • Target type (a snapshot or a bucket) • Backup frequency • Size of backup data • Backup retention period If you use a bucket as a target, the following is also considered: • Use of copies of backup data • Use of data archives, configured archive tiers and their retention periods • Enabled restore of individual files or folders
Temporary resources	Cost of temporary resources that HYCU for Google Cloud creates in Google Cloud when performing the following tasks: • Instance rediscovery after assigning a credential group • Instance rediscovery after selecting the Enable restore of individual files option • Backup of instances • Backup of applications • Backup of buckets • Restore of instances or entire instance disks • Restore of individual files or folders • Restore of applications • Restore of buckets

A HYCU for Google 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 Google as usual.

For more details on pricing, see Google Cloud Marketplace.

Backup and data retention pricing

The HYCU for Google 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 HYCU for Google Cloud. For details on the subscription plans, see "HYCU for Google Cloud subscription plans" on the next page.

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, application, or bucket. HYCU for Google 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 25.

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 Google 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.

HYCU for Google Cloud subscription plans

HYCU for Google 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.

  For details on the pay-as-you-go subscription plan, see Google Cloud Marketplace.
- Annual subscription plans

You can choose among different annual subscription plans with token-based pricing.

For more information, contact your HYCU sales representative.

Subscribing to the service

You subscribe to HYCU for Google Cloud online from the Google Cloud Marketplace and HYCU then automatically activates the service for you. This is usually done by one user for an entire organization.

Prerequisites

- You have a Google Account.
- Your billing account meets the following requirements:
  - Your Google Account is granted the Billing Account Administrator (`roles/billing.admin`) role on the billing account. This role is required for purchasing solutions on the Google Cloud Marketplace.
  - The billing account has at least one linked Google Cloud project that your Google Account has access to.
- You are signed in to Google, and your currently selected project in the Google Cloud Console is linked to the billing account.

For details, see Google and Google Cloud documentation.

Considerations

- You cannot change the billing account for an active HYCU for Google Cloud subscription. If you chose a wrong billing account and want to change it, contact HYCU Customer Support.
- If you violate the terms of use of HYCU for Google 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 TO HYCU**.
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 HYCU Data Protection as a Service for Google Cloud webpage, enter the required information.
7. Verify that the provided information is correct, and then click **Sign up with Google**.

8. In the Choose an account page, specify or select the email address of the same Google Account with which you were already signed in to Google. If needed, enter the corresponding password. Click **Next**.
9. In the Apps with access to your account webpage, review the listed actions for which HYCU for Google Cloud must be granted permissions in scope of access to your Google Account. If you consent to grant the permissions, click **Allow**.

   HYCU for Google Cloud requires these permissions for determining which Google Cloud projects are linked to the billing account of the subscription. The permission for sending you notifications related to your subscription by email is implied.

   > Note You can revoke permissions for HYCU for Google Cloud to access your Google Account at any time. For instructions, see "Preventing account access" on page 138.

10. On the HYCU Data Protection as a Service for Google Cloud webpage, do one of the following:
    - Review the only billing account.
    - From the drop-down menu, select the billing account that you chose for your subscription.
11. Under Projects, review the list of the linked Google Cloud projects. If the billing account is correct, click **Submit**.

    Provisioning of HYCU for Google Cloud for your subscription starts.
12. Click **Close & Return to Google**.
13. On the New HYCU subscription page, click the left arrow icon to return to the solution's main page on the Google Cloud Marketplace.
14. Click **Refresh to check status** to track the provisioning progress until the following message appears under the Your HYCU subscription heading:

    `You activated your HYCU service on <Date>.`

    > Tip On the Google Cloud Marketplace page, under the Get started with HYCU heading, click **HYCU, Inc** to go to the sign-in page of the HYCU for Google Cloud web user interface.

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

Signing in to HYCU for Google Cloud

After successfully subscribing to HYCU for Google Cloud, you can sign in to the HYCU for Google Cloud web user interface.

Prerequisites

- Your Google Account has at least the Viewer (`roles/viewer`) role granted on at least one Google Cloud project that is linked to the billing account of a HYCU for Google Cloud subscription.
- The Google Cloud projects with instances, applications, and buckets that you plan to protect are linked to the assigned billing account.
- In Google Cloud, the Compute Engine default service account must be present on the project that you plan to protect. If this service account is not available, you must set up an alternative service account in the following format: `hycu-<projectNumber>@<projectId>.iam.gserviceaccount.com`
- In Google Compute Engine, your Google Account has 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 has 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 is enabled on the Google Cloud projects with instances, applications, and buckets that you plan to protect. For instructions, see Google Cloud documentation.
- You are using a supported web browser. For a list of supported web browsers, see the *HYCU for Google Cloud Compatibility Matrix*.

For details, see Google and Google Cloud documentation.

Procedure

1. Open a web browser and go to the HYCU Data Protection as a Service for Google Cloud webpage.
2. Click **Next**.
3. On the sign-in webpage, click **Sign in with Google**.
4. Specify or select the email address of your Google Account. If you are not signed in with that account yet, enter the password, and then click **Next**.
5. Review the listed actions for which HYCU for Google Cloud must be granted permissions in scope of access to your Google Account. If you consent to grant the permissions, click **Allow**.

   HYCU for Google Cloud requires these permissions for determining which Google Cloud projects are linked to the billing account of the subscription, for accessing Google Compute Engine instances and their disks during backup and restore, for accessing Google Cloud Storage buckets to store your backup data, and so on. The permission for sending you notifications related to your subscription by email is implied.

> **Note** You can revoke permissions for HYCU for Google Cloud to access your Google Account any time. For instructions, see "Preventing account access" on page 138.

After you sign in to the HYCU for Google 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 HYCU for Google 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 HYCU for Google 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 121.

Tasks

Establishing a data protection environment involves the following tasks:

Task	Instructions
1. Enable the HYCU Managed Service Account (HMSA).	"Enabling the HYCU Managed Service Account" on the next page
2. *Only if you plan to use multiple protection sets.* Configure a protection set and select it.	"Managing protection sets" on page 123 and "Selecting a HYCU for Google Cloud protection set" on page 22
3. *Only if you plan to use manually created targets.* Add Google Cloud Storage buckets to HYCU for Google Cloud as targets.	"Setting up targets" on page 23
4. Decide for predefined policies or create custom ones.	"Defining your backup strategy" on page 25
5. *Required only in special data protection scenarios.* Configure credential groups and assign them to instances.	"Enabling access to data" on page 36

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

Enabling the HYCU Managed Service Account

The HYCU Managed Service Account (HMSA) is a special type of account that is designed specifically for HYCU for Google 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.

To enable the HMSA, you must perform all the steps in the HYCU Managed Service Account configuration wizard that appears after you sign in to HYCU for Google Cloud or that you can access by clicking **Enable** in the Subscription Information dialog box. To open the Subscription Information dialog box, click **?** in the toolbar, and then select **Subscription Information**.

Prerequisite

You must have the Administrator role assigned.

Considerations

- When selecting the projects in the HYCU Managed Service Account configuration wizard, keep in mind that the project list includes all the HYCU for Google Cloud projects that are linked to your billing account. The protected projects are preselected and, if required, you can add additional projects by selecting the check box before the name of each project.
- After you enable the HMSA, you can no longer access the HYCU Managed Service Account configuration wizard, and on each project that you plan to protect, you must manually grant the following permissions to the HMSA in Google Cloud:
  - 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.

Procedure

The HYCU Managed Service Account configuration wizard guides you through all the required steps of enabling the HMSA. You first select projects that are already protected or that you plan to protect with HYCU for Google Cloud, and then allow the HMSA to access all the selected projects, which you can do either manually (by granting the required permissions to the HMSA on each project) or automatically (by using Google Cloud Shell).

Selecting a HYCU for Google Cloud protection set

An environment for which HYCU for Google Cloud provides data protection consists of one or more protection sets that join together Google Cloud projects. When you subscribe to HYCU for Google Cloud, a default protection set is created automatically (represented by the icon) and all the projects that are linked to the billing account of your HYCU for Google Cloud subscription are included in it.

Depending on your business needs, you can create additional protection sets and distribute your projects 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 123.

If no multiple protection sets are 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.

Considerations

- Regardless of your protection set configuration, you can see only projects linked to the billing account that was selected when subscribing to HYCU for Google Cloud and projects that you can access with your user account.
- *Only if multiple protection sets are available in your data protection environment.* The currently selected protection set has the icon next to it.

Procedure

1. On the toolbar, click next to the name of the selected protection set.
2. In the Protection Set Picker dialog box, from the Subscription — billing account drop-down menu, select the HYCU for Google Cloud subscription that contains the protection set for which you want to perform data protection tasks.
3. From the list of available protection sets, select the scope of your data protection by selecting the preferred protection set.

   > Tip You can also search for a protection set by entering its name or the name of an included project and then pressing **Enter** in the Protection set search field.

4. Click **Select**.

The HYCU for Google Cloud web user interface switches the context to the selected scope of data protection. The protection set that you selected last is remembered for the next time you sign in.

Setting up targets

Targets are locations where backup data is stored. HYCU for Google Cloud allows you to define either a bucket or a snapshot as a location for storing your data.

Target	Description
Bucket	Backup data is stored in Google Cloud Storage buckets that you create yourself or HYCU for Google Cloud creates for you automatically: • Manually created targets You can create your own buckets in Google Cloud Storage and add them to HYCU for Google Cloud as targets. For instructions, see "Adding a bucket to HYCU for Google Cloud as a target" on the next page. • Automatically created targets *Applicable only if you are protecting instances or Google Kubernetes Engine applications.* HYCU for Google Cloud creates Google Cloud Storage buckets automatically while backing up data and uses them as targets. For increasing restore speed and minimizing costs, these targets are created in the same Google Cloud project and at the same location as the instances you are backing up or the clusters on which the applications you are backing up are deployed. The same target is used for storing the backup data of multiple instances and applications where possible. You can use these targets also for storing your data (for example, for individual files that you restore). For the target naming conventions, see "Objects created by the service" on page 140. **Caution** Never delete any targets used by HYCU for Google Cloud because this may result in data loss. Additionally, within targets, ensure that the `hycu/backups/` folders are always kept intact.
Snapshot	*Available only if you are protecting instances or Google Kubernetes Engine applications using persistent volumes.* Backup data is stored as a snapshot in a Google Cloud project that contains the instances you want to protect or the clusters on which the applications you want to protect are deployed. **Note** If snapshots created by HYCU for Google Cloud are deleted from Google Cloud, you will not be able to restore backup data from this location. However, you can still restore your data

Target	Description
	from targets if copies of backup data or data archives exist. For the snapshot naming conventions, see "Objects created by the service" on page 140.

Adding a bucket to HYCU for Google Cloud as a target

Prerequisites

- Your HYCU Managed Service Account (HMSA) must have access to the bucket.
- *Only if you plan to select a specific service account for performing all operations on the target.* The service account must have access to at least one of the projects linked to the selected billing account and the bucket.
- If a retention policy or the default event-based hold property is enabled on the bucket, a service account must be imported to invoke operations on the target. For instructions on how to import service accounts, see "Importing service accounts" on page 128.

Limitations

- Publicly available buckets cannot be added as targets.
- *Only if you plan to protect buckets.* Restoring the original access control list is supported only if a service account is used for invoking operations on the target. For instructions on how to import service accounts, see "Importing service accounts" on page 128.

Considerations

- You can set up the same target in multiple protection sets.
- The exclude policy is automatically assigned to the bucket that is added to HYCU for Google Cloud as a target. It is highly recommended that you do not change this default configuration.
- *Only if you plan to select a specific service account 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.

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**. The Add Target dialog box opens.
2. In the Target field, enter the name of the bucket that you want to add to HYCU for Google Cloud as a target.
3. In the Size field, specify the amount of storage space that should be used for storing backup data (in MiB, GiB, or TiB).

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

4. From the Cloud account drop-down menu, select the service account that you want to be used for performing all operations on this target. If None is selected, the HMSA will be used.
5. 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 111.

Defining your backup strategy

HYCU for Google 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 "Taking advantage of predefined policies" on the next page.
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 "Creating custom policies" 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 34.

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.

HYCU for Google 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 instances, applications, or buckets 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 23.

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 instances, Google Kubernetes Engine applications, 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 "Creating backup windows" on page 30.
Copy[a]	Create a copy of backup data.
Archiving[a]	Preserve your data for future reference. For details, see "Creating data archives" on page 32.
Labels	Set up automatic policy assignment based on the labels or tags

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

[a] *For GKE applications:* This policy option is available only for applications using persistent volumes.

Prerequisites

- *Only if you plan to select a manually created target.* A bucket must be added to HYCU for Google Cloud as a target. For instructions, see "Setting up targets" on page 23.
- *Only if you plan to enable the Backup Window policy option.* A backup window must exist for the selected HYCU for Google Cloud protection set. For instructions, see "Creating backup windows" on page 30.
- *Only if you plan to enable the Archiving policy option.* A data archive must exist for the selected HYCU for Google Cloud protection set. For instructions, see "Creating data archives" on page 32.
- *Only if you plan to enable the Labels policy option.*
  - 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 HYCU for Google Cloud must be added to instances in Google Compute Engine as labels (preferred) or custom metadata tags, to GKE applications in Google Kubernetes Engine as metadata labels, or to buckets in Google Cloud Storage as bucket labels.

    For instructions on how to do this, see Google Cloud or Kubernetes documentation.

Considerations

- HYCU for Google 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 that 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 HYCU for Google 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 instance in Google Compute Engine, to the same application in Google Kubernetes Engine, or to the same bucket in Google Cloud Storage. For more information on the `hycu-policy` tag, see "Setting up automatic policy assignment" on page 34.
- *Only if you plan to store backup data on a target.* Backup and restore speed depends on the region of the chosen target and the regions of the instances or Kubernetes clusters with your GKE applications. The optimum speed is achieved when the target and the instances or clusters reside in the same region.

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
Backup Window	✓	×	✓
Copy[a]	✓	×	✓
Archiving[a]	✓	×	✓
Labels	✓	×	✓

[a] *For GKE applications:* This policy option is available only for applications using persistent volumes.

4. In the Backup section, do the following:
   a. In the Backup every fields, 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.

b. In the Retention fields, set a retention period (in months, weeks, or days) for the backup data.

c. Select one of the following backup target types:

- *Applicable only if protecting instances or GKE applications using persistent volumes.* **Snapshot**

  Under Snapshot Location, select **Regional** or **Multi-regional**.

  > Example
  >
  > If your instance resides in the `us-central1-a` 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 `us-central1` region.

- **Target**

  From the Target drop-down menu, select the target that you want to use for storing data.

  If you select the **Automatically selected** option, HYCU for Google Cloud creates a bucket in the region of the instance or the Kubernetes cluster and uses it as a target for storing the data. If an automatically created bucket already exists, it is used instead.

  > ⚠ Important Automatically created targets can be selected only if you plan to protect instance data or GKE application data (and not SAP HANA application or bucket data).

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

Policy option	Instructions
Backup Window	In the Backup Window section, from the Backup window drop-down menu, select a backup window for backup tasks. 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.
Copy[a]	In the Copy section, do the following: a. Set a retention period (in months, weeks, or days) for the copy of backup data. b. From the Target drop-down menu, select a target that you

Policy option	Instructions
	want to use for storing data. If you want the target to be selected automatically, make sure the **Automatically selected** option is selected. In this case, HYCU for Google Cloud creates a bucket in the region of the instance or the Kubernetes cluster and uses it as a target for storing the data. If an automatically created bucket already exists, it is used instead. If you want to select a manually created target, make sure that this target is different from the one you selected for the backup. ⚠ Important Automatically created targets can be selected only if you plan to protect instance data or GKE application data (and not SAP HANA application or bucket data).
Archiving[a]	In the Archiving section, from the Data archive drop-down menu, select a data archive.
Labels	In the Labels section, enter a label key and value, and then click **Add**. If required, repeat the action as appropriate. For details on automatic policy assignment, see "Setting up automatic policy assignment" on page 34.

[a] *For GKE applications:* This policy option is available only for applications using persistent volumes.

6. Click **Save**.

The policy is created and added to the list of policies. For details on managing policies, see "Managing policies" on page 114.

Creating backup windows

HYCU for Google 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 instances, GKE applications, 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 ×.
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 26.
- 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.

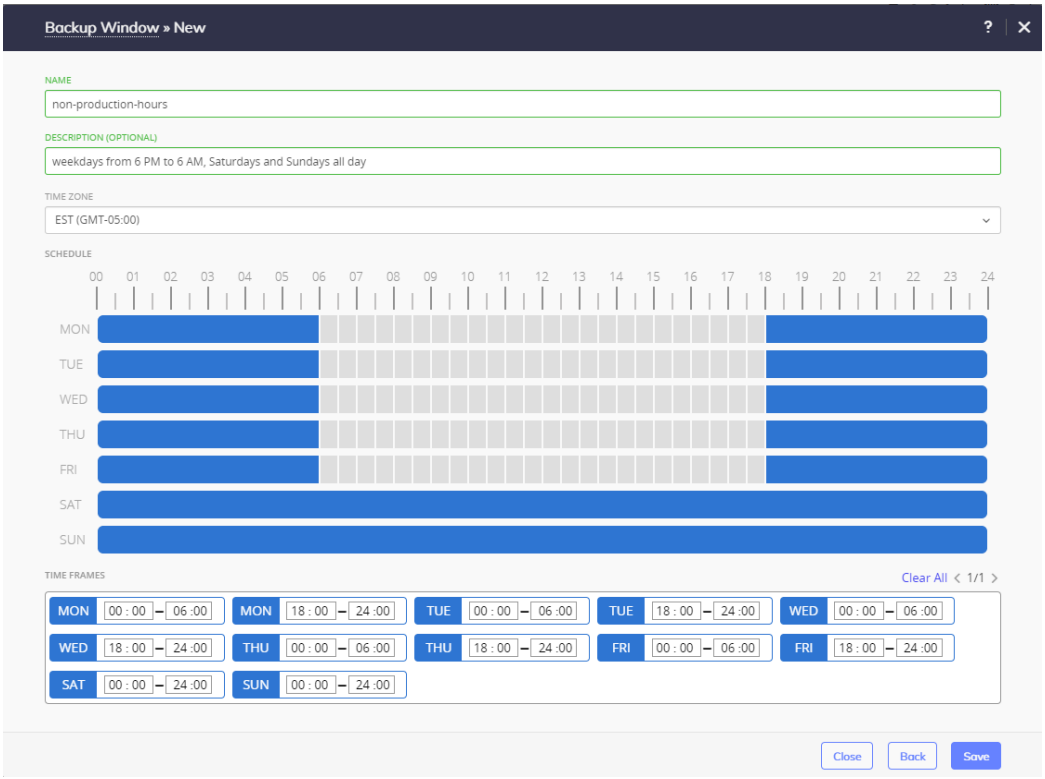

In this case, the backup tasks can be run every 24 hours at any point of time within the specified time frames.

Creating data archives

HYCU for Google 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. You have created a bucket and added it to targets of the selected protection set in HYCU for Google Cloud.

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:

Daily	Allows you to create a daily archive of data.
Weekly	Allows you to create a weekly archive of data.
Monthly	Allows you to create a monthly archive of data.
Yearly	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 a target that you want to use for storing the data archive.

      If you select the **Automatically selected** option, HYCU for Google Cloud creates a bucket in the region of the instance or the Kubernetes cluster and uses it as a target for storing the data. If an automatically created bucket already exists, it is used instead.

   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, a storage class is automatically selected depending on the specified retention.

      For details on storage classes, see Google Cloud 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 if an archiving task is in progress on that target.

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 26.
- 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" below.

> ### 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**

   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 instances, Google Kubernetes Engine (GKE) applications, or buckets by using one of the following methods:

Resources	Method 1	Method 2
Instances	By first adding labels (preferred) or custom metadata tags to instances in Google Compute Engine, and then specifying the corresponding label names and values in HYCU for Google Cloud policies. For details, see "Creating custom policies" on page 26.	By adding the `hycu-policy` tag to instances in Google Compute Engine, applications in Google Kubernetes Engine, or buckets in Google Cloud Storage. Use the following name/value pair: Name: `hycu-policy` Value: <*PolicyName*> In this case, <*PolicyName*> is the name of a HYCU for Google Cloud policy (for example, Gold).
GKE applications	By first adding metadata labels to applications in Google Kubernetes Engine, and then specifying the corresponding label names and values in HYCU for Google Cloud policies. For details, see "Creating custom policies" on page 26.	
Buckets	By first adding bucket labels to buckets in Google Cloud Storage, and then specifying the corresponding label names and values in HYCU for Google Cloud policies. For details, see "Creating custom policies" on page 26.	

The corresponding policies are automatically assigned to the instances, GKE applications, or buckets during the next instance, application, or bucket synchronization in HYCU for Google Cloud.

Prerequisites

- All relevant prerequisites that apply also for manual policy assignment are fulfilled. For details, see "Backing up instances" on page 45.
- *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` is 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 or the tag added to the preferred instance, GKE application, or bucket defines which policy is assigned to it, even if the same instance, application, or bucket already has an assigned policy.
- If you want to assign a new policy to an instance, a GKE application, or a bucket for which automatic policy assignment has been set up, do one of the following:
  - Define new tags or labels as described in this section.
  - Assign the policy to the instance, the application, or the bucket as described in "Backing up instances" on page 45, "Backing up applications" on page 74, or

"Backing up buckets" on page 85. In this case, the manually assigned policy will not be overridden by the automatically assigned one again.

Enabling access to data

HYCU for Google Cloud uses the following default parameters to connect to instances:

Guest OS	Authority user name	Network service protocol	Port	Transport protocol
Linux	<*UserName*>[a]	SSH	22	N/A
Windows[b]	hycu	WinRM	5986	HTTPS

[a] The email address of the authority that is running the task in HYCU for Google Cloud is <*UserName*>@<*DomainName*>.
[b] HYCU for Google Cloud automatically configures a credential group named auto-<*InstanceName*> and assigns it to the instance.

The default connection parameters are suitable for the majority of data protection scenarios. However, in the following cases, you must manually enable access to the instances by assigning credential groups to them in HYCU for Google Cloud:

Guest OS	Data protection scenario
any	• You plan to restore individual files using a user account that you specify. • You plan to use a specified user account for the restore, either to reuse an already existing user account or to comply with policies that impose restrictions on the utilized user names and passwords.
Linux	• You plan to protect SAP HANA applications. • You plan to use pre-snapshot or post-snapshot scripts and run them with a user account that you specify. • The SSH server is configured to use a non-default TCP port. • The SSH server is configured to use public key authentication. • OS Login is enabled on the instance in Google Compute Engine. For more information on OS Login as the access method, see Google Cloud documentation.
Windows	• You plan to use pre-snapshot or post-snapshot scripts. • The WinRM server is configured to use the HTTP transport protocol or a non-default TCP port.

Configuring and assigning credential groups manually

Prerequisites

- A user account with sufficient privileges is configured within each instance. For details on how to do this, see Google Cloud documentation.
- *For Linux instances:*
  - *Only if the Authentication option in HYCU for Google Cloud is set to either Password authentication or Public key authentication:* Ensure the following within the instance:
    - The specified user account is a member of the `sudo` user group.
    - The following line is included in the `/etc/sudoers` file:

      ```
      <UserName> ALL=(ALL) NOPASSWD: /bin/lsblk, /bin/ls, /bin/mkdir,
      /bin/mv, /bin/umount, /bin/cp, /bin/rm, /bin/mount
      ```
  - *Only if you want HYCU for Google Cloud to access the instance by using a specific user account with password authentication.* The SSH server is 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**. Alternatively, in the Dashboard panel, click the **Instances** widget.

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:

<table>
<tr><th>Protocol option</th><th>Instructions</th></tr>
<tr><td>Automatic</td><td>Select this option if you want HYCU for Google 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.

Use the following format for the user name:

• Linux: <LocalOrDomainUserName>
• Windows: <LocalUserName>, <Domain>\<DomainUserName>, <DomainUserName>@<Domain></td></tr>
<tr><td>SSH</td><td>Select this option if you want to use the SSH protocol for accessing the instance, and then do the following:

a. In the Port field, enter the SSH server port number.
b. From the Authentication drop-down menu, select the type of authentication you want to be used, and then provide the required information:

<table><tr><td>Automatic</td><td>This option provides the same behavior as if no credential group is assigned to the instance, but adds the possibility to adjust the port number used when accessing to the instance.

⚠ Important Do not select this option if OS Login is enabled on your instance.</td></tr><tr><td>Password authentication</td><td>Enter the user name (in the <LocalOrDomainUserName> format) and password of a user account that has required permissions to access the instance.</td></tr><tr><td>Public key authentication</td><td>Do the following:

i. Enter the user name (in the <LocalOrDomainUserName> format) and password of a user account that has required permissions to access the instance.</td></tr></table></td></tr>
</table>

<table>
<tr><th>Protocol option</th><th>Instructions</th></tr>
<tr><td></td><td>
<table>
<tr><td></td><td>ii. Click Browse. Browse for and then select the file with the private key, and click Open.

For information on how to obtain the private key, see Google Cloud documentation.

iii. Only if the private key is encrypted. Enter the private key passphrase.

⚠ Important This selection is mandatory for utilization of the OS Login access method in Google Compute Engine connection to an instance in which the SSH server is configured to use public key authentication. For more information, see Google Cloud documentation.</td></tr>
</table>
</td></tr>
<tr><td>WinRM</td><td>Select this option to use the WinRM protocol for instance access and to enable the credential group adjustment for the actual WinRM server configuration.

a. From the Transport drop-down menu, select the transport protocol of the WinRM server in the instance.

b. In the Port field, enter the WinRM server port number.

c. Enter the user name (in the <code><<i>LocalOrDomainUserName</i>></code> format) and password of a user account that has required permissions to access the instance.</td></tr>
</table>

6. Click **Save**.
7. Click **Assign**.

The name of the assigned credential group appears in the Credential group column of the Instances panel. HYCU for Google Cloud performs instance and application discovery after you assign the credentials to the instance and 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**.

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**).

Chapter 4

Protecting instances

HYCU for Google Cloud enables you to protect your instance data with fast and reliable backup and restore operations.

Prerequisites

- The HYCU Managed Service Account (HMSA) must have the following roles granted on the projects with the instances that you plan to protect:
  - Compute Admin (`roles/compute.admin`)
  - Service Account User (`roles/iam.serviceAccountUser`)
  - Storage Admin (`roles/storage.admin`)

  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

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

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 HYCU for Google 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, HYCU for Google 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 116.

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 121.
- HYCU for Google 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.
- Data in instance backup images, copies of backup images, and data archives that HYCU for Google 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 applications that store their data on a single disk, and 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 more information about Windows VSS snapshot prerequisites, see "Backing up instances" on page 45.

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

- "Configuring instance backup options" below
- "Backing up instances" on page 45
- "Restoring instances" on page 47
- "Restoring individual files or folders" on page 60

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.

Backup option	Description
Running pre/post scripts	You can use the pre-snapshot and post-snapshot scripts to perform necessary actions before and 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.
Excluding disks from the backup	You can specify any disk to be excluded from the instance backup.

Backup option	Description
Allowing the restore of individual files	You can allow the restore of individual files if your data protection needs require that only individual files are restored, and not the entire instance. As an alternative to allowing the restore of individual files by using the Configuration option described in this procedure, you can also tag an instance in Google Cloud, and by doing so, instruct HYCU for Google Cloud to allow it automatically. For details, see "Allowing the restore of files by tagging the instance in Google Cloud" on page 45.
Specifying the temporary instance location and subnet	You can specify the region, the zone, and the subnet where you want HYCU for Google Cloud to create a temporary instance during the backup. By default, the temporary instance is created in the project of the original instance.

Prerequisites

- *Only if you plan to use pre-snapshot and post-snapshot scripts.*
  - Access to the instance file system must be enabled. For instructions, see "Enabling access to data" on page 36.
  - 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.
- *Only if 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.

Considerations

- *Only if you plan to use pre-snapshot and post-snapshot scripts.* The scripts are run from the home directory of the user account that HYCU for Google Cloud uses for running the scripts.

  Depending on the operating system on the instance, the following user accounts are used:
  - GNU/Linux:
    - *The instance has no credential group assigned in HYCU for Google Cloud:* The HYCU Managed Service Account (HMSA).
    - *The instance has a credential group assigned:* The user account specified in the credential group.

    - Microsoft Windows: The user account that is assigned to the instance in HYCU for Google Cloud by means of a credential group.
- *Only if you plan to exclude the boot disk from the backup.* When restoring the instance whose boot disk was excluded from the backup, the Restore Instance and Clone Instance options are not available.

> Accessing the Instances panel
>
> To access the Instances panel, in the navigation pane, click **Instances**. Alternatively, in the Dashboard panel, click the **Instances** widget.

Procedure

1. In the Instances panel, select the instance for which you want to configure backup options.

   > Tip To configure the same backup options for multiple instances at once, select the preferred instances.
   >
   > Keep in mind that you cannot configure disk exclusion from backup for multiple instances at the same time. You can edit other backup options only if they currently have the same settings for all selected instances.

2. Click **Configuration**. The Instance Configuration dialog box opens.
3. Depending on what you want to do, perform the required action:
   - Run the pre-snapshot and post-snapshot scripts:

     On the Pre/post scripts tab, do the following:
     - In the Pre-snapshot script field, enter the script that HYCU for Google Cloud runs before it creates a snapshot of the instance. The following are examples of the scripts:
       - GNU/Linux: `bash /home/<UserName>/freeze_db.sh`
       - Microsoft Windows: `%USERPROFILE%\quiesce_db.bat`
     - In the Post-snapshot script field, enter the script that HYCU for Google Cloud runs after it creates a snapshot of the instance. The following are examples of the scripts:
       - GNU/Linux: `bash /home/<UserName>/thaw_db.sh`
       - Microsoft Windows: `%USERPROFILE%\resume_db.bat`
   - Exclude disks from the backup:

     On the Exclude from backup tab, select the disks that you want to exclude from the backup.
   - Allow the restore of individual files or folders:

     On the Restore individual files tab, enable the **Enable restore of individual files** switch.

- Specify the region, the zone, and the subnet where you want HYCU for Google Cloud to create a temporary instance:

  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.

4. Click **Save**.

Allowing the restore of files by tagging the instance in Google Cloud

As an alternative to allowing the restore of individual instance files in HYCU for Google Cloud, you can add the `hycu-enable-flr` tag as the label or the custom metadata tag to the instance in Google Cloud, and by doing so, instruct HYCU for Google Cloud to allow it automatically.

Procedure

In Google Cloud, use the following name/value pair for the instance:

Name	Value
`hycu-enable-flr`	True[a]

[a] By setting the value to `False`, you disallow the restore of individual files for the specific instance.

If the instance has credentials assigned, HYCU for Google Cloud automatically allows the restore of its individual files. Otherwise, you must assign the credentials to the instance. For details on how to do this, see "Enabling access to data" on page 36.

Backing up instances

With HYCU for Google Cloud, you can back up your instances in a fast and efficient way.

Prerequisites when planning to restore individual files or folders

- *For instances running Microsoft Windows:* The following prerequisites are fulfilled:
  - In the Google Cloud Console, there is a network firewall rule applied to the instances—either to the entire network or to individual instances through the use of network tags. For each instance, the rule must allow ingress network traffic through a TCP port configured for WinRM communication (by default, 5986) from the entire subnetwork that the instance belongs to.

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.

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

- On the instances that you plan to protect, WinRM is configured to use a TCP port (by default, 5986) for HTTPS or HTTP communication. The firewall is configured to enable inbound network traffic through this port.

- The restore of individual files or folders is enabled for the instance. For instructions on how to enable the restore of individual files or folders, see "Configuring instance backup options" on page 42.
- *Only if you want to use custom credentials for the restore.* The correct credential group is assigned to the original instance, and the corresponding credentials belong to a user account with sufficient privileges. For instructions on how to assign access credentials, see "Enabling access to data" on page 36.

Prerequisites when planning to use pre-snapshot or post-snapshot scripts

- In the Google Cloud Console, there is a network firewall rule applied to the instances—either to the entire network or to individual instances through the use of network tags. For each target instance, the rule must allow ingress network traffic through a specific TCP port from the entire subnetwork that the instance belongs to. The port number depends on the guest operating system of the instance and connection server configuration:
  - GNU/Linux: TCP port 22 (or a different port if configured for SSH communication)
  - Microsoft Windows: TCP port 5986 (or a different port if configured for WinRM communication)

  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.

  For instructions on how to configure and apply the network firewall rule, see Google Cloud documentation.
- *For instances running GNU/Linux:* On the instances that you plan to protect, an SSH server is installed and configured to use a TCP port (by default, 22) for SSH communication. The firewall is configured to enable inbound network traffic through this port.
- *For instances running Microsoft Windows:* On the instances that you plan to protect, WinRM is configured to use a TCP port (by default, 5986) for HTTPS or HTTP communication. The firewall is configured to enable inbound network traffic through this port.
- *For instances running Microsoft Windows, and for instances running GNU/Linux with non-default configuration of SSH server or if you want to use custom user accounts for*

running the script: The correct credential group is assigned to the instance and the corresponding credentials belong to a user account with sufficient privileges. For instructions on how to assign access credentials, see "Enabling access to data" on page 36.

Consideration

When backing up an instance with multiple disks, HYCU for Google 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.

> **Accessing the Instances panel**
>
> To access the Instances panel, in the navigation pane, click **Instances**. Alternatively, in the Dashboard panel, click the **Instances** widget.

Procedure

1. Select the instances that you want to back up. You can update the instance list by clicking **Synchronize**. In a protection set with a large number of projects, 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 in panels" on page 105.
2. Click **Policies**. The Policies dialog box opens.
3. From the list of available policies, select the desired policy.
4. Click **Assign** to assign the policy to the selected instances.

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 image 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 115.

Restoring instances

HYCU for Google 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.

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 36.

- A script must be available in an accessible folder.
- The user account must have permissions to run a script on the instance.

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 HYCU for Google Cloud uses for running the scripts.

  Depending on the operating system on the instance, the following user accounts are used:
  - GNU/Linux:
    - *The instance has no credential group assigned in HYCU for Google Cloud:* The HYCU Managed Service Account (HMSA).
    - *The instance has a credential group assigned:* The user account specified in the credential group.
  - Microsoft Windows: The user account that is assigned to the instance in HYCU for Google Cloud by means of a credential group.

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

Restore option	Description	Instructions
Restore Instance	Enables you to restore an instance and its disks to the original location with the same settings.	"Restoring an instance" on the next page
Clone Instance	Enables you to restore an instance and its disks by creating a clone of the instance.	"Cloning an instance" on page 50
Restore Disks	Enables you to restore instance disks and attach them to the same instance.	"Restoring disks" on page 54
Clone Disks	Enables you to restore instance disks by creating their clones and attaching them to the same or a different instance.	"Cloning disks" on page 55
Export Disks	Enables you to restore instance disks to the same or a different projects or zone without attaching them to an instance.	"Exporting disks" on page 57

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

Restore option	Description	Instructions
Restore Instances	Enables you to restore multiple instances	"Restoring multiple

Restore option	Description	Instructions
	by creating clones of the instances.	instances in a single session" on page 58
Restore Disks	Enables you to restore multiple instance disks on multiple instances at once.	"Restoring multiple disks in a single session" on page 59

Accessing the Instances panel

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

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.

Considerations

- Any data changes after the last successful backup are not protected and therefore cannot be restored.
- If you are restoring an instance with three or more disks, HYCU for Google Cloud by default performs the parallel restore, which speeds up the entire instance restore process. For instances with fewer than three disks, the sequential restore is performed.

Procedure

1. In the Instances panel, click the instance that you want to restore to open the Details section.

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

2. In the Details section that appears at the bottom of the screen, select the desired restore point.
3. Click **Restore Instance**. The Restore Options dialog box opens.
4. Select **Restore Instance**, and then click **Next**.
5. 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)**

6. 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.

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

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

8. 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 project, region, and zone, and the instance network configuration.

Limitation

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

Considerations

- If you are restoring an instance with three or more disks, HYCU for Google Cloud by default performs the parallel restore, which speeds up the entire instance restore process. For instances with fewer than three disks, the sequential restore is performed.
- *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 Details section.

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

2. In the Details section that appears at the bottom of the screen, select the desired restore point.
3. Click **Restore Instance**. The Restore Options dialog box opens.
4. Select **Clone Instance**, and then click **Next**.
5. In the New instance name field, specify a new name for the instance.
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. *Optional.* In the Post-restore script field, enter the path to the script or a command that HYCU for Google 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.
8. From the Target project drop-down menu, select the project to which you want to restore the instance. The original project of the instance is preselected. You can choose from projects that belong to the currently selected protection set and that your user account can access.
9. From the Target region and Target 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.
10. 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 disks, do the following:
       i. Click **Edit Disk**.
       ii. *Only if you do not want HYCU for Google 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, 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 Replica zones drop-down menu, select to which two zones 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. If you want to add labels to the restored disk, click **Add Tags**, 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 ✕ next to it.

vi. Click **Save**.

11. 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:

- 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
- Network type: Subnet for VPC networks and shared VPC networks, Legacy for legacy networks

For each configured network interface, you can separately adjust its external and internal IP address types. By default, the external 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 Target network drop-down menu, select the target network.
  b. *Only if you are adding a network interface.* From the Virtual network drop-down menu, select the virtual network for the network interface.

  > Note The list of available virtual networks includes only the ones within the region you selected for the restored instance.

  c. In the External address type field, select the external IP address for the network interface. You can select among the following options:

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

d. In the Internal address type field, select the internal IP address for the network interface. You can select between the following options:

Option	Description
Ephemeral (Automatic)	The network interface uses an automatically allocated internal IP address. This option is selected by default for the preselected network interfaces.
Ephemeral (Custom)	The network interface uses an internal IP address that is defined by you. ⚠ Important Use of this option might result in IP address conflicts.
Static (Reserved)	*Not available for legacy networks.* The network interface uses a static internal IP address that was reserved in Google Compute Engine in advance.
Static (New)	*Not available for legacy networks.* The network interface uses a new static internal IP address that is defined by you. Note 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.

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.

12. If you want to add tags and/or labels to the restored instance, click **Add Tags** and then for each and/or tag that you want to add, do the following:
    a. From the Available options drop-down menu, select whether you want to add a label, a network tag, or a custom metadata tag.
    b. 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. If you want to delete any of the added labels, network tags, and/or custom metadata tags, click ✕ next to it.

13. Click **Restore**.

Restoring disks

You can restore instance 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 Details section.

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

2. In the Details section that appears at the bottom of the screen, select the desired restore point.
3. Click **Restore Instance**. The Restore Options dialog box opens.
4. Select **Restore Disks**, and then click **Next**.
5. 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.

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. *Optional.* In the Post-restore script field, enter the path to the script or a command that HYCU for Google 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.

8. Click **Restore**.

Cloning disks

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

Limitation

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.

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.
- For details on how the restored disks are named, see "Objects created by the service" on page 140.

Procedure

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

   > Note The Details section appears only if you click an instance. Selecting the check box before the name of the instance will not open the Details section.

2. In the Details section that appears at the bottom of the screen, select the desired restore point.
3. Click **Restore Instance**. The Restore Options dialog box opens.
4. Select **Clone Disks**, and then click **Next**.
5. From the list of disks that are available for the restore, select the ones that you want to restore, 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. Select the project and the zone that contain the instance to which you want to attach the restored disks, and then select the instance to which you want to attach the restored disks.
8. *Optional.* In the Post-restore script field, enter the path to the script or a command that HYCU for Google 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. Edit the disks as required. For each selected disks, do the following:
   a. Click **Edit Disk**.
   b. *Only if you do not want HYCU for Google 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, 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 (Standard persistent disk, Balanced persistent disks, or SSD persistent disk). By default, the original disk type is selected.
   d. 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 Replica zones drop-down menu, select to which two zones 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. If you want to add labels to the restored disk, click **Add Tags**, 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 ✕ next to it.

   f. Click **Save**.
10. Click **Restore**.

Exporting disks

You can export instance disks by restoring them to the same or a different project or zone. In this case, the disks will not be attached to any instance.

Prerequisite

Only if you plan to restore disks to a different project. The default network must be set for the project to which you plan to restore disks, or the project to which you plan to restore disks must have the same network as the instance whose disks you plan to restore.

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.
- For details on how the restored disks are named, see "Objects created by the service" on page 140.

Procedure

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

   > Note The Details section appears only if you click an instance. Selecting the check box before the name of the instance will not open the Details section.

2. In the Details section that appears at the bottom of the screen, select the desired restore point.
3. Click **Restore Instance**. The Restore Options dialog box opens.
4. Select **Export Disks**, and then click **Next**.
5. From the list of disks that are available for the restore, select the ones that you want to restore, 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 Target project drop-down menu, select the project to which you want to restore the disks. You can choose from the projects that belong to the currently selected protection set.
8. From the Target region and Target zone drop-down menus, select the Google Cloud region and zone to which you want to restore the disks.
9. Edit the disks as required. For each selected disks, do the following:
    a. Click ✎ **Edit Disk**.
    b. *Only if you do not want HYCU for Google 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, 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 (Standard persistent disk, Balanced persistent disks, or SSD persistent disk). By default, the original disk type is selected.
    d. 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 Replica zones drop-down menu, select to which two zones 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. If you want to add labels to the restored disk, click **Add Tags**, 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 ✕ next to it.

    f. Click **Save**.
10. Click **Restore**.

Restoring multiple instances in a single session

You can restore 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.

Limitation

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

> Accessing the Instances panel

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

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**. The Restore Instances dialog box opens.
4. From the Target project drop-down menu, select the project to which you want to restore the instances. You can choose from projects that belong to the currently selected protection set and that your user account can access.
5. From the Target region and Target zones drop-down menus, select the Google Cloud region and zone to which you want to restore the instances.
6. *Optional.* Enter the target instance postfix and the target disk postfix to add a postfix to the names of the target instances and disks.
7. *Optional.* In the Post-restore script field, enter the path to the script or a command that HYCU for Google 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 target region and zone.
9. *Only if you want to start the restore immediately.* Click **Restore**.
10. *Only if you want to edit the restore specification before executing it.*
    a. Click **Edit Restore Spec**. The restore specification generated by HYCU for Google Cloud for all selected instances is displayed, together with the POST method URL. You can edit the specification or copy it to the clipboard by clicking **Copy to clipboard**, and then edit it with a text editor of your choice.

       For details on the generated restore specification and the options used, see "Bulk restore specifications" on page 146.

    b. Click **Execute** to execute the modified restore specification.

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.

Limitation

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

> **Accessing the Instances panel**
>
> To access the Instances panel, in the navigation pane, click **Instances**. Alternatively, in the Dashboard panel, click the **Instances** widget.

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 Target project drop-down menu, select the project to which you want to restore the disks. You can choose from projects that belong to the currently selected protection set and that your user account can access.
5. From the Target region and Target zones drop-down menus, select the Google Cloud region and 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 start the restore immediately.* Click **Restore**.
9. *Only if you want to edit the restore specification before executing it.*
   a. Click **Edit Restore Spec**. The restore specification generated by HYCU for Google Cloud for all selected instances is displayed, together with the POST method URL. You can edit the specification or copy it to the clipboard by clicking **Copy to clipboard**, and then edit it with a text editor of your choice.

      For details on the generated restore specification and the options used, see "Bulk restore specifications" on page 146.
   b. Click **Execute** to execute the modified restore specification.

Restoring individual files or folders

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

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 target" on page 65

> **Accessing the Instances panel**

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

Restoring files or folders to an instance

You can restore one or more individual files or folders to the same or a new location on the original instance, or to a custom location on a different instance.

Prerequisites

- The instance to which you are restoring data is up and running.
- The target disk volume uses one of the supported file systems. For details, see the *HYCU for Google Cloud Compatibility Matrix*.
- In the Google Cloud Console, there is a network firewall rule applied to the instances—either to the entire network or to individual instances through the use of network tags. For each target instance, the rule must allow ingress network traffic through a specific TCP port from the entire subnetwork that the instance belongs to. The port number depends on the guest operating system of the instance and connection server configuration:
  - GNU/Linux: TCP port 22 (or a different port if configured for SSH communication)
  - Microsoft Windows: TCP port 5986 (or a different port if configured for WinRM communication)

  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.

  For instructions on how to configure and apply the network firewall rule, see Google Cloud documentation.
- *For Linux instances:*
  - On the original instance, an SSH server is installed and configured to use a TCP port (by default, 22) for SSH communication. The firewall is configured to enable inbound network traffic through this port.
  - *Only if the SSH server is configured to use a non-default TCP port or public key authentication, or OS Login is enabled on the instance in Google Compute Engine.* An appropriate credential group is assigned to the original instance.
- *For Windows instances:*
  - On the original instance, WinRM is configured to use a TCP port (by default, 5986) for HTTPS or HTTP communication. The firewall is configured to enable inbound network traffic through this port.
  - An appropriate credential group is assigned to the original instance, and the supplied credentials belong to a user account with sufficient privileges. Credential group assignment is performed automatically by HYCU for Google Cloud. For

instructions on how to manually assign credential groups, see "Enabling access to data" on page 36.

- *Only if you plan to restore individual files or folders to a different instance.* The discovery status of the instance to which you want to restore the individual files or folders is ✅.

Limitations

- You cannot restore individual files or folders located on an extended Master Boot Record (MBR) partition to their original location.
- *For restoring files or folders to a different instance:*
  - You can restore individual files or folders 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.
  - *Only if you plan to enable the Restore ACL option.* An LDAP directory service or any similar directory information service is configured.

Considerations

- HYCU for Google Cloud considers folders as containers of the file system objects. This means that in a restore task:
  - Folders are never renamed.
  - Folder access control lists (ACLs) are never restored and the original folder ACLs are kept on the file system.
- For details on how the restored individual files or folders are named, see "Objects created by the service" on page 140.
- *For Linux instances:* Depending on whether the instance has a credential group assigned in HYCU for Google Cloud, the following user accounts are used for the restore task:
  - *No credential group is assigned:* The HYCU Managed Service Account (HMSA).
  - *A credential group is assigned:* The user account that is specified in the credential group.

Restoring files or folders to the original instance

Procedure

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

   > Note The Details section appears only if you click an instance. Selecting the check box before the name of the instance does not open the Details section.

2. In the Details section, select the desired restore point, and then click **Restore Files**.

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. In the File Restore Options dialog box, select **Restore to original instance**, and then click **Next**.
4. In the Restore Settings 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 and most cost-effective restore.
      - **Backup (Snapshot)**
      - **Backup (Target)**
      - **Copy**
      - **Archive—(daily, weekly, monthly, yearly)**
   b. In the Disks drop-down menu, make sure only the disks with the files or folders that you want to restore are selected.
   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, 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. In the Restore to Instance dialog box, do the following:
   a. Select the location on the instance where you want to restore the files or folders, and provide the required information:
      - **Original location**

        Select how the restore should save the files when there is a file with the same name at the original location (overwrite the file, rename the original file, or rename the restored file).

        For naming conventions, see "Objects created by the service" on page 140.
      - **Alternate location**

        Specify the path to an alternate location on the instance in the following format:
        - Linux:

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

        - Windows:

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

The restored file overwrites the file with the same name that might exist at the alternate location.

b. Use the **Restore ACL** switch if you want to restore the original access control list. If enabled, HYCU for Google Cloud preserves original ACLs. If disabled, HYCU for Google Cloud applies inherited ACLs on the restored files (according to the file system ACL inheritance rules).

7. Click **Restore**.

Restoring files or folders to a different instance

Procedure

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

   > Note The Details section appears only if you click an instance. Selecting the check box before the name of the instance does not open the Details section.

2. In the Details section, select the desired restore point, and then click **Restore Files**.

   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. In the File Restore Options dialog box, select **Restore to different instance**, and then click **Next**.
4. In the Restore Settings 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 and most cost-effective restore.
      - **Backup (Snapshot)**
      - **Backup (Target)**
      - **Copy**
      - **Archive—(daily, weekly, monthly, yearly)**

   b. In the Disks drop-down menu, make sure only the disks with the files or folders that you want to restore are selected.
   c. From the Project drop-down menu, select the project that contains the instance to which you want to restore data.

      > Note The list of available projects includes only the ones that belong to the currently selected protection set.

   d. From the Zone drop-down menu, select the zone to which the instance to which you want to restore data belongs.
   e. From the Instances drop-down menu, select the instance to which you want to

restore data.

f. 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, 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. In the Restore to Different Instance dialog box, do the following:

   a. Specify the path to a custom location on the instance to which you want to restore data in the following format:

      - Linux:

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

      - Windows:

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

      The restored file overwrites the file with the same name that might exist at the custom location on the instance to which you want to restore data.

   b. Select how the restore should save the files when there is a file with the same name on the different instance (overwrite the file, rename the original file, or rename the restored file).

      For naming conventions, see "Objects created by the service" on page 140.

   c. Use the **Restore ACL** switch if you want to restore the original access control list. If enabled, HYCU for Google Cloud preserves original ACLs. If disabled, HYCU for Google Cloud applies inherited ACLs on the restored files (according to the file system ACL inheritance rules).

7. Click **Restore**.

Restoring files or folders to a target

Prerequisite

At least one target is set up in the protection set that includes the project of the original instance. For information on how to add manually created targets, see "Adding a bucket to HYCU for Google Cloud as a target" on page 24.

Consideration

For details on how the restored individual files or folders are named, see "Objects created by the service" on page 140.

Procedure

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

   > Note The Details section appears only if you click an instance. Selecting the check box before the name of the instance does not open the Details section.

2. In the Details section, select the desired restore point, and then click **Restore Files**.

   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. In the File Restore Options dialog box, select **Restore to target**, and then click **Next**.
4. In the Restore Settings 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 and most cost-effective restore.
      - **Backup (Snapshot)**
      - **Backup (Target)**
      - **Copy**
      - **Archive—(daily, weekly, monthly, yearly)**

   b. In the Disks drop-down menu, make sure only the disks with the files or folders that you want to restore are selected.
   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, 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. From the Target name drop-down menu, select a target to which you want to restore data.
7. Click **Restore**.

Chapter 5

Protecting applications

HYCU for Google Cloud enables you to protect your 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 for Google 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 121.
- HYCU for Google 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 application data efficiently, see the following sections:

- "Preparing for SAP HANA application protection" below
- "Preparing for Google Kubernetes Engine application protection" on page 70
- "Backing up applications" on page 74
- "Restoring SAP HANA applications" on page 75
- "Restoring Google Kubernetes Engine applications" on page 77

Preparing for SAP HANA application protection

Before you start protecting SAP HANA applications, perform the following tasks:

Task	Instructions
Mandatory. Make sure HYCU for Google Cloud can access applications that you want to protect.	"Enabling access to application data" below
Optional. Specify the location and the subnet for the temporary instance that HYCU for Google Cloud creates during the backup.	"Specifying the temporary instance location and subnet" on the next page
Optional. Configure the backup chain length.	"Configuring the backup chain length" on page 70

Enabling access to application data

After you assign credentials to instances as described in "Enabling access to data" on page 36, 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
✅	HYCU for Google 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 "Backing up applications" on page 74.
❌	The instance credentials do not have proper permissions and HYCU for Google Cloud cannot access applications. To enable HYCU for Google 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 "Enabling access to data" on page 36. 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.

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. In the Credentials section, 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 applications" on page 74.

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 36. 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.

Specifying the temporary instance location and subnet

You can specify the region, the zone, and the subnet where you want HYCU for Google Cloud to create a temporary instance during the backup. By default, the temporary instance is created in the original project of the application.

Prerequisite

VPC Network Peering must be configured. For details on how to configure VPC Network Peering, see Google Cloud documentation.

> **Accessing the Applications panel**
>
> To access the Applications panel, in the navigation pane, click **Applications**.

Procedure

1. In the Applications panel, select the application for which you want to select the temporary instance location.
2. Click **Configuration**. The Application Configuration dialog box opens.
3. In the Temporary Instance Configuration section, specify the region, the zone, and the subnet:
   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**.

Configuring the backup chain length

You can adjust application protection to the needs of your data protection environment by configuring 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**.

Procedure

1. In the Applications panel, select the application for which you want to configure the backup chain length.
2. Click **Configuration**. The Application Configuration dialog box opens.
3. In the Backups section, specify when you want a new backup chain to be started.
4. Click **Save**.

Preparing for Google Kubernetes Engine application protection

Before you start protecting your Google Kubernetes Engine (GKE) applications, get familiar with prerequisites, limitations, and procedures in this section to prepare your environment for application data protection.

Prerequisite

The HYCU Managed Service Account (HMSA) must have the following roles granted on the projects with the Kubernetes clusters on which the GKE applications that you plan to protect are deployed:

- Compute Admin (`roles/compute.admin`)
- Service Account User (`roles/iam.serviceAccountUser`)
- Storage Admin (`roles/storage.admin`)
- Kubernetes Engine Admin (`roles/container.admin`)

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.

- HYCU for Google Cloud does not support protecting applications that are running on a public GKE cluster or a private GKE cluster with the selected Access control plane using its external IP address option, and that are at the same time configured in a subnet where Google Private Access is enabled.
- *For applications using volumes:* Only GCE persistent disk volumes and CSI volumes are supported.

Before you start protecting GKE applications, perform the following tasks:

Task	Instructions
Mandatory. Make sure appropriate labels are applied on all resource objects.	"Applying labels on resource objects" below
Mandatory. Make sure your GKE applications are discovered in HYCU for Google Cloud.	"Discovering applications" on the next page
Optional. Specify the location and the subnet for the temporary instance that HYCU for Google Cloud creates during the backup.	"Configuring Google Kubernetes Engine application backup options" on the next page
Optional. Specify pre-backup and post-backup scripts.	"Configuring Google Kubernetes Engine application backup options" on the next page

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 HYCU for Google 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/<project-id>/zones/<zone>/disks/<disk-name>`.)

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 **Synchronize**.

> **Note** Before a GKE application can be discovered, the Kubernetes cluster on which it is deployed must be discovered by HYCU for Google Cloud. This is an automated task that is performed every 15 minutes.

Configuring Google Kubernetes Engine application backup options

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

Backup option	Description
Specifying the temporary instance location and subnet	You can select the region, the zone, and the subnet where you want HYCU for Google 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.
Running pre/post scripts	These options allow you to specify the pre-backup and post-backup scripts to perform necessary actions before and after the backup of the application is performed.

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'
```

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 application 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 temporary instance location and subnet.*

     On the Temporary instance configuration tab, specify the region, the zone, and the subnet:

     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.
   - *Only if specifying pre-snapshot and post-snapshot scripts.* Specify the scripts to perform necessary actions before and/or after the backup of the application is performed:
     - In the Pre-Backup Script field, enter the path to the script that HYCU for Google Cloud will run just before it performs the backup of the application.
     - In the Post-Backup Script field, enter the path to the script that HYCU for Google Cloud will run immediately after it performs the backup 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-backup script:

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

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

4. Click **Save**.

Backing up applications

With HYCU for Google Cloud, you can back up SAP HANA and GKE applications in a fast and efficient way.

Prerequisites

- *Only if you plan to back up applications running on instances or 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`.
- *For backing up SAP HANA applications:*
  - The minimum required SAP HANA privileges of the configured SAP HANA database user are BACKUP ADMIN and CATALOG READ.
  - The configured SAP HANA database user has access permissions to all databases that are being backed up.
  - *For SAP HANA systems with the same SID:* A separate target is configured for each SAP HANA system.

Considerations

- *Only if backing up SAP HANA applications.* Consider the following:
  - 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, HYCU for Google 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 in panels" on page 105.

2. Click **Policies**. The Policies dialog box opens.
3. From the list of available policies, select the desired 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.

> Note If required, you can also perform a manual backup of any application at any time. For details, see "Performing manual backups" on page 115.

Restoring SAP HANA applications

HYCU for Google 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 is 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 exist.
  - A system database is online and tenant databases are 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 to open the Details section.

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

2. In the Details section that appears at the bottom of the screen, select the desired 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 Instances drop-down menu, select the instance to which you want to restore application data.
6. Select the **All 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**.

Restoring Google Kubernetes Engine applications

HYCU for Google 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 not supported for applications not using persistent volumes.
- *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	"Restoring a whole application" on the next page
Restore application storage to the original or a different location.	Restore Storage	"Restoring storage" on page 79
Restore specific resource objects to the original or a different location.	Restore Resource Objects	"Restoring resource objects" on page 80

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 to open the Details section.

   > Note The Details section appears only if you click an application. Selecting the check box before the name of the application will not open the Details section.

2. In the Details section that appears at the bottom of the screen, select the desired restore point, and then click **Restore**. The Restore Kubernetes Application 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 HYCU for Google 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 to open the Details section.

   > Note The Details section appears only if you click an application. Selecting the check box before the name of the application will not open the Details section.

2. In the Details section that appears at the bottom of the screen, select the desired restore point, and then click **Restore**. The Restore Kubernetes Application 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 HYCU for Google 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 **All 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 to open the Details section.

   > Note The Details section appears only if you click an application. Selecting the check box before the name of the application will not open the Details section.

2. In the Details section that appears at the bottom of the screen, select the desired restore point, and then click **Restore**. The Restore Kubernetes Application 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 HYCU for Google 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 buckets

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

Prerequisites

- The HYCU Managed Service Account (HMSA) must have the following roles granted on the projects with the buckets that you plan to protect:
  - Compute Admin (`roles/compute.admin`)
  - Service Account User (`roles/iam.serviceAccountUser`)
  - Storage Admin (`roles/storage.admin`)

  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.

Limitation

Bucket data (backup data, copies of backup data, and data archives) can be stored only to manually created targets, and not to automatically created targets or as a snapshot. For instructions on how to add a bucket to HYCU for Google Cloud as a target, see "Adding a bucket to HYCU for Google Cloud as a target" on page 24.

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 121.
- HYCU for Google 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

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

Backup option	Description
Specifying the temporary instance location and subnet	You can specify the location and the subnet where you want HYCU for Google Cloud to create a temporary instance during the backup. By default, the temporary instance is created in the original project of the bucket.
Running pre/post scripts	You can specify the pre-backup and post-backup scripts to perform necessary actions before and after the backup of the bucket is performed.

Prerequisites

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

- The HYCU Managed Service Account (HMSA) must have access to the bucket where the script is located.
- The `#!/usr/bin/env python3` header must be specified in the script.
- *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'
```

Limitations

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

- Currently, only Python scripts are supported.
- Only the `googleapiclient` Python library can be used for making Google Cloud API calls.

Consideration

Only when specifying the location or the subnet for a temporary instance. If not specified otherwise, the temporary instance will be created in the following region (based on the location type of the bucket):

- *A region:* In the same region as the bucket (for example, US-CENTRAL1).
- *A dual-region:*

Dual-region name	Temporary instance region
ASIA1	ASIA-NORTHEAST1
EUR4	EUROPE-NORTH1
NAM4	US-CENTRAL1

- *A 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 bucket 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:
   - Specify the region, the zone, and the subnet where you want HYCU for Google Cloud to create a temporary instance:
     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 Google Cloud pricing.

     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.

     > **Important** A policy cannot be assigned to a bucket on which HYCU for Google Cloud could not detect the subnet.

   - 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 HYCU for Google 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 HYCU for Google 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:
>
> `gs://bucket-name/script.py parameter1 parameter2 ...`

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

```
#!/usr/bin/env python3
import os
import googleapiclient.discovery

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

storage = googleapiclient.discovery.build('storage', 'v1')
```

4. Click **Save**.

Backing up buckets

With HYCU for Google Cloud, you can back up data that is stored in Google Cloud Storage buckets in a fast and efficient way.

Prerequisite

Only 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`.

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 **Synchronize**.

   > Tip To narrow down the list of displayed buckets, you can use the filtering options as described in "Filtering and sorting data in panels" on page 105.

2. Click **Policies**. The Policies dialog box opens.
3. From the list of available policies, select the desired policy.
4. Click **Assign** to assign the policy to the selected buckets.

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.

> Note If required, you can also perform a manual backup of any bucket at any time. For details, see "Performing manual backups" on page 115.

Restoring buckets

HYCU for Google Cloud enables you to restore one or more individual files or folders inside a Google Cloud Storage bucket to the original or a different bucket.

Prerequisite

Only 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`.

Consideration

For details on how the restored individual files or folders are named, see "Objects created by the service" on page 140.

Procedure

1. In the Buckets panel, click the bucket that contains the files or folders that you want to restore. The Details section appears at the bottom of the screen.

   > Note The Details section appears only if you click a bucket. Selecting the check box before the name of the bucket does not open the Details section.

2. In the Details section, select the desired restore point, and then click **Restore Files**.

   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. In the Choose Files and Folders dialog box, 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**.

5. Depending on where you want to restore data, select the desired restore option, and then follow the instructions:

Restore option	Instructions
Restore to original bucket	a. Select the location on the bucket where you want to restore the files or folders, and provide the required information: • **Original location** 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). • **Alternate location** 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. b. Use the **Restore ACL** switch if you want to restore the original access control list. If enabled, HYCU for Google Cloud preserves original ACLs. If disabled, HYCU for Google Cloud applies inherited ACLs on the restored files (according to the ACL permissions at the bucket or project level). c. If you want to add custom metadata tags to the restored bucket objects, click **Add Tags**, enter a key and value, and then click **Add** for each custom metadata tag that you want to add. > Note If you want to delete any of the added custom metadata tags, click ✕ next to it.
Restore to different	a. From the Project drop-down menu, select the project that contains the bucket to which you want to restore data.

Restore option	Instructions
bucket	Note You can select only among the projects inside the selected protection set. b. From the Bucket name drop-down menu, select the name of the bucket to which you want to restore data, and then click **Next**. c. Select the location on the bucket where you want to restore the files or folders, and provide the required information: • **Original location** 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). • **Alternate location** 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. d. Use the **Restore ACL** switch if you want to restore the original access control list. If enabled, HYCU for Google Cloud preserves original ACLs. If disabled, HYCU for Google Cloud applies inherited ACLs on the restored files (according to the ACL permissions at the bucket or project level). e. If you want to add custom metadata tags to the restored bucket objects, click **Add Tags**, enter a key and value, and then click **Add** for each custom metadata tag that you want to add. Note If you want to delete any of the added custom metadata tags, click ✕ next to it.

6. Click **Restore**.

Chapter 7

Performing daily tasks

To ensure your data protection environment is in the optimal state in terms of security, reliability, and efficiency, HYCU for Google 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 121.

I want to ...	Instructions
Get an at-a-glance overview of the data protection environment state, identify eventual bottlenecks, and inspect different areas of the data protection environment.	"Using the HYCU for Google Cloud dashboard" on the next page
Track tasks that are running in the data protection environment and get an insight into the status of a specific task.	"Checking task statuses" on page 91
View all events that occurred in my data protection environment.	"Viewing events" on page 92
Obtain HYCU for Google Cloud reports on different aspects of the data protection environment.	"Using HYCU for Google Cloud reports" on page 96
View instance, application, and bucket details.	"Viewing instance, application, and bucket details" on page 101
Narrow down the list of displayed items by applying filters and sort the items in panels.	"Filtering and sorting data in panels" on page 105
View target information, activate or deactivate a target, and edit or remove a target.	"Managing targets" on page 111
View policy information, edit a policy, or delete a policy.	"Managing policies" on page 114
Back up data manually.	"Performing manual backups" on

I want to ...	Instructions
	page 115
Mark a restore point as expired.	"Expiring backups manually" on page 116
Export data that I can view in a table in any of the panels to a JSON or CSV file.	"Exporting the contents of the panel" on page 95

Using the HYCU for Google Cloud dashboard

The HYCU for Google Cloud dashboard provides you with an at-a-glance overview of the data protection status in your environment. This intuitive dashboard enables you to monitor all data protection activity and to 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
Instances	Percentage of protected instances, and the exact number of protected and unprotected instances in the protection set. An instance is considered protected if it has a policy assigned and at least one valid backup within the retention period specified in the policy. Instances that have the exclude policy assigned are omitted from the figures depicted in this widget. For details about instances, see "Protecting instances" on page 41.
Applications	Percentage of protected applications, and the exact number of protected and unprotected applications in the protection set. An application is considered protected if it has a policy assigned and at least one valid backup within the retention period specified in the policy. For details about applications, see "Protecting applications" on page 67.
Subscription	Information about your HYCU for Google Cloud subscription.
Backups	Backup success rate for the last seven days.
Targets	Number of targets in the protection set, and the information about how much space is used and available for storing backup data. For details, see "Setting up targets" on page 23.

Widget	Description
Policies	Percentage of policies that are compliant, and the number of compliant and non-compliant policies in the protection set. A policy is considered compliant if all instances, applications, or buckets to which this policy is assigned are compliant with the policy settings. For details on policies, see "Defining your backup strategy" on page 25.
Tasks	Total number of tasks in the protection set, and the number of tasks according to their status (Success, Warning, Failed, In progress) in the last 48 hours. For details on tasks, see "Checking task statuses" below.
Events	Total number of events in the protection set, and the number of events according to their severity level in the last 48 hours. For details on events, see "Viewing events" on the next page.

> 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.

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 Details section that appears at the bottom of the screen after you select the task.

  > Tip To minimize the Details section, click ▼ **Minimize** or press the Spacebar. To return it 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 Report**. To copy the report to the clipboard, in the Task Report dialog box that opens, click **Copy to clipboard**.

- Cancel any currently running task by selecting it, and then clicking **Abort Task**. Keep in mind that you cannot abort tasks related to retention maintenance.

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).
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 Details section that appears at the bottom of the screen after you select the event.

  > Tip If you click the related task link in the Details section, 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 HYCU for Google Cloud to send notifications when new events occur in your data protection environment. For details, see "Configuring event notifications" on the next page.

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: • (Info): Events representing regular service operation. • (Warning): Potentially harmful situations that do not represent an immediate threat to service operation. • (Error): Errors that immediately affect service operation.
Message	Description of the event.
Category	Functional area of HYCU for Google Cloud to which the event belongs (for example, Targets, Credentials, Policies, System for an internal event, and so on).
Timestamp	Event creation date and time.

Configuring event notifications

You can configure HYCU for Google 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 the previous page.
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 the previous page.
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.

You can later edit settings for existing email notifications (click **Edit** and make the required modifications) or delete the ones that you do not need anymore (click **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 categories. To include all categories, click **Select All**.
4. From the Status drop-down menu, select one or more statuses. To include all statuses, click **Select All**.
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 HYCU for Google 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 HYCU for Google Cloud sends to the specified URL, see "Webhook data format" below.
9. Click **Save**.

Your changes take effect immediately and webhook notifications are sent to any URL that you specified in the notification settings.

You can later edit settings for existing webhook notifications (click **Edit** and make the required modifications) or delete the ones that you do not need anymore (click **Delete**).

Webhook data format

The webhook data format is defined by:

- HTTP request header sent by HYCU for Google Cloud
- HTTP request body sent by HYCU for Google Cloud

- HTTP response code sent by the webhook endpoint and received by HYCU for Google 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`.

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 box 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
Export to JSON (Current)	Exports the current table page to a JSON file.
Export to JSON (All)	Exports all table data to a JSON file.
Export to CSV (Current)	Exports the current table page to a CSV file.
Export to CSV (All)	Exports all table data to a CSV file.

Using HYCU for Google Cloud reports

HYCU for Google 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 98.
- Generate reports. For details, see "Generating reports" on page 99.
- Schedule reports. For details, see "Scheduling reports" on page 99.

  > 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 100.

Accessing the Reports panel

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

> Tip To minimize the Details section, click **Minimize** or press the Spacebar. To return it to its original size, click **Maximize** or press the Spacebar.

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 the next page.

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 instance disks and the total size of instance backup 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.
protected-data-on-targets-per-vm	Amount of protected data on targets for each protected instance.
protected-data-on-targets-per-policy	Amount of protected data on targets for each policy.
protected-data-on-targets-per-storage-class	Amount of protected data on targets for each storage class.
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.
total-protected-data-on-targets-trend	Total amount of protected data on targets 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.

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.	1. Click + **New**. The Report Configuration dialog box opens.

I want to...	Procedure
	2. Enter a report name and, optionally, its description. 3. Select the type of report (a table or a chart). 4. Specify the time range for the report. 5. Select the aggregation value that you want to use to perform a calculation on a set of collected data. 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. 7. Click **Save**.
Edit an existing report and save it as a new report.	1. From the list of reports, select the one that you want to edit and save as a new report, and then click **Edit**. The Report Configuration dialog box opens. 2. Enter a new name for the report, and then make the required modifications. 3. Click **Save as**.

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.

I want to...	Procedure
View a report on the current state of my data protection environment.	From the list of reports, select the desired report, and then double-click it or click **Preview**.
View a saved report version.	1. From the list of reports, select the desired report. 2. In the Details section that appears at the bottom of the screen, select the desired report version, and then double-click it or click **View Report**. For instructions on how to generate report versions manually or automatically, see "Generating reports" on the next page or "Scheduling reports" on the next page.

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.

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 97.

2. In the Details section 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 Details section 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 desired 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.

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 97.

2. In the Schedule date box, 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 **Save**.

> 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

HYCU for Google Cloud enables you to share user-created reports among different HYCU for Google 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.

Viewing instance, application, and bucket details

You can view the details about each instance, discovered application, or bucket in the Detail view section of the Instances, Applications, or Buckets panel.

> **Note** The Details section appears only if you click an instance, an application, or a bucket. Selecting the check box before its name will not open the Details section.

The following details are available:

Instance, application, or bucket information	Description
Summary	Shows detailed information about the selected instance, application, or bucket.
Restore point	Shows the following information for the restore point: • Creation date and time. • Available tiers from which you can restore data: ◦ *For instances and Google Kubernetes Engine applications:* ▪ SNAP or S: Snapshot. Displayed if a snapshot of the instance or the Google Kubernetes Engine application using persistent volumes exists. Snapshots allow faster completion of restore tasks. ▪ BCKP or B: Backup data on a target. Displayed if backup data is stored on a target. ▪ COPY or C: Copy of a backup image. Displayed if a copy of a backup image (snapshot or backup data on a target) exists on another target. ▪ ARCH-D or D: Data archive—daily. Displayed if a daily data archive exists on a target. ▪ ARCH-W or W: Data archive—weekly. Displayed if a weekly data archive exists on a target. ▪ ARCH-M or M: Data archive—monthly. Displayed if a monthly data archive exists on a target. ▪ ARCH-Y or Y: Data archive—yearly. Displayed if a yearly data archive exists on a target. ▪ CTLG or C: Catalog. Displayed if a restore of

	individual files or folders is available. *(Available only for instances.)* > **Note** A restore point may or may not include backup data of the entire instance. This depends on the disks included in the corresponding backup. Visual labels of the tiers may be specially marked to denote different statuses. For more information, see "Tier statuses" on page 104. ◦ *For SAP HANA applications:* ▪ FULL: Full backup. ▪ INCR: Incremental backup. ◦ *For buckets:* ▪ BCKP or B: Backup data on a target. ▪ COPY or C: Copy of backup data. Displayed if a copy of a backup data exists on another target. ▪ ARCH-D or D: Data archive—daily. Displayed if a daily data archive exists on a target. ▪ ARCH-W or W: Data archive—weekly. Displayed if a weekly data archive exists on a target. ▪ ARCH-M or M: Data archive—monthly. Displayed if a monthly data archive exists on a target. ▪ ARCH-Y or Y: Data archive—yearly. Displayed if a yearly data archive exists on a target.
Compliance	Shows the compliance status of the backup (and the resulting restore point): • The ✓ icon: The backup is compliant (the RPO setting in the policy assigned to the instance, the application, or the bucket was met). • The ✗ icon: The backup is not compliant (the RPO setting in the policy assigned to the instance, the application, or the bucket was not met). • The ? icon: The backup compliance status is undefined (the backup is still running). 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.

Backup status	Shows the backup status of your instance, application, or bucket. For more information, see "Viewing the backup status of instances, applications, and buckets" below.
Restore status	Shows a progress bar indicating the progress of the instance, application, or bucket restore. **Tip** 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 Details section, click ▼ **Minimize** or press the Spacebar. To return it to its original size, click ▲ **Maximize** or press the Spacebar.

Viewing the backup status of instances, applications, and buckets

The backup status of your instance, application, or bucket determines whether it is possible to restore it.

Backup status	Restore an instance, a GKE application, or disks?	Restore files?	Restore a SAP HANA application?	Restore a bucket?
(Done)	✓	✓[a]	✓	✓
(Done with warnings)	✓	✓[a]	✓	✓
(Done with errors)	✓[b]	?[c]	✓[d]	✓[e]
(Failed)	×	×	×	×
Aborted	×	×	×	×
(Expired / Inaccessible on Google Cloud / Deleted from Google Cloud)	×	×	×	×

[a] All instance 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.

[b] This backup status may indicate one of the following:

- Not all instance or application disks were backed up successfully, therefore the instance or application 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 instance or application 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.* Some actions specified by the scripts might not be performed.

[c] This backup status may indicate one of the following:

- Not all instance 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 instance disks were backed up successfully, therefore only the files that belong to the disks displayed in the Restore Settings dialog box can be restored.

[d] An application can be partially restored (only the databases that are displayed in the Application Restore dialog box).

[e] *Applicable only if you are using the pre-backup and post-backup scripts.* Some actions specified by the scripts might not be performed.

> 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 instance, an application, or a bucket. The following is an example of possible marks:

Tier status	Restore an instance, a GKE application, or a bucket?
BCKP or B (Done)	✓
BCKP or B (Done with warnings or Done with errors)	✓ For details on what data can be restored if one of these backup statuses is shown, see "Viewing the backup status of instances, applications, and buckets" on the previous page.
BCKP or B (Failed)	×
BCKP or B (Aborted)	×
BCKP or B (Expired)	×
BCKP or B (Inaccessible on Google Cloud)	×

Tier status	Restore an instance, a GKE application, or a bucket?
BCKP or B (Deleted from Google Cloud)	×

Filtering and sorting data in panels

HYCU for Google Cloud enables you to filter data in the panels so you can easily find what you need. Each panel contains different filtering options and it can display only the entries that meet the specified filter criteria. For example, filtering the data in the Instances panel helps you to focus only on the instances that you are interested in. In addition, you can sort displayed items in ascending or descending order based on an alphabetical value or a label. For example, sorting data in the Policies panel by the Compliance label helps you easily track non-compliant policies.

> Tip After selecting a set of items in the filtered view, you can easily clear the list of selected items by clicking the × icon next to the number of displayed items.

Filtering data in panels

Procedure

1. Go to the web user interface panel of interest.
2. *Optional.* On the left side of the main pane, in the Search field, enter your main filter keyword. Which property can be used as the main filter keyword depends on the panel you are in.
3. To filter the data set (when no main filter keyword is specified) or filter the resulting data set further, follow the steps:
   a. On the right side of the main pane, click **Filters**. The Filters side pane opens.
   b. In the Filters pane, specify your filtering options.
   c. Click **Apply Filters**.

Depending on the panel the contents of which you want to filter, see one of the following sections for information on the available filtering options:

- "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 107
- "Filtering options in the Policies panel" on page 108
- "Filtering options in the Targets panel" on page 108
- "Filtering options in the Tasks panel" on page 109
- "Filtering options in the Events panel" on page 110

Filtering options in the Applications panel

You can enter an application name (or a part of it) as the main filter keyword.

In the Filters side panel, select one or more filtering options:

Filtering option	Action
Project	From the drop-down menu, select the Google Cloud projects to which the instances or Google Kubernetes Engine clusters with the applications you want to filter belong.
Type	From the drop-down menu, select the application type.
Policy	From the drop-down menu, select the policies that are assigned to the applications.
Compliance	Select one or more options to filter by the compliance status: • **Success**: The application is compliant. • **Failure**: The application is not compliant. • **Undefined**: The exclude policy is assigned to the application or the application does not have a policy assigned.

Protection	Select one or more options to filter by the protection status:
	• **Yes**: The application is protected.
	• **No**: The application is not protected.
	• **Deleted**:
	  • *For SAP HANA applications:* The instance with the application is deleted, the status of the instance with the application is PROTECTED_DELETED, or the application is deleted from the instance.
	  • For GKE applications: The cluster with the application is deleted, or the application deployment is deleted from the cluster (or can no longer be discovered).

Discovery	Select one or more options to filter by the application discovery status: • **Success**: One or more applications are discovered. • **Failure**: No applications were discovered. • **Warning:** Application discovery failed because the instance is offline or not reachable. • **Undefined**: The status of the discovered application is PROTECTED_DELETED.

Filtering options in the Instances panel

You can enter an instance name (or a part of it) as the main filter keyword.

In the Filters side panel, you can select one or more filtering options:

Filtering option	Action
Project	From the drop-down menu, select the Google Cloud projects of interest.
Policy	From the drop-down menu, select the policies that are assigned to the instances.
Credential group	From the drop-down menu, select the credential groups that are assigned to the instances.
Zone	From the drop-down menu, select the Google Compute Engine instance zones.
Compliance	Select one or more options to filter by the compliance status: • **Success**: The instance is compliant. • **Failure**: The instance is not compliant. • **Undefined**: The exclude policy is assigned to the instance, or the instance does not have a policy assigned.
Protection	Select one or more options to filter by the protection status: • **Yes**: The instance is protected. • **No**: The instance is not protected. • **Deleted**: The instance no longer exists, but at least one of its backup images does. • **Undefined**: The exclude policy is assigned to the instance.
Discovery	Select one or more options to filter by the instance discovery status: • **Success**: Connection to the instance was established (as part of checking the connectivity after assigning a credential group to the instance, selecting the Enable restore of individual files option, or specifying the pre-snapshot or post-snapshot scripts). • **Failure**: The instance could not be connected to. • **Warning**: The project has moved to another protection set. • **Undefined**: Connectivity to the instance has not been checked.

Filtering options in the Buckets panel

You can enter a bucket name (or a part of it) as the main filter keyword.

In the Filters side panel, you can select one or more filtering options:

Project	From the drop-down menu, select the Google Cloud projects to which the buckets that you want to filter belong.
Policy	From the drop-down menu, select the policies that are assigned to the buckets.

Location	From the drop-down menu, select the Cloud Storage location of the buckets.
Storage class	Select one or more options to filter by the Google Cloud storage class: • **Standard** • **Regional** • **Coldline** • **Nearline** • **Archive** • **Multi-regional** • **Durable reduced availability**
Compliance	Select one or more options to filter by the compliance status: • **Success**: The bucket is compliant. • **Failure**: The bucket is not compliant. • **Undefined**: The exclude policy is assigned to the bucket or the bucket does not have a policy assigned.
Protection	Select one or more options to filter by the protection status: • **Yes**: The bucket is protected. • **No**: The bucket is not protected. • **Deleted**: The bucket is deleted or the status of the bucket is PROTECTED_DELETED.

Filtering options in the Policies panel

You can enter a policy name (or a part of it) as the main filter keyword.

In the Filters side panel, you can select one or more filtering options:

Filtering option	Action
Compliance	Select one or more options to filter by the compliance status: • **Success**: All instances, applications, or buckets to which the policy is assigned are compliant. • **Failure**: Not all instances, applications, or buckets to which the policy is assigned are compliant. • **Undefined**: The policy is not assigned to any instance, application, or bucket, or this is the exclude policy.

Filtering options in the Targets panel

You can enter a target name (or a part of it) as the main filter keyword.

In the Filters side panel, you can select one or more filtering options:

Filtering option	Action
Storage class	Select one or more options to filter by the Google Cloud storage class: • **Standard** • **Regional** • **Coldline** • **Nearline** • **Archive** • **Multi-regional** • **Durable reduced availability**
Health	Select one or more options to filter by the status of the target: • **Ok** • **Warning** • **Error** • **Undefined**

Filtering options in the Tasks panel

You can enter a task description (or a part of it) or a task ID as the main filter keyword.

In the Filters side panel, you can select one or more filtering options:

Filtering option	Action
Project	From the drop-down menu, select the Google Cloud projects of interest.
Username	From the drop-down menu, select items to filter the list to include only the tasks started by any of the selected user accounts.
Type	From the drop-down menu, select one or more items to filter the list to include only the selected task types.
Status	Select one or more options to filter by the status of the task: • **Ready** • **Running** • **Aborting** • **Aborted** • **Done** • **Failed** • **Done with errors** • **Done with warnings** • **Skipped**
Time range	Specify a time range to limit your search for tasks. You can select one of the predefined time ranges (Last 1 hour, Last 24 hours, or

Filtering option	Action
	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

You can enter a text string as the main filter keyword.

In the Filters side panel, you can select one or more filtering options:

Filtering option	Action
Project	From the drop-down menu, select the Google Cloud projects of interest.
Category	From the drop-down menu, select items to filter the list to include only the selected event categories.
Username	From the drop-down menu, select items to filter the list to include only the events resulting from the selected user account actions.
Severity	Select one or more options to filter by the event severity: • **Success** • **Warning** • **Failed**
Time range	Specify a time range to limit your search for events. You can select one of the predefined time ranges (Last 1 hour, Last 24 hours, 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.

Sorting data in panels

Procedure

1. Go to the web user interface panel of interest.
2. Click the table column heading of the property that you want to sort the data in table rows by.

   The ^ icon appears in the heading cell, indicating that the column data is sorted in ascending order.
3. Click the column heading again to toggle the sort order.

   The ⌄ icon appears in the heading cell, indicating that the column data is sorted in descending order.

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.

Consideration

Only Google Cloud Storage buckets that were added to HYCU for Google Cloud as targets either automatically or manually are listed in the Targets panel. Snapshots are not included in this list.

> **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	Target name (globally unique). A target that has Object Lock (WORM) enabled is represented by the icon in the list of targets. For information on how automatically created targets are named, see "Objects created by the service" on page 140. **Tip** You can click the target name to open the target details page of the Google Cloud Console in your web browser.
Location	Name of the Google Cloud Storage region in which the target resides.
Storage Class	Default object storage class of the target in the Google Cloud Storage service: Multi-regional, Regional, Durable Reduced Availability, Standard, Nearline, Coldline, or Archive.
Status	Status of the target: • Active: You can use the target for backing up data, creating data archives, and restoring data. • Inactive: The target has been deactivated within HYCU for Google Cloud. As long as it is not activated you can use it only for restoring data.

Property name	Description
	• Inaccessible on Google Cloud: Insufficient permissions are set on the target in the Google Cloud Storage service. HYCU for Google Cloud cannot access the target. • Deleted from Google Cloud: The target no longer exists in the Google Cloud Storage service. For instructions on how to change the status of active or inactive targets, see "Deactivating and activating targets" on the next page.
Size Limit	Maximum amount of the target storage space (expressed in MiB, GiB, or TiB) that is allowed to be used by backup data created by HYCU for Google Cloud. The amount represents a soft limit, therefore actual usage may exceed it.
Health	Health status of the target: • The ? icon: Indicates one of the following: ◦ The target health has not been determined yet. ◦ The target is inactive. • 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. • The ! icon: Utilization of storage space for backup data in the target is over 90 percent and under 100 percent of the configured size limit, or the target is publicly accessible in Google Cloud. • The ✕ icon: Indicates one of the following: ◦ Target storage space occupied by backup data exceeds the configured size limit. ◦ The target is not accessible due to an I/O error, insufficient permissions, or some other reason.
Utilization	Ratio (expressed in percentage) between the target storage space occupied by backup data and the configured size limit.
Automatic	Indicator of whether the target was created automatically by HYCU for Google Cloud (✓) or not (✕).

To open the Details section where you can find more details about the target, click the desired target.

> Tip To minimize the Details section, click ▼ **Minimize** or press the Spacebar. To return it 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

Deactivation of a target makes the target unavailable for backup operations in HYCU for Google Cloud. The target remains registered with HYCU for Google Cloud with all the contained backup data intact. Restore of data from the target is still possible.

> **Note** You cannot deactivate targets that were created automatically by HYCU for Google Cloud.

Prerequisite

For target deactivation: The target is not specified in the Target option of any policy or data archive.

Consideration

After deactivating a target, the target cannot be selected for the Target option of a policy until it is activated again.

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 for deactivation.* Click **Yes** to confirm that you want to deactivate the selected target.

Removing targets

Removal of a target deregisters the target from HYCU for Google Cloud. After deregistration, the target and its contained data other than backup data continue to be available in your Google Cloud project.

Prerequisites

- The target contains no backup data.
- The target is not specified in the Target option of any policy or data archive.

Considerations

- After removing a target, no backup operations that include this target are possible anymore.
- You cannot remove targets that were created automatically by HYCU for Google Cloud unless they have been deleted from Google Cloud.

Procedure

1. In the Targets panel, select the target that you want to remove, and then click **Remove**.
2. Click **Yes** to confirm that you want to remove the selected target.

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.

Consideration

You cannot delete the exclude policy.

> **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: • The icon: The policy is compliant. • The icon: The policy is non-compliant. • The icon: Policy compliance is undefined. The policy is not assigned to any instance, application, or bucket, or this is the exclude policy.
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.

> **Note** To open the Details section where you can find more details about the policy, click the desired policy.

> **Tip** To minimize the Details section, click ▼ **Minimize** or press the Spacebar. To return it to its original size, click ▲ **Maximize** or press the Spacebar.

Creating a policy

See "Creating custom policies" on page 26.

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 26.
3. Click **Save**.

Deleting a policy

Procedure

1. In the Policies panel, select the policy that you want to delete, and then click **Delete**.
2. Click **Yes** to confirm that you want to delete the selected policy.

Performing manual backups

HYCU for Google Cloud backs up your data automatically after you assign a policy to the selected instances, applications, or buckets. 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 is assigned to the instance, the application, or the bucket.

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 Instances, Applications, or Buckets panel, select which instances, applications, or buckets you want to back up.
2. Click **Backup** to perform the backup of the selected instances, applications, or buckets.
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

HYCU for Google 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 an instance, an application, or a bucket, it can contain one or more tiers that you can mark as expired:

- *For instances and Google Kubernetes Engine applications:* Snapshot, Backup, 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 old backups for an instance, an application, or a bucket, access one of the following panels:

- Accessing the Instances panel

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

- Accessing the Applications panel

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

- Accessing the Buckets panel

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

Procedure

1. In the Instances, Applications, or Buckets panel, click the instance, the application, or the bucket for which you want to expire a backup. The Details section appears at the bottom of the screen.

   > Note The Details section appears only if you click an instance, an application, or a bucket. Selecting the check box before its name does not open the Details section.

2. In the Details section, select the restore point that you want to mark as expired.
3. Click **Expire**. The Expire dialog box opens
4. *Only if marking an instance, a GKE application, or a bucket 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 instances and GKE applications using persistent volumes.*
   - Backup (Target)
   - Copy
   - Archive - daily
   - Archive - weekly
   - Archive - monthly
   - Archive - 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 first next retention maintenance task in HYCU for Google Cloud removes the corresponding data from Google Compute Engine (snapshots) or Google Cloud Storage (other tiers).

Viewing subscription information

This section describes the HYCU for Google Cloud subscription information that is provided in the HYCU for Google Cloud web user interface. You can check the information

about each subscription which you can use with your user account (corresponding to each billing account that is linked from any Google Cloud project which you can access).

Accessing the Subscription Information dialog box

To access the Subscription Information dialog box, click ? in the toolbar, and then select **Subscription Information**.

The following information is displayed in the Subscription Information dialog box for the HYCU for Google Cloud subscription:

Subscription — Billing Account	Name of the billing account that is used for the subscription. Note With multiple HYCU for Google Cloud subscriptions that you can use, choose which one you want to check the information about by selecting its corresponding billing account from the drop-down menu.
Subscriber	
First name	Information about the person who subscribed to HYCU for Google Cloud.
Last name	
Company	
Notification email recipients	A list of recipients to whom notifications related to the selected HYCU for Google Cloud subscription will be sent. If this field is empty, all important notifications related to the HYCU for Google 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.
Subscription details	
Subscription ID	An identification that is automatically assigned to the subscription by Google.
Subscription plan	The plan that your HYCU for Google 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 "Backup and data retention pricing" on page 14.
Subscribed on	The date of subscribing to HYCU for Google Cloud.

Trial period until	The trial period end date, provided that your use of the service does not exceed the fee that is initially credited to you by HYCU.
Billing account details	
Billing account name	Information about the billing account that is billed for the subscription cost.
Billing account ID	
Billing account viewer	Email address of the authority (the Google Account or the Google Cloud service account) that is the billing account viewer for the HYCU for Google Cloud subscription.
Linked projects	Names and IDs of all Google Cloud projects that are linked to the billing account of the HYCU for Google Cloud subscription. This list may include projects which your user account does not have access to.

Chapter 8

Customizing HYCU for Google Cloud

After you subscribe to HYCU for Google Cloud, you can perform various tasks to customize HYCU for Google 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 the next page.

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 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

- 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. The Context Picker dialog box opens.
2. In the Context Picker dialog box, select the context.
3. Click **Save**.

   The HYCU for Google Cloud web user interface switches the context. The context that you select is remembered for the next time you sign in.

Tasks

Task	Instructions
Change roles, set the default role for users and service accounts, and delete users.	"Managing roles" on the next page

Task	Instructions
Manage HYCU for Google Cloud protection sets.	"Managing protection sets" on page 123
Import service accounts.	"Importing service accounts" on page 128
Hide instances from HYCU for Google Cloud.	"Excluding instances from synchronization by tagging the instance in Google Cloud" on page 129
Stop protecting individual projects.	"Stopping protection for individual projects" on page 129

Managing roles

A role determines the scope of actions that can be performed in the HYCU for Google 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 authority.

Considerations

- Each user that signs in to HYCU for Google Cloud or each configured service account has by default the Administrator role assigned unless set otherwise. For details on changing the default role, see "Changing the default role" on the next page.
- At least one user with the Administrator role assigned must exist in the data protection environment.
- If multiple protection sets are available in your data protection environment, a user or a service account has the same role in all protection sets within the same subscription.
- If a user or a service account has access to multiple subscriptions, they can have different roles assigned in different subscriptions. The user can also switch among these subscriptions while being signed in to HYCU for Google Cloud.

Accessing the Roles dialog box

To access the Roles dialog box, click ⚙ **Administration** in the toolbar, and then select **Roles**.

HYCU for Google Cloud roles

A user or a service account can be assigned one or more of the following roles:

Role	Allowed actions
Viewer	Acquire information about instances, applications, buckets, policies, targets, tasks, events, reports, service accounts, and protection sets in the data protection environment.
Backup Operator	Acquire the same information as Viewer, define backup strategies, and back up instances, applications, and buckets.
Restore Operator	Acquire the same information as Viewer and restore instances, applications, and buckets.
Protégé Operator	*Reserved for service accounts.* Migrate protected data from the on-premises environment to Google Cloud and the other way round by using the HYCU SpinUp functionality. For details on how to employ HYCU Protégé, see HYCU documentation.
Administrator	Perform all actions in the data protection environment.

Changing a role

Consideration

If you plan to change your own role, keep in mind that you will not be able to change it back to Administrator yourself.

Procedure

1. In the Roles dialog box, from the list of available authorities (users and service accounts), select the one to which you want to assign a different role.

   > Tip You can also search for an authority by entering its name in the Search field.

2. Click **Change Role**. The Role Change dialog box opens.
3. From the Role drop-down menu, select the role that you want to assign to the user or the service account.

   > Note You can assign multiple roles to the same user or service account if the needs of your data protection environment require it.

Changing the default role

You can at any time change the default role for users and service accounts. This means that all new users that sign in to HYCU for Google Cloud and all newly configured service accounts will automatically acquire the new default role.

Procedure

1. Click **Change Role** next to Default Role at the upper right of the Roles dialog box. The Default Role Change dialog box opens.
2. From the Role drop-down menu, select which role you want to be the default one.
3. Click **Save**.

Deleting a user

Considerations

- Deleting a user from HYCU for Google Cloud does not remove it from Google Cloud.
- You cannot delete the billing account viewer or yourself from HYCU for Google Cloud.
- Any upcoming data protection tasks related to the user that you delete will be automatically assigned to you.

Procedure

1. In the Roles dialog box, 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 dialog box opens.
3. Click **Yes** to confirm that you want the selected user to be deleted from HYCU for Google Cloud.

Managing protection sets

By default, a predefined protection set is created automatically (named default-protection-set and represented by the icon) and all the projects that are linked to the billing account of your HYCU for Google Cloud subscription are included in it. You can adjust the default setup to better suit your needs by creating additional protection sets and distributing your projects among them.

You can perform the following tasks related to protection sets:

Task	Instructions
Create a protection set and include preferred projects in it.	"Creating protection sets" on the next page
Edit an existing protection set.	"Editing protection sets" on page 125
Add a project to a protection set by using a label.	"Adding projects to a protection set by using a label" on page 125
Exclude a project from a protection set.	"Excluding projects from protection sets" on page 126

Task	Instructions
Delete a protection set that you no longer need.	"Deleting protection sets" on page 127

Consideration

All projects that are linked to the billing account of the currently selected HYCU for Google Cloud subscription are displayed.

> Tip Automatic protection set synchronization is performed every hour. However, you can update the protection set list manually at any time by clicking **Synchronize**.

Accessing the Protection Sets dialog box

To access the Protection Sets dialog box, click **Administration** in the toolbar, and then select **Protection Sets**.

Creating protection sets

If you have the required permissions granted, you can create additional protection sets that allow you to have different data protection setup for different groups of projects.

Considerations

If you move a project to a different protection set, consider the following:

- Policies will be automatically unassigned from the instances, the applications, and the buckets in the projects.
- Credential groups that were manually assigned to instances in the projects will be automatically unassigned from those instances.

Procedure

1. In the Protection Sets dialog box, from the Subscription — billing account drop-down menu, select the HYCU for Google Cloud subscription for which you want to create a new protection set.
2. Click **New**. The New Protection Set dialog box opens.
3. Enter a name for your protection set and, optionally, its description.
4. From the list of available projects, select one or more projects that you want to include in the protection set.

   > Tip You can search for a project by entering its name in the Project search field and then pressing **Enter**. By selecting the Project check box, you select all projects at once.
5. Click **Save**.

The protection set is created and added to the list of protection sets.

Editing protection sets

If you have the required permissions granted, you can change the name of a protection set, and include or exclude projects from the protection set. You can exclude the projects from the protection set also directly by following the procedure described in "Excluding projects from protection sets" on the next page.

Consideration

If you move a project to a different protection set, consider the following:

- Policies will be automatically unassigned from the instances, the applications, and the buckets in the projects.
- Credential groups that were manually assigned to instances in the projects will be automatically unassigned from those instances.

Procedure

1. In the Protection Sets dialog box, from the Subscription — billing account drop-down menu, select the HYCU for Google Cloud subscription with the protection set that you want to edit.
2. From the list of protection sets, select the one that you want to edit, and then click **Edit**.
3. Edit the selected protection set as required.
4. Click **Save**.

Adding projects to a protection set by using a label

As an alternative to adding a project to a protection set by using the HYCU for Google 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 HYCU for Google Cloud.

Procedure

In Google Cloud, attach the label to the project as the following key-value pair:

Key	Value
`hycu-protection-set`	*<ProtectionSetName>* In this case, *<ProtectionSetName>* 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.

Excluding projects from protection sets

If you have the required permissions granted, you can exclude a project from a protection set by using the HYCU for Google Cloud web user interface or by adding the `hycu-project-exclude` label to the project in Google Cloud.

I want to exclude the project...	Consideration	Instructions
By using the HYCU for Google Cloud web user interface.	When you exclude the project from the protection set other than the default one, it is automatically moved to the default protection set. If you want to completely remove the project from HYCU for Google Cloud and to stop protecting its resources, you must exclude the project from the default protection set.	"Excluding a project by using the HYCU for Google Cloud web user interface" below
By adding the `hycu-project-exclude` label to the project in Google Cloud.	When you exclude the project from the protection set, it is completely removed from HYCU for Google Cloud and its resources are no longer protected.	"Excluding a project by using a label" on the next page

Excluding a project by using the HYCU for Google Cloud web user interface

Procedure

1. In the Protection Sets dialog box, from the Subscription — billing account drop-down menu, select the HYCU for Google Cloud subscription with the protection set from which you want to exclude one or more projects.
2. *Only if the project belongs to a protection set other than the default one.* Do the following:
   a. Click ❯ next to the protection set with the project that you want to exclude. The list of all included projects is displayed.
   b. Select the project that you want to exclude from the protection set, and then click 🗑 **Remove**.
   c. Click **Yes** to confirm that you want to exclude the selected *project*.

   The excluded project is added to the default protection set.
3. *Only if you want to exclude the project from the default protection set.* Do the following:
   a. Click ❯ next to the default protection set. The list of all included projects is displayed.
   b. Select the *project* that you want to exclude from the default protection set, and

then click **Remove**.

c. Click **Yes** to confirm that you want to exclude the selected *project*.

The project is no longer included in any protection set and HYCU for Google Cloud no longer retrieves the project information from Google Cloud.

Excluding a project by using a label

Consideration

If after excluding a project from a protection set and HYCU for Google Cloud by using the `hycu-project-exclude` label, you need to add the same project to HYCU for Google Cloud again, contact HYCU Customer Support.

Procedure

In Google Cloud, add the label to the project as the following key-value pair:

Key	Value
`hycu-project-exclude`	true

After you add the label to the project, the project is no longer included in the protection set and HYCU for Google Cloud no longer retrieves the project 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 is empty with no included projects.
- The current data protection scope is set to a protection set other than the protection set that you want to delete.

Consideration

The default protection set created by HYCU for Google Cloud cannot be deleted (represented by the icon).

Procedure

1. In the Protection Sets dialog box, from the Subscription — billing account drop-down menu, select the HYCU for Google Cloud subscription that contains the protection set that you want to delete.

2. From the list of protection sets, select the one that you want to delete from HYCU for Google Cloud, and then click **Delete**.
3. Click **Yes** to confirm that you want to delete the selected protection set.

Importing service accounts

You can use a specific service account for performing all operations on a target. For details, see "Adding a bucket to HYCU for Google Cloud as a target" on page 24 and "Managing targets" on page 111.

Prerequisites

- The service account must be configured in Google Cloud and must have access to at least one of the projects linked to the selected billing account.
- The service account must be granted the Service Account User (`roles/iam.serviceAccountUser`) and Service Account Token Creator (`roles/iam.serviceAccountTokenCreator`) roles on at least one of the projects in the protection set.
- You must have access to a valid JSON file that stores the service account information, including its private key.

Consideration

Imported service accounts are available only for the currently selected HYCU for Google Cloud subscription.

> ### Accessing the Service Accounts dialog box
>
> To access the Service Accounts dialog box, click **Administration** in the toolbar, and then select **Service Accounts**.

Procedure

1. In the Service Accounts dialog box, click **Import**.
2. Click **Browse**. The Choose File to Upload dialog box opens.
3. Select the JSON file with the service account information, and then click **Open**.
4. Review the service account information, and then click **Upload**.

   The service account's email address appears in the list of service accounts.
5. Click **Close**.

You can at any time delete any service account that you no longer need from HYCU for Google Cloud by selecting it, and then clicking **Delete**. Keep in mind that deleting the service account from HYCU for Google Cloud does not remove it from Google Cloud and that you cannot delete a service account that is used as the billing account viewer of a HYCU for Google Cloud subscription.

Stopping protection for individual projects

This section provides instructions that you must follow to stop protecting individual projects in HYCU for Google Cloud.

> Note If you want to stop using HYCU for Google Cloud completely, see "Unsubscribing from HYCU for Google Cloud" on page 136.

Procedure

1. In HYCU for Google Cloud, unassign policies from all protected instances, applications, and buckets in the project. For instructions, see "Stopping service charges" on page 136.
2. In HYCU for Google Cloud, manually mark restore points of all instances, applications, and buckets in the project as expired. For instructions, see "Expiring backups manually" on page 116.
3. Exclude the project from any protection set. For instructions, see "Excluding projects from protection sets" on page 126.

When a project is no longer protected, irrelevant notifications are prevented, and the unneeded associated charges are avoided.

Excluding instances from synchronization by tagging the instance in Google Cloud

This section provides information on how to make selected instances invisible to HYCU for Google Cloud. The needs of your environment may require that some instances are not protected by HYCU for Google 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 HYCU for Google Cloud. The invisible instances cannot be assigned policies in any way.

Procedure

1. In the Google Cloud Console, choose a Google Cloud project to which the instances that you want to leave unprotected belong.
2. Within the project, choose an instance and add it the `hycu-vm-sync` custom metadata tag in Google Compute Engine. Use the following data:

Key	Value
hycu-vm-sync	false

   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.

3. Repeat step 2 for each additional instance that you want to make invisible to HYCU for Google Cloud.
4. Sign in to the HYCU for Google Cloud web user interface.
5. Select the protection set that includes the same Google Cloud project as you selected in step 1 of the procedure. For instructions on selecting protection sets in HYCU for Google Cloud, see "Selecting a HYCU for Google Cloud protection set" on page 22.
6. In the navigation pane, click **Instances**.
7. Click **Synchronize** 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 9

Troubleshooting

If you encounter a problem while using HYCU for Google 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 HYCU for Google 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 92 and "Checking task statuses" on page 91.
   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 Customer Support. It is recommended that you collect and send the following information to HYCU Customer Support:
   - Description of your data protection environment
   - Description of your problem
   - Results of any testing you have done (if available)

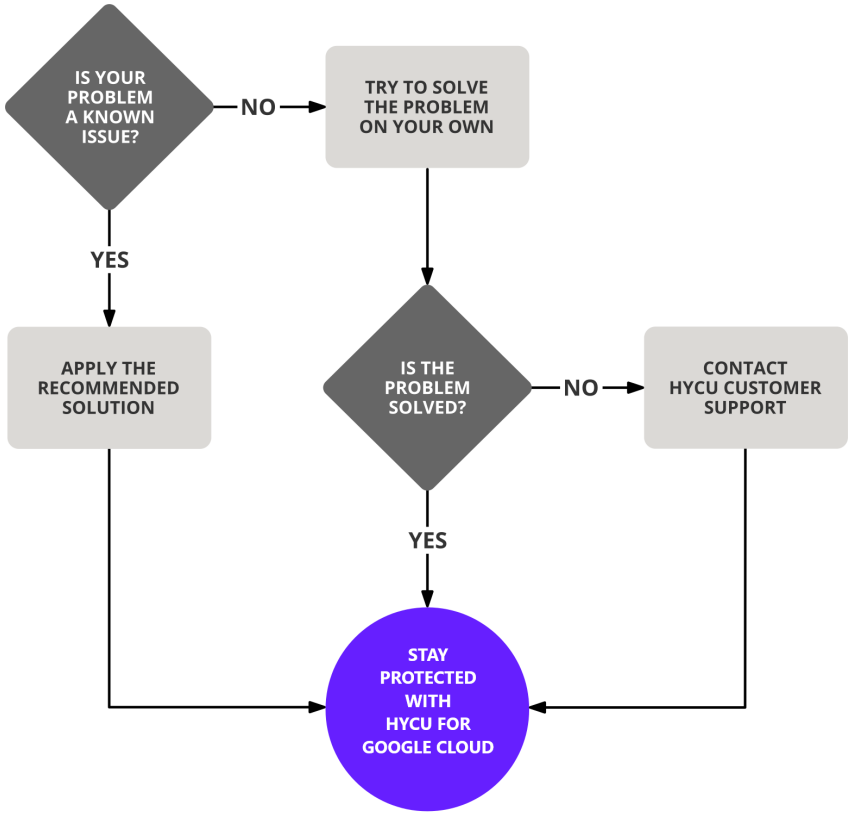

Figure 9–1: Major steps of the troubleshooting process

Known problems and solutions

This section lists all known problems that you may encounter while using HYCU for Google Cloud, along with their solutions.

Missing Google Cloud projects

Problem

When configuring protection sets in HYCU for Google Cloud, not all of your Google Cloud projects are listed in the Protection Sets dialog box. Switching to a billing account of another HYCU for Google Cloud subscription does not show the missing projects.

Cause

The missing projects are not linked to any Google Cloud billing account that was selected when subscribing to the service.

Solution

To solve this problem, do one of the following:

- In Google Cloud, link the missing projects to a billing account that was selected for the HYCU for Google Cloud subscription. For instructions, see Google Cloud documentation.
- Subscribe to HYCU for Google Cloud again and select the billing account to which your missing projects are linked.

Inability to set up manually created targets

Problem

When you try to add a manually created target, HYCU for Google Cloud reports that the target is inaccessible.

Solution

In the Google Cloud Storage service, grant your Google Account the Storage Admin (`roles/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 HYCU for Google Cloud" on page 17.

Policy assignment 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 HYCU for Google 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 HYCU for Google 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

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 (`roles/compute.admin`) role on the Google Cloud project.

For information on the required roles for general use of the service, see "Signing in to HYCU for Google Cloud" on page 17.

Task progress indicator remains at 0%

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.

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 objects have not been restored.

Cause

The original volume no longer exists, or the credential group that is assigned to the original instance in HYCU for Google 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 target, or update the configuration of the credential group that is assigned to the original instance in HYCU for Google Cloud.

Restore of individual files fails

Problem

The restore of individual files to the original instance fails because of unsuccessful mounting of the original disk.

Cause

HYCU cannot connect to the original instance because no credential group is assigned to the instance in HYCU for Google Cloud or the credential group contains incorrect settings.

Solution

Assign a credential group to the instance or make the necessary adjustments to the credential group configuration. For instructions, see "Configuring and assigning credential groups manually" on page 37.

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 HYCU for Google Cloud, only the currently selected protection set is available in the Protection Set Picker dialog box. After your web user interface session ends, you are unable to sign in again.

Solution

Contact HYCU Customer Support.

Instance backup option reconfiguration fails

Problem

After you enable the restore of individual files in the Instance Configuration dialog box for an instance running Microsoft Windows, automatic assignment of a credential group to the instance fails. HYCU for Google Cloud is therefore unable to update the configuration of the instance backup options.

Solution

Manually create a credential group and assign it to the instance, and then retry updating its configuration. For instructions on manual credential group assignment, see "Enabling access to data" on page 36.

Chapter 10

Unsubscribing from HYCU for Google Cloud

If for whatever reason you decide that you no longer want to use HYCU for Google Cloud for protecting your data, you can easily unsubscribe from the service.

Unsubscribing from HYCU for Google Cloud includes the following tasks:

Task	Instructions
1. Stop being charged for using HYCU for Google Cloud.	"Stopping service charges" below
2. Prevent HYCU for Google Cloud to access your Google Account.	"Preventing account access" on page 138
3. *Optional.* Remove the HYCU Managed Service Account permissions.	"Removing the HYCU Managed Service Account permissions" on page 138
4. Cancel your HYCU for Google Cloud subscription in Google Cloud.	"Canceling your HYCU for Google Cloud subscription" on page 139

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.	In HYCU for Google Cloud, unassign policies from all protected instances, applications, and buckets: • To unassign policies from instances: 1. In the navigation pane, click **Instances**. 2. Select all instances with assigned policies, and then click **Policies**. 3. Click **Unassign**, and then click **Yes** to confirm that you want to unassign the policies from the selected instances.

	• To unassign policies from applications: 1. In the navigation pane, click **Applications**. 2. Select all applications with assigned policies, and then click **Policies**. 3. Click **Unassign**, and then click **Yes** to confirm that you want to unassign the policies from the selected applications. • To unassign policies from buckets: 1. In the navigation pane, click **Buckets**. 2. Select all buckets with assigned policies, and then click **Policies**. 3. Click **Unassign**, and then click **Yes** to confirm that you want to unassign the policies from the selected buckets. ⚠ Important *Only if multiple protection sets are available in your data protection environment.* Make sure to follow these steps for each protection set separately.
2. Stop charges for backup data storage.	1. Manually mark restore points of all instances, applications, and buckets as expired. For instructions, see "Expiring backups manually" on page 116. ⚠ Important *Only if multiple protection sets are available in your data protection environment.* Make sure to do this for each protection set separately. 2. *Only if SAP HANA application data was backed up using the Backint agent.* Disable log backups and remove all existing log backups from the Google Cloud Storage buckets from the following location: `<SAPHANAAppName>/usr/sap/<SAPHANAAppName>/SYS/global/hdb/backint/<DatabaseName>` For details on how to disable log backups, see SAP HANA documentation. 3. Remove all backup data created by HYCU for Google Cloud from Google 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 "Objects created by the service" on page 140. For instructions on how to delete

	targets and remove backup data from targets, see Google Cloud documentation. 4. Remove all snapshots created by HYCU for Google Cloud from Google Cloud. For the snapshot naming conventions, see "Objects created by the service" on page 140. For instructions on how to remove snapshots, see Google Cloud documentation.

Preventing account access

When you subscribed to HYCU for Google 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 for Google 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 HYCU for Google 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:

Method	Instructions
	Open Google Cloud Shell

> ⚠ Important If you remove the HMSA permissions by using either of these methods, keep in mind that to add the HMSA back to HYCU for Google Cloud, you will have to grant the HMSA the following roles in Google Cloud on each project 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.

Canceling your HYCU for Google Cloud subscription

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 of the HYCU for Google Cloud subscription.
- Your currently selected project in the Google Cloud Console is linked to the billing account of the HYCU for Google Cloud subscription.

Procedure

1. Open a web browser and go to the HYCU | Marketplace - Google Cloud webpage.
2. Click **Cancel service**.
3. In the Cancel HYCU subscription dialog box, click **Cancel auto-renewal** to confirm your choice.

After you cancel your HYCU for Google 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 HYCU for Google Cloud, resubscribe to HYCU for Google Cloud and specify the billing account of the canceled subscription.

Appendix A

Objects created by the service

During data protection tasks, HYCU for Google Cloud creates temporary and persistent HYCU objects in your Google Cloud projects. Temporary HYCU objects exist only for the duration of a task, and persistent HYCU objects 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 objects.

Names or location path templates of persistent HYCU objects created during backup tasks

- Snapshot:

  `hycu-snap-<TaskUUID>-<DiskName>`

- Automatically created target:

  `hycu-<CloudStorageRegionName>-<UUID>`

- Target folder with a backup, a backup copy, or a data archive:

  `hycu/backups/<ProjectName>/<ZoneName>/<InstanceName>/disks/<DiskName>/<StorageClass>`

- Target folder with a disk catalog, instance metadata, and instance disk metadata:

  `hycu/backups/<ProjectName>/<ZoneName><InstanceName>/tasks/<TaskUUID>`

Names or location path templates of persistent HYCU objects 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):

  `<OriginalFileName>.hycu.restored[.<OriginalFileExtension>]`

- Target folder with restored files or folders:

  `hycu/restores/<ProjectName>/<ZoneName>/<InstanceName>/<TaskUUID>/<DiskName>/<VolumeName>/<PathName>`

- Restored file:

  `<FileName>.<FileExtension>.<TimeStamp>.restored`

- External IP address resource automatically allocated by HYCU for Google Cloud during cloning:

  `hycu-static-external-<UUID>`

- Internal IP address resource automatically allocated by HYCU for Google Cloud during cloning:

  `hycu-static-internal-<UUID>`

- Cloned disk:

  `hycu-disk-<TaskUUID>-<UUID>-<DiskName>`

- Exported disk:

  `hycu-disk-<TaskUUID>-<UUID>-<DiskName>`

Name templates of temporary HYCU objects created during backup and restore tasks

- Temporary disk:

  `hycu-disk-tmp-<TaskUUID>-<OriginalDiskName>`

Appendix B

Deploying a HYCU backup controller

If you are employing HYCU Protégé, you can use the HYCU for Google Cloud web user interface to deploy a HYCU backup controller instance to Google Cloud in the event of a disaster in the on-premises data protection environment.

For details on the supported on-premises infrastructures and how to employ HYCU Protégé, see HYCU for Enterprise Clouds documentation.

Prerequisites

- You must own the HYCU and HYCU Protégé licenses. For details on how to obtain these licenses, see HYCU for Enterprise Clouds 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 8 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** on the toolbar, and then select **HYCU Controller Deployment**.

Procedure

1. From the Subscription — billing account drop-down menu, select the Google subscription that is connected with HYCU for Google Cloud and to which you want to deploy the HYCU backup controller.
2. From the Project 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 box, 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 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 exactly corresponds to the specified values, the list is empty and you need to adjust the values in the vCPU and Memory fields.

10. Under Network interface, 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/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 Target networks drop-down menu, select the target network.

         > Note The list of available target networks includes only the ones within the region you selected for the HYCU backup controller.

      b. In the External address type field, select the external IP address for the network interface. You can select among the following options:

Option	Description
None	The network interface does not use an external IP address. This option is preselected if the network interface of the

	original instance did not use an external IP address.
Ephemeral	The network interface uses an automatically allocated external IP address. This option is preselected if the network interface of the original instance used an external IP address.
Static (Reserved)	The network interface uses a static external IP address that was reserved in Google Compute Engine in advance.
Static (New)	The network interface uses a static external IP address that is allocated at the time of the deployment. If the allocation fails, the instance is assigned a temporary external IP address. Such a fallback also sets the deployment task status to Done with errors.

c. In the Internal address type field, select the internal IP address for the network interface. You can select between the following options:

Option	Description
Ephemeral (Automatic)	The network interface uses an automatically allocated internal IP address. This option is selected by default for the preselected network interfaces.
Ephemeral (Custom)	The network interface uses an internal IP address that is defined by you. ⚠ Important Use of this option might result in IP address conflicts.
Static (Reserved)	*Not available for legacy networks.* The network interface uses a static internal IP address that was reserved in Google Compute Engine in advance.
Static (New)	*Not available for legacy networks.* The network interface uses a new static internal IP address that is defined by you. Note 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** or **Save**.

- 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.

11. Click **Next**.
12. From the Available versions drop-down menu, select the version of the HYCU backup controller. By default, the latest version is selected.
13. 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.

Appendix C

Bulk restore specifications

Based on the bulk restore options you specify when restoring multiple instances, HYCU for Google Cloud generates a bulk restore specification. If you choose to further modify the bulk restore specification instead of running the restore immediately, a dialog box opens, containing the POST method URL and the generated restore specification.

You can edit the specification directly or copy it to the clipboard by clicking **Copy to clipboard** and then further editing it with a text editor of your choice. After you update the specification, click **Execute** to run the restore specification.

> Note The bulk restore specification API is not available through the REST API Explorer.

Elements of a bulk restore specification

The basic elements of a bulk restore specification include the type of the specification (`type`), a flag whether to overwrite existing items (`overwriteExisting`), and the items to restore.

Syntax

```
{
  "type": "VMS" | "DISKS",
  "overwriteExisting": false | true,
  "items": [
    {
      "source":
      {
        "path": "<path>",
        "disks": [<disk>,...]
      },
      "destination":
      {
        "path": "<path>",
        "disks": [<disk>, ...],
        "networkInterfaces": [<networkInterface>, ...],
        "metadata": {},
        "labels": {},
```

```
        "tags": []
      },
    },
    ...
  ],
}
```

Basic elements

- `type: "VMS" | "DISKS"`

  The restore type, `VMS` for instances and `DISKS` for disks.

- `overwriteExisting: false | true`

  If set to `true`, instances in the target 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:

  - `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 (`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 target 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:

  - `path`

    The path to the network device.

  - `externalIpType`

    The external IP type for the network interface. 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. 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).

Appendix D

Least-privilege permissions used by HYCU for Google Cloud

To perform data protection tasks, HYCU for Google Cloud uses the permissions that you granted to the Google Account or Google Service Account or the HMSA in Google Cloud. If the needs of your data protection environment require you to create a custom role, you can use the HYCU for Google Cloud role template with a predefined set of least-privilege permissions to grant the required permissions to the created role.

Using a role template with a predefined set of permissions

Prerequisite

Your account has 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 HYCU for Google 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

2. Create the role and grant it the following permissions by running 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.

Permissions required by HYCU for Google Cloud

The following is a list of permissions required by HYCU for Google Cloud:

Service	Permissions
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 compute.instances.create compute.instances.delete compute.instances.deleteAccessConfig compute.instances.detachDisk compute.instances.get

	compute.instances.getSerialPortOutput
	compute.instances.list
	compute.instances.setLabels
	compute.instances.setMachineType
	compute.instances.setMetadata
	compute.instances.setServiceAccount
	compute.instances.setTags
	compute.instances.start
	compute.instances.stop
	compute.instances.update
	compute.machineImages.useReadOnly
	compute.machineTypes.get
	compute.machineTypes.list
	compute.networks.get
	compute.networks.list
	compute.networks.updatePolicy
	compute.networks.use
	compute.networks.useExternalIp
	compute.projects.get
	compute.regionOperations.get
	compute.regions.get
	compute.regions.list
	compute.snapshots.create
	compute.snapshots.delete
	compute.snapshots.get
	compute.snapshots.list
	compute.snapshots.setLabels
	compute.snapshots.useReadOnly
	compute.subnetworks.get
	compute.subnetworks.list
	compute.subnetworks.use
	compute.subnetworks.useExternalIp

	compute.zoneOperations.get compute.zones.get compute.zones.list
Google Kubernetes Engine	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 resourcemanager.projects.get
Google Cloud Storage	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 storage.objects.update

Provide feedback

For any suggestions and comments regarding this product or its documentation, send us an e-mail to:

info@hycu.com

We will be glad to hear from you!

www.hycu.com

HYCU®