




HYCU Protégé



User Guide

June 2023

Legal notices

Copyright notice

© 2023 HYCU. All rights reserved.

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

Trademarks

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

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

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

1 About HYCU Protégé	12
Key features and benefits	12
Data protection environment overview	14
HYCU Protégé data protection	15
2 Starting with HYCU Protégé	16
Service pricing	16
Backup and data retention pricing	18
Subscribing to the service	19
Subscribing from Google Cloud Marketplace	20
Subscribing from AWS Marketplace	21
Signing in to HYCU Protégé	22
3 Establishing a data protection environment	24
Switching the user interface context	25
Selecting a HYCU Protégé protection set	26
Setting up targets	27
Adding a bucket to HYCU Protégé as a target	29
Defining your backup strategy	31
Taking advantage of predefined policies	31
Creating custom policies	32
Creating backup windows	37
Creating data archives	39
Setting default policies	41
Setting up automatic policy assignment	42
Enabling access to data	44

Manually enabling access to data	45
4 Protecting SaaS applications	51
Configuring SaaS application backup options	52
Backing up SaaS applications	53
Restoring SaaS applications	54
5 Protecting Google Cloud applications	56
Protecting SAP HANA applications	57
Preparing for SAP HANA application protection	57
Backing up SAP HANA applications	60
Restoring SAP HANA applications	62
Protecting Google Kubernetes Engine applications	64
Preparing for Google Kubernetes Engine application protection	64
Backing up Google Kubernetes Engine applications	69
Restoring Google Kubernetes Engine applications	70
6 Protecting instances	76
Planning instance protection	76
Preparing your data protection environment	76
Configuring instance backup options	78
Backing up instances	82
Restoring instances	86
Restoring an instance	88
Cloning an instance	89
Moving an instance	99
Restoring disks	108
Cloning disks	109
Moving disks	112

Restoring multiple instances in a single session	114
Restoring multiple disks in a single session	116
Restoring multiple instances or disks from a JSON file	118
Restoring individual files or folders	118
Restoring files or folders to an instance	118
Restoring files or folders to a target	124
7 Protecting buckets	126
Configuring bucket backup options	127
Backing up buckets	130
Restoring buckets	131
8 Performing daily tasks	136
Using the HYCU Protégé dashboard	137
Viewing entity details	138
Viewing the backup status of entities	141
Tier statuses	143
Managing policies	143
Viewing policy information	144
Creating a policy	144
Editing a policy	145
Deleting a policy	145
Managing targets	145
Viewing target information	146
Editing targets	148
Deactivating and activating targets	148
Removing targets	149
Checking task statuses	149

Viewing events	150
Configuring event notifications	152
Creating email notifications	152
Creating webhook notifications	153
Using HYCU Protégé reports	155
Getting started with reporting	156
Viewing reports	158
Generating reports	158
Scheduling reports	159
Exporting and importing reports	160
Filtering and sorting data in panels	161
Filtering data in panels	161
Filtering options in the SaaS panel	162
Filtering options in the Applications panel	163
Filtering options in the Instances panel	164
Filtering options in the Buckets panel	165
Filtering options in the Policies panel	166
Filtering options in the Targets panel	166
Filtering options in the Tasks panel	167
Filtering options in the Events panel	168
Filtering options in the IAM panel	169
Sorting data in panels	169
Performing manual backups	169
Expiring backups manually	170
Exporting the contents of the panel	172
Viewing subscription information	173
9 Customizing HYCU Protégé	175

Managing sources	176
Managing AWS accounts	177
Managing Google Cloud projects	178
Managing Protégé SaaS modules	180
Discovering services	182
Configuring service discovery	182
Exploring R-Graph	183
Managing identity and access	187
Managing identity providers	188
Managing users	190
Managing roles	192
Managing protection sets	193
Creating protection sets	194
Editing protection sets	195
Deleting protection sets	197
Importing service accounts	198
Stopping protection for individual sources	199
Excluding instances from synchronization by tagging the instance in AWS or Google Cloud	199
10 Troubleshooting	202
Known problems and solutions	203
Missing Google Cloud projects	203
Inability to set up manually created Google Cloud targets	203
Assigning a policy to a Google Cloud instance fails	204
Snapshot creation fails for instances in a specific Google Cloud project	204
Task progress indicator remains at 0% during the backup of a Google Cloud instance	205

Restore of individual files ends with errors or fails	205
Restore of individual files fails	206
Inability to change the protection set or to sign in	206
Instance backup option reconfiguration fails	206
Problem with sorting data in the Events panel	207
11 Unsubscribing from HYCU Protégé	208
Stopping service charges	208
Preventing account access	210
Preventing access to an AWS account	211
Preventing access to a Google Cloud account	211
Removing the HYCU Managed Service Account permissions	212
Canceling your HYCU Protégé subscription	213
Cancelling the HYCU Protégé subscription in the AWS Marketplace	213
Cancelling the HYCU Protégé subscription in the Google Cloud Marketplace	214
A Objects created by HYCU Protégé	215
B Bulk restore specifications	218
Elements of a bulk restore specification	218
C Least-privilege permissions used by HYCU Protégé	222
Using a role template for AWS	222
AWS permissions required by HYCU Protégé	223
Using a role template for Google Cloud	226
Google Cloud permissions required by HYCU Protégé	227
D Deploying a HYCU backup controller	231
Deploying a HYCU backup controller to AWS	231

Accessing the HYCU web user interface	235
Deploying a HYCU backup controller to Google Cloud	235
Accessing the HYCU web user interface	239

Chapter 1

About HYCU Protégé

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

The following are the key elements of HYCU Protégé:

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

Key features and benefits

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

- **Protection against data loss**

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

- **Data protection in a few minutes**

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

- **SaaS discovery and protection**

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

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

- **Application discovery and protection**

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

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

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

- **Scheduled backups**

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

- **Centralized data protection management and monitoring**

You can join your cloud projects or accounts into protection sets to establish centralized data protection management and monitoring.

- **Lower impact on the environment**

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

- **Use of data archives**

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

- **Restore of individual files**

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

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

The HYCU Protégé dashboard helps you to identify potential problems and bottlenecks to improve the performance of your data protection environment.

- **Optimized consumption of storage space**

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

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

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

Data protection environment overview

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

Term	Description
HYCU Protégé web user interface	An interface for protecting entities and administering the data protection environment.
Sources	Environments for which HYCU Protégé provides data protection—Google Cloud projects, AWS accounts, and SaaS modules.
Protection sets	Groups that join together sources that you have successfully added to HYCU Protégé.
Entities	Objects to which you can assign a policy and for which you therefore provide data protection—SaaS applications, SAP HANA and GKE applications, Amazon EC2 and Google Cloud instances, and Amazon S3 and Google Cloud Storage buckets. Data is always protected at a granular level, allowing you to restore either the whole entities or their parts.

Term	Description
Targets	Buckets that HYCU Protégé uses for storing backup data. Backup data can also be stored as snapshots.

HYCU Protégé data protection

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

The AWS accounts, Google Cloud projects, and the Protégé modules for SaaS applications that you add as sources define the scope of data protection.

HYCU Protégé enables you to protect instances, applications, data in buckets, and SaaS application data. 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 Protégé

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

Task	Instructions
Getting familiar with HYCU Protégé pricing concepts	“Service pricing” below
Subscribing to HYCU Protégé	“Subscribing to the service” on page 19
Signing in to the HYCU Protégé web user interface	“Signing in to HYCU Protégé” on page 22

Service pricing

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

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




Figure 2-1: Data protection cost

Cost	Details
Backup and data retention	Cost of backing up data and data retention. For details, see “Backup and data retention pricing” on

Cost	Details
	page 18.
Backup data storage	<p>Cost of storing backup data. The following factors are considered:</p> <ul style="list-style-type: none"> • Target type (a snapshot or a bucket) • Backup frequency • Size of backup data • Backup retention period <p>If you use a bucket as a target, the following is also considered:</p> <ul style="list-style-type: none"> • 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	<p>Cost of temporary resources that HYCU Protégé creates in the cloud when performing the following tasks:</p> <ul style="list-style-type: none"> • 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 Protégé subscription includes a 14-day free trial period. During this time, HYCU does not charge you for the backup and data retention cost. The cost of backup data storage and temporary resources is charged by your cloud provider as usual.

Backup and data retention pricing

The HYCU Protégé 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 Protégé. For details on the subscription plans, see [“HYCU Protégé 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 Protégé 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 31.](#)

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

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

Considerations

- If an instance, an application, or a bucket is deleted from the cloud, but it still has at least one valid restore point available, it is considered protected (its status is PROTECTED_DELETED) and HYCU automatically associates such an entity with the bronze pricing tier. In the case of instances, it charges you for protecting only the included disks.
- If you unassign a policy from an instance, an application, or a bucket that still has at least one valid restore point available, such an entity is

considered protected and HYCU automatically associates it with the bronze pricing tier. In the case of instances, it charges you for protecting only the included disks.

- *Applicable for instances and applications running on them.* If you assign policies to an instance and an application running on the same instance, keep in mind that you will be charged for both protecting the instance and protecting the application.

HYCU Protégé subscription plans

HYCU Protégé offers you the following subscription plans:

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

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

Subscribing to the service

You subscribe to HYCU Protégé online from your cloud provider marketplace and HYCU then automatically activates the service for you. This is usually done by one user for an entire organization.

Service	Description
Google Cloud	“Subscribing from Google Cloud Marketplace” on the next page
Amazon Web Services	“Subscribing from AWS Marketplace” on page 21

Subscribing from Google Cloud Marketplace

Prerequisite

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

Consideration

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

Procedure

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

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.

Subscribing from AWS Marketplace

Prerequisites

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

For details, see AWS documentation.

Consideration

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

Procedure

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

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 Protégé

After successfully subscribing to HYCU Protégé, you can sign in to the HYCU Protégé web user interface.

Prerequisites

- You are using a supported web browser. For a list of supported web browsers, see the *HYCU Protégé Compatibility Matrix*.

Only if you want to protect a Google Cloud project:

- 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 Protégé 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.

For details, see Google and Google Cloud documentation.

Procedure

1. Open a web browser and go to the HYCU Protégé webpage, by using the link you received when you subscribed to HYCU Protégé.


Alternatively, open the [HYCU Protégé](#) webpage and enter the HYCU account ID you received when you subscribed to HYCU Protégé.

 **Tip** You can set a sign-in alias for your HYCU account. For details, see [“Viewing subscription information” on page 173](#).

2. Click **Next**.
3. On the sign-in webpage, depending on how you want to sign-in to HYCU Protégé, do one of the following:
 - *By using dedicated sign-in credentials for HYCU.* Enter your sign-in name and password.
 - *By using an identity provider.* Click the preferred identity provider, and then, if required, enter your credentials.

For details on how to integrate HYCU Protégé with identity providers, see [“Managing identity providers” on page 188](#).

After you sign in to the HYCU Protégé 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 Protégé 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 Protégé, 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 192](#).

If you have the Administrator role assigned, you can switch between the Subscription and Protection set contexts. Depending on the scope of the tasks that you want to perform, click ▼ next to the name of the currently selected context to switch to another one. The Subscription context enables you to perform administration tasks related to the selected subscription such as adding identity providers, adding or removing users, and changing roles, whereas the Protection set context enables you to perform data protection tasks related to the selected protection set. See [“Managing identity and access” on page 187](#) and [“Managing protection sets” on page 193](#) for details.

Tasks

Establishing a data protection environment involves the following tasks:

Task	Instructions
1. Add the sources (SaaS applications, accounts, or projects) to HYCU Protégé.	“Managing sources” on page 176
2. <i>Only if you plan to use multiple protection sets.</i> Configure a protection set and select it.	“Managing protection sets” on page 193 and “Selecting a HYCU Protégé protection set” on page 26
3. <i>Only if you plan to use manually created</i>	“Setting up targets” on

Task	Instructions
<i>targets</i> . Add Google Cloud Storage or Amazon S3 buckets to HYCU Protégé as targets.	page 27
4. Decide for predefined policies or create custom ones.	“Defining your backup strategy” on page 31
5. <i>Required only in special data protection scenarios</i> . Configure credential groups and assign them to instances.	“Enabling access to data” on page 44

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

Switching the user interface context

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

- **Subscription**


In Subscription context, only the IAM panel is active. Use this context to perform administration tasks related to your subscription, such as adding identity providers, adding or removing users, or changing roles. See [“Managing identity and access” on page 187](#).

- **Protection set**

In the protection set context, you select the scope of data protection by selecting a specific protection set.

The HYCU Protégé web user interface switches the context to the selected scope of data protection. See [“Managing protection sets” on page 193](#).

Procedure

1. On the toolbar, click  next to the name of the selected protection set or subscription.
2. From the drop-down menu, select the context:
 - **Subscription**

- **Protection Set**

The Select Protection Set dialog box opens. Select the protection set and click **Confirm**.

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

Selecting a HYCU Protégé protection set

An environment for which HYCU Protégé provides data protection consists of one or more protection sets that join together sources—Google Cloud projects, AWS accounts, and Protégé SaaS modules. When you subscribe to HYCU Protégé, a default protection set is created automatically.

Depending on your business needs, you can create additional protection sets and distribute your projects or accounts 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 193](#).

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

- *For Google Cloud:* Regardless of your protection set configuration, you can see only projects linked to the billing account that was selected when subscribing to HYCU Protégé 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. From the list of available protection sets, select the scope of your data protection by selecting the preferred protection set.
3. Click **Confirm**.

The HYCU Protégé 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 Protégé allows you to define either a bucket or a snapshot as a location for storing your data.

Target	Description
Bucket	<p>Backup data is stored in Google Cloud Storage or Amazon S3 buckets that you create yourself or HYCU Protégé creates for you automatically:</p> <ul style="list-style-type: none"> • Manually created targets <p>You can create your own buckets in Google Cloud Storage or Amazon S3 and add them to HYCU Protégé as targets. For instructions, see “Adding a bucket to HYCU Protégé as a target” on page 29.</p> • Automatically created targets <p><i>Applicable only if you are protecting instances or Google Kubernetes Engine applications.</i> HYCU Protégé creates Google Cloud Storage or Amazon S3 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 or AWS account and at the same location as the instances you are backing up or the GKE clusters on which the applications you are backing up are deployed.</p> <p>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).</p> <p>For the target naming conventions, see “Objects created by HYCU Protégé” on page 215.</p> <p>⚠ Caution Never delete any targets used by HYCU Protégé because this may result in data loss.</p>

Target	Description
	<p>Additionally, within targets, ensure that the <code>hycu/backups/</code> folders are always kept intact.</p>
Snapshot	<p><i>Available only if you are protecting instances, Google Kubernetes Engine applications using persistent volumes, or SaaS applications.</i> Backup data is stored as a snapshot in the Google Cloud project, the AWS account that contains the instances you want to protect, the staging targets for SaaS applications, or in the clusters on which the applications you want to protect are deployed.</p> <p>Note If snapshots created by HYCU Protégé are deleted from Google Cloud or AWS, you will not be able to restore backup data from this location. However, you can still restore your data from targets if copies of backup data or data archives exist.</p> <p>For the snapshot naming conventions, see “Objects created by HYCU Protégé” on page 215.</p>

Staging targets for SaaS backup

Staging targets are targets used when backing up SaaS application data. They serve as the temporary staging area for application data during the backup process. A target becomes a staging target when you select it during source configuration. See [“Managing sources” on page 176.](#)

Targets are valid staging targets, if they meet the following requirements:

- Only Google Cloud Storage buckets are supported.
- Targets that are attached to policies cannot be used as staging targets.
- Targets with Object Lock (WORM) enabled and data archives cannot be used as staging targets.

Keep in mind the following:

- If a target from another protection set is selected as a staging target, it will be moved to the active protection set.
- If a target is selected as a staging target, it cannot be changed to a non-staging target.

Adding a bucket to HYCU Protégé as a target

Prerequisites

- *For AWS:*
 - The AWS account where the bucket resides must be added to HYCU Protégé as a source. For instructions on how to add accounts as sources, see [“Managing sources” on page 176](#).
 - *For adding a bucket with Object Lock (WORM) enabled:* The Object Lock option must be enabled for the bucket. For details on how to configure buckets, see AWS documentation.
- *For Google Cloud:*
 - 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 198](#).

Limitations


- Publicly available buckets cannot be added as targets.
- *For AWS:* Only copies of backup data can be stored to a bucket with the S3 Glacier Flexible Retrieval or S3 Glacier Deep Archive storage tier. Keep in mind that AWS can charge you additionally for premature removal of data if the retention period specified in your policy is shorter than the recommended (minimum) retention period in AWS.
- *Only if protecting buckets in Google Cloud:* 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 198](#).

Considerations

- You can set up the same target in multiple protection sets.
- *For Google Cloud:*


- The exclude policy is automatically assigned to the bucket that is added to HYCU Protégé 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. Select **Amazon S3** or **Google Cloud Storage**, and then click **Next**.
3. In the Target field, enter the name of the bucket that you want to add to HYCU Protégé as a target.
4. 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.

5. Specify the account that you want to be used for performing all operations on this target:
 - *For AWS:* In the Account ID field, enter the AWS account ID of that target.
 - *For Google Cloud:* From the Cloud Account drop-down menu, select the service account that has access to this target. If None is selected, the HMSA will be used.
6. *Only if you are adding an AWS target.* From the Storage class drop-down, select the storage class of the objects that are uploaded during backup or copy.
7. Click **Save**.

The target is added to the list of targets in the Targets panel. For details on managing targets, see [“Managing targets” on page 145](#).

Defining your backup strategy

HYCU Protégé 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” below.
Creating a custom policy	If none of the predefined policies meets the needs of your environment, you can create a new policy and tailor it to your needs. For details, see “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 41.

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 Protégé comes with the following predefined policies:

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

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

Consideration

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

Creating custom policies

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

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

Policy option	Allows you to...
Backup Window	Start all backup tasks within specified time frames to improve efficiency and avoid an overload of your environment. For details, see “Creating backup windows” on page 37 .
Copy ^{ab}	Create a copy of backup data.
Archiving ^a	Preserve your data for future reference. For details, see “Creating data archives” on page 39 .
Labels ^b	Set up automatic policy assignment based on the labels or

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

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

^b This policy option is not available for all SaaS applications. For more information, see the [SaaS application guides](#).

Prerequisites

- *Only if you plan to select a manually created target.* A bucket must be added to HYCU Protégé as a target. For instructions, see [“Setting up targets” on page 27](#).
- *Only if you plan to enable the Backup Window policy option.* A backup window must exist for the selected HYCU Protégé protection set. For instructions, see [“Creating backup windows” on page 37](#).
- *Only if you plan to enable the Archiving policy option.* A data archive must exist for the selected HYCU Protégé protection set. For instructions, see [“Creating data archives” on page 39](#).
- *Only if you plan to enable the Labels policy option.*
 - *Google Cloud specifics:* The HYCU Managed Service Account (HMSA) must have the following roles granted on the projects with the instances that you plan to protect, the clusters on which the GKE applications that you plan to protect are deployed, or the buckets that you plan to protect:
 - Compute Admin (`roles/compute.admin`)
 - Service Account User (`roles/iam.serviceAccountUser`)
 - Storage Admin (`roles/storage.admin`)
 - *Required only if protecting GKE applications.* Kubernetes Engine Admin (`roles/container.admin`)

For instructions on how to grant permissions to service accounts, see [Google Cloud documentation](#).

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


Cloud Storage or Amazon S3 as bucket labels.

For instructions on how to do this, see the [SaaS application guides](#), or the Kubernetes, AWS, or Google Cloud documentation.


Considerations

- HYCU Protégé automatically associates the resource with one of the pricing tiers based on the value of the Backup every option that you set in the policy. However, if you are storing data as a snapshot and have enabled the Archiving option, the pricing tier is automatically set to bronze regardless of the specified RPO.
- If you want your data to be stored as a snapshot and on a target, make sure to select the Snapshot backup target type and also enable the Copy policy option.
- *Only if you plan to enable the Labels policy option.*
 - Labels that you specify in policies in HYCU Protégé must be unique within the selected protection set.
 - When matched, the `hycu-policy` custom metadata tag takes precedence over other labels or tags that might be added to the same SaaS application, to the same application in Google Kubernetes Engine, to the same instance in Google Compute Engine or Amazon EC2, or to the same bucket in Google Cloud Storage or Amazon S3. For more information on the `hycu-policy` tag, see [“Setting up automatic policy assignment” on page 42](#).
- *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.
- *Only if you plan to back up SaaS application data.* SaaS application data cannot be stored to the automatically created targets. Therefore, the policy that you plan to assign to SaaS applications and/or resources cannot have the Automatically selected option specified for the Target.

Accessing the Policies panel

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

Procedure


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

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

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

^b This policy option is not available for all SaaS applications. For more information, see the [SaaS application guides](#).

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 SaaS applications, GKE applications using persistent volumes, or instances.* **Snapshot**

Google Cloud only. 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 Protégé creates a bucket in the region of the Kubernetes cluster or the instance 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 GKE application data or instance data (and not SaaS application data, SAP HANA application data, or bucket data).

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

Policy option	Instructions
Backup Window	<p>In the Backup Window section, from the Backup window drop-down menu, select a backup window for backup tasks.</p> <p>If you do not select a backup window, the Always value is shown, which means that your backups are allowed to run at any time.</p>
Copy ^{ab}	<p>In the Copy section, do the following:</p> <ol style="list-style-type: none"> Set a retention period (in months, weeks, or days) for the copy of backup data. From the Target drop-down menu, select a target that you want to use for storing data. <p>If you want the target to be selected automatically,</p>

Policy option	Instructions
	<p>make sure the Automatically selected option is selected. In this case, HYCU Protégé creates a bucket in the region of the Kubernetes cluster or the instance 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.</p> <p>ⓘ Important Automatically created targets can be selected only if you plan to protect GKE application data or instance data (and not SaaS application data, SAP HANA application data, or bucket data).</p>
Archiving ^a	In the Archiving section, from the Data archive drop-down menu, select a data archive.
Labels ^b	<p>In the Labels section, enter a label key and value, and then click Add. If required, repeat the action as appropriate.</p> <p>For details on automatic policy assignment, see “Setting up automatic policy assignment” on page 42.</p>

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

^b This policy option is not available for all SaaS applications. For more information, see the [SaaS application guides](#).

6. Click **Save**.

The policy is created and added to the list of policies. For details on managing policies, see [“Managing policies” on page 143](#).

Creating backup windows


HYCU Protégé enables you to define time frames when backup tasks are allowed to start. If you use a backup window, the backup tasks are started only within the hours you specify, which improves effectiveness and prevents overloading your data protection environment. For example, you can schedule

your backup tasks to run on non-production hours to reduce the load during peak hours.



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

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

Accessing the Policies panel

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

Procedure

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



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

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


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

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

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 32](#).
- Assign the backup window to an existing policy. To do so, select the policy, click  **Edit**, and then make the required modifications.

Example

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



Backup Window > New ? X

Name
non-production-hours

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

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

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

MON [18:00-06:00] [00:00-24:00]

TUE [18:00-06:00] [00:00-24:00]

WED [18:00-06:00] [00:00-24:00]

THU [18:00-06:00] [00:00-24:00]

FRI [18:00-06:00] [00:00-24:00]

SAT [00:00-24:00]

SUN [00:00-24:00]

Time Frames Clear All

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

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

Close Back **Save**

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


Creating data archives

HYCU Protégé 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 Protégé.

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 Protégé creates a bucket in the region of the Kubernetes cluster or the instance and uses it


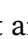
as a target for storing the data. If an automatically created bucket already exists, it is used instead.

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

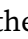
If you select the **Automatically selected** option, a storage class is automatically selected depending on the specified retention.

For details on storage classes, see Google Cloud or AWS 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 32](#).
- Include the data archive into an existing policy. To do so, select the policy, click  **Edit**, and then make the required modifications.

Setting default policies


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

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

Consideration

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

Accessing the Policies panel


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

Procedure

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

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

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

The default policy is represented by the  icon. If you later decide not to use this policy as the default one, click **Clear Default**. Keep in mind that by doing so, you do not unassign this policy from the resources to which it was assigned.

Setting up automatic policy assignment

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

Entities	Method 1	Method 2
SaaS applications ^a	By adding labels to SaaS applications, and then specifying the corresponding label names and values in HYCU Protégé policies. For details, see “Creating custom policies” on page 32 .	By adding the hycu-policy tag to SaaS applications, applications in Google Kubernetes Engine, instances in Amazon EC2 or Google Compute Engine, or buckets in Amazon S3 or Google Cloud Storage. Use the

Entities	Method 1	Method 2
GKE applications	By first adding metadata labels to applications in Google Kubernetes Engine, and then specifying the corresponding label names and values in HYCU Protégé policies. For details, see “Creating custom policies” on page 32 .	following name/value pair: Name: hycu-policy Value: <PolicyName> In this case, <PolicyName> is the name of a HYCU Protégé policy (for example, gold).
Instances	By first adding custom labels to instances in Amazon EC2, or labels (preferred) or custom metadata tags to instances in Google Compute Engine, and then specifying the corresponding label or tag names and values in HYCU Protégé policies. For details, see “Creating custom policies” on page 32 .	
Buckets	By first adding bucket labels to buckets in Google Cloud Storage or Amazon S3, and then specifying the corresponding label names and values in HYCU Protégé policies. For details, see “Creating custom policies” on page 32 .	

^a For details about how to set up automatic policy assignment for specific SaaS applications and for which SaaS applications this feature is supported, see [SaaS application guides](#).

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

Prerequisites

- All relevant prerequisites that apply also for manual policy assignment must be fulfilled. For details, see [“Backing up SaaS applications” on page 53](#), [“Backing up Google Kubernetes Engine applications” on page 69](#), [“Backing up instances” on page 82](#), or [“Backing up buckets” on page 130](#).

- *For Google Kubernetes Engine applications:* The resource objects for which you want to set up automatic policy assignment must be deployed as applications (the resource object of `kind: Application` must be defined in the application deployment).

Considerations

- Assigning policies automatically takes precedence over assigning policies manually or setting a default policy. This means that the label or the tag added to the preferred SaaS application, GKE application, instance, or bucket defines which policy is assigned to it, even if the same entity already has an assigned policy.
- If you want to assign a new policy to a SaaS application, a GKE application, an instance, or a bucket for which automatic policy assignment has been set up, do one of the following:
 - Define new tags or labels as described in this section.
 - Assign the policy to the entity as described in [“Backing up SaaS applications” on page 53](#), [“Backing up Google Kubernetes Engine applications” on page 69](#), [“Backing up instances” on page 82](#), or [“Backing up buckets” on page 130](#). In this case, the manually assigned policy will not be overridden by the automatically assigned one again.

Enabling access to data

Depending on your cloud platform, one of the following applies:

- *For AWS:* You must manually enable access to instances by assigning credential groups to them in HYCU Protégé.
- *For Google Cloud:* HYCU Protégé 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
			5985	HTTP

^a The email address of the authority that is running the task in HYCU Protégé is

`<UserName>@<DomainName>`.

^b HYCU Protégé 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 Protégé.

Data protection scenarios where you must manually enable access to instances

Guest OS	Data protection scenario
any	<ul style="list-style-type: none"> ◦ 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	<ul style="list-style-type: none"> ◦ 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. <p>For more information on OS Login as the access method, see Google Cloud documentation.</p>
Windows	<ul style="list-style-type: none"> ◦ 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.

Manually enabling access to data

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

Task	Instructions
1. Configure credentials groups.	“Configuring credential groups” below
2. Assign credential groups to instances.	“Assigning credential groups” on page 49

Configuring credential groups

Prerequisites

- A user account with sufficient privileges is configured within each instance. For details on how to do this, see Google Cloud or AWS documentation.
- *For Linux instances:*
 - *For AWS or Google Cloud if you do not plan to use the Automatic authentication option.* 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 Protégé 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**.

Procedure

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



Protocol option	Instructions		
Automatic	<p>Select this option if you want HYCU Protégé to automatically select a protocol for accessing the instance—the SSH protocol (TCP port 22) or the WinRM protocol (TCP port 5985, HTTP transport)—, and then enter the user name and password of a user account that has required permissions to access the instance.</p> <p>Use the following format for the user name:</p> <ul style="list-style-type: none"> • Linux: <code><LocalOrDomainUserName></code> • Windows: <code><LocalUserName></code>, <code><Domain>\<DomainUserName></code>, <code><DomainUserName>@<Domain></code> 		
SSH	<p>Select this option if you want to use the SSH protocol for accessing the instance, and then do the following:</p> <ol style="list-style-type: none"> 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 border="1" data-bbox="568 1581 1327 1868"> <tr> <td data-bbox="568 1581 807 1868">Automatic <i>(Only for Google Cloud)</i></td> <td data-bbox="807 1581 1327 1868">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.</td> </tr> </table>	Automatic <i>(Only for Google Cloud)</i>	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.
Automatic <i>(Only for Google Cloud)</i>	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.		

Protocol option	Instructions
	<div data-bbox="820 344 1302 472" style="border-left: 2px solid purple; padding-left: 10px;"> <p>ⓘ Important Do not select this option if OS Login is enabled on your instance.</p> </div> <hr/> <div data-bbox="571 566 778 640"> <p>Password authentication</p> </div> <div data-bbox="820 501 1289 707"> <p>Enter the user name (in the <code><LocalOrDomainUserName></code> format) and password of a user account that has required permissions to access the instance.</p> </div> <hr/> <div data-bbox="571 1140 778 1214"> <p>Public key authentication</p> </div> <div data-bbox="820 741 1315 1615"> <p>Do the following:</p> <ol style="list-style-type: none"> i. Enter the user name (in the <code><LocalOrDomainUserName></code> format) and password of a user account that has required permissions to access the instance. 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 or AWS documentation. iii. <i>Only if the private key is encrypted.</i> Enter the private key passphrase. <div data-bbox="820 1451 1283 1615" style="border-left: 2px solid purple; padding-left: 10px;"> <p>ⓘ Important This selection is mandatory in cases where the SSH server is configured to use public key authentication.</p> </div> </div>

Protocol option	Instructions
	<p>transport protocol of the WinRM server in the instance.</p> <p>b. In the Port field, enter the WinRM server port number.</p> <p>c. Enter the user name (in the <code><LocalOrDomainUserName></code> format (for Google Cloud) or <code><localuser></code>, <code><domain>\<user></code>, or <code><user>@<domain></code> format for AWS) and the password of a user account that has required permissions to access the instance.</p>

6. Click **Save**.

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

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


Assigning credential groups

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

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

Assigning credential groups by using the HYCU Protégé web user interface

Procedure

1. In the Instances panel, select the instances to which you want to assign a credential group.
2. Click  **Credentials**. The Credential Groups dialog box opens.
3. From the list of credential groups, select the credential group that you want to assign to the selected instances, and then click **Assign**.

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

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

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

Assigning credential groups by using labels or metadata tags

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


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

The credential group is automatically assigned to the instance during the next instance synchronization in HYCU Protégé.

Chapter 4

Protecting SaaS applications

HYCU Protégé enables you to protect your SaaS application data with fast and reliable backup and restore operations.



For a list of SaaS applications for which HYCU provides data protection, see HYCU Protégé Marketplace. To access it, in the toolbar, click  **HYCU Protégé Marketplace**.

Prerequisite

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

Consideration

Keep in mind that the role you have assigned determines what kind of actions you can perform. For details on roles, see [“Managing roles” on page 192](#).

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

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

- [“Configuring SaaS application backup options” on the next page](#)
- [“Backing up SaaS applications” on page 53](#)
- [“Restoring SaaS applications” on page 54](#)


Configuring SaaS application backup options

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


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

Backup option	Description
Exclude resources	Enables you to specify one or more resources to be excluded from the backup.
Options	Enables you to use backup options specific to each SaaS application or SaaS application resource (for example, if you are protecting Google Cloud SQL, you can set the offload option that enables HYCU Protégé to delegate the export operation to a separate, temporary instance).

Accessing the SaaS panel

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

Procedure

1. In the SaaS panel, select the SaaS application or the resource for which you want to configure backup options.
2. Click  **Configuration**. The SaaS Configuration dialog box opens.

3. Depending on what you want to do, perform the required action:

I want to...	Instructions
Exclude resources from the backup.	On the Exclude resources tab, select the resources that you want to exclude from the backup.
Use a backup option specific to my SaaS application or resource.	On the Options tab, specify which of the available backup options you want to use and provide the required information.


4. Click **Save**.

Backing up SaaS applications


With HYCU Protégé, you can back up your SaaS application data securely and efficiently.

Limitation

SaaS application data cannot be stored to the automatically created targets. Therefore, the policy that you plan to assign to SaaS applications and/or resources cannot have the Automatically selected option specified for the Backup Target Type.


 **Note** All of the predefined policies have the Automatically selected option enabled.


Accessing the SaaS panel

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

Procedure

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

 **Note** If you want to narrow down the list of displayed SaaS applications, use the filtering options as described in “[Filtering and sorting data in panels](#)” on page 161.

2. Click  **Assign Policy**. The Assign Policy dialog box opens.
3. From the list of available policies, select the preferred policy.

4. Click **Assign** to assign the policy to the selected SaaS applications and/or resources.

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

If required, you can also perform a manual backup of individual SaaS applications and/or resources at any time. For details, see [“Performing manual backups” on page 169](#).


Restoring SaaS applications

HYCU Protégé enables you to restore an entire SaaS application or its resources to a specific point in time.

Considerations


- Only one restore task can run at the same time for a SaaS application or its resource.
- *Only if a SaaS application resource is deleted from cloud.* If the deleted resource has at least one valid restore point available in HYCU Protégé, it is considered protected and its status is PROTECTED_DELETED.


Accessing the SaaS panel

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

Procedure

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

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

2. In the Detail view, select the preferred restore point.
3. Click  **Restore**. The Restore dialog box opens.
4. Select at what level you want to restore your SaaS application or resource (for example, at an instance, database, attachment, or story level), and then click **Next**.

ⓘ **Important** The list of restore options varies depending on the type of your SaaS application. For details about restore options for your SaaS application, see [SaaS application guides](#).

5. Select the SaaS application data 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. *Only if restore options specific to your SaaS application are available.* Specify which of the available restore options you want to use and provide the required information.
8. Click **Restore**.

Chapter 5

Protecting Google Cloud applications

HYCU Protégé enables you to protect your Google Cloud application data with fast and reliable backup and restore operations. After you prepare your application for data protection and back it up, you can choose to restore either the whole application or only specific application items. For a list of supported applications, see the *HYCU Protégé 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 192](#).
- HYCU Protégé uses an external IP address to access Google Cloud APIs if Private Google Access is disabled on subnets. If your data protection environment requires the use of an internal IP address, make sure Private Google Access is enabled on subnets. For details, see Google Cloud documentation.

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

I plan to protect...	Instructions
SAP HANA applications	“Protecting SAP HANA applications” on the next page

I plan to protect...	Instructions
Google Kubernetes Engine applications	“Protecting Google Kubernetes Engine applications” on page 64

Protecting SAP HANA applications

Protecting SAP HANA application data consist of the following tasks:

Task	Instructions
Preparing SAP HANA applications for data protection, which includes enabling access to application data and configuring backup options.	“Preparing for SAP HANA application protection” below
Backing up SAP HANA applications.	“Backing up SAP HANA applications” on page 60
Restoring SAP HANA application data.	“Restoring SAP HANA applications” on page 62

Preparing for SAP HANA application protection




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

Task	Instructions
1. <i>Mandatory.</i> Make sure HYCU Protégé can access applications that you want to protect.	“Enabling access to application data” on the next page
2. <i>Optional.</i> Configure SAP HANA application backup options.	“Configuring SAP HANA application backup options” on page 59


Enabling access to application data

After you assign credentials to instances as described in “[Enabling access to data](#)” on page 44, 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
	<p>HYCU Protégé 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 SAP HANA applications” on page 60.</p>
	<p>The instance credentials do not have proper permissions and HYCU Protégé cannot access applications. To enable HYCU Protégé 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 44.</p> <p>After the discovery status of your application is , make sure to provide also application-specific credentials if your application requires database-level authentication. In this case, follow the procedure described in this section.</p>

Accessing the Applications panel

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

Applications.

Procedure

1. In the Applications panel, select the applications that you want to back up.
2. Click  **Configuration**. The Application Configuration dialog box opens.

3. On the Credentials tab, make sure the **Use instance credentials** switch is disabled, and then enter credentials for a user account with required permissions and access to the applications.
4. Click **Save**.

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


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

Configuring SAP HANA application backup options

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

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

Accessing the Applications panel

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


Applications.

Prerequisite

Only if specifying the temporary instance location and subnet. VPC Network Peering must be configured. For details on how to configure VPC Network Peering, see

Google Cloud documentation.

Procedure

1. In the Applications panel, select the 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, perform the required action:

I want to...	Instructions
Specify the temporary instance location and subnet.	<p>On the Temporary instance configuration tab, do the following:</p> <ol style="list-style-type: none"> 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.
Configure the backup chain length.	<p>On the Backups tab, in the Backup chain length field, specify when you want a new backup chain to be started.</p>

4. Click **Save**.

Backing up SAP HANA applications

With HYCU Protégé, you can back up your SAP HANA application data securely and efficiently.

Prerequisites


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

- The minimum required SAP HANA privileges of the configured SAP HANA database user must be BACKUP ADMIN and CATALOG READ.
- The configured SAP HANA database user must have access permissions to all databases that are being backed up.
- *For SAP HANA systems with the same SID:* A separate target must be configured for each SAP HANA system.

Considerations

- Application data can be stored only to manually created targets, and not to automatically created targets or as a snapshot.
- Before each backup task, the Backint agent is configured to use the service account that you specified when setting up the target for storing backup data. If you are using the default instance service account, the access scope for storage must be Read Write. For details on Cloud API access scopes, see Google Cloud documentation.
- During each backup task, HYCU Protégé 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 161.

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

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

If required, you can also perform a manual backup of any application at any time. For details, see [“Performing manual backups” on page 169](#).

Restoring SAP HANA applications

HYCU Protégé enables you to restore either a whole application or only individual application items to a specific point in time.

Prerequisites

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

Limitation

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

Considerations


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

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

Recommendation

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


Accessing the Applications panel


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

Applications.

Procedure

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

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

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

Protecting Google Kubernetes Engine applications

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

Task	Instructions
Preparing GKE applications for data protection, which includes applying labels on resource objects, discovering applications, and configuring backup options.	“Preparing for Google Kubernetes Engine application protection” below
Backing up GKE applications.	“Backing up Google Kubernetes Engine applications” on page 69
Restoring GKE application data.	“Restoring Google Kubernetes Engine applications” on page 70

Preparing for Google Kubernetes Engine application protection

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

Prerequisite

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

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

Limitations

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

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

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

Task	Instructions
1. <i>Mandatory.</i> Make sure appropriate labels are applied on all resource objects.	“Applying labels on resource objects” below
2. <i>Mandatory.</i> Make sure your GKE applications are discovered in HYCU Protégé.	“Discovering applications” on the next page
3. <i>Optional.</i> Configure GKE application backup options.	“Configuring GKE 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 Protégé. However, you can also use `app: <AppName>`.
 - Namespace in the metadata of the resource object
- *Persistent volume objects:* By applying labels, you ensure that persistent volumes can be discovered and linked to Google Compute Engine disks, which is required for zone/region identification:

Example

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

 **Note** For persistent volumes that use a Container Storage Interface (CSI) provider, the zone/region is specified in the volume handle (for example, volumeHandle:


```
projects/<ProjectID>/zones/<Zone>/disks/<DiskName>).
```


The following deprecated Kubernetes labels are also supported:

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

For details on labels, see Kubernetes documentation.

Discovering applications

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

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

Configuring GKE application backup options

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

Backup option	Description
Pre/post scripts	Enables you to specify the pre-snapshot and post-snapshot scripts to perform necessary actions before and/or after the snapshot of an application is created.
Temporary instance	Enables you to select the region, the zone, and the subnet where you want HYCU Protégé to create a temporary

Backup option	Description
configuration	instance during the backup. By default, the temporary instance is created in the project of the GKE cluster on which the application is running.

Prerequisites

Only if you plan to use pre-snapshot and post-snapshot scripts.


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

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

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 pre-snapshot and post-snapshot scripts.* On the Pre/post scripts tab, specify the scripts to perform necessary actions before and/or after the snapshot of the application is created:

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

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

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

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

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

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

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


Backing up Google Kubernetes Engine applications

With HYCU Protégé, you can back up your GKE application data securely and efficiently.

Prerequisite

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

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 161](#).

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

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

If required, you can also perform a manual backup of any application at any time. For details, see [“Performing manual backups” on page 169](#).

Restoring Google Kubernetes Engine applications

HYCU Protégé 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 72
Restore specific resource objects to the original or a different location.	Restore Resource Objects	“Restoring resource objects” on page 74

Accessing the Applications panel

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


Applications.


Restoring a whole application

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

Procedure

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

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

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

- 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).
- Optional.* In the Post-restore script field, enter the path to the script that HYCU Protégé 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'
```

- 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

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

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

- In the Detail view, select the preferred restore point, and then click  **Restore**. The Application Restore dialog box opens.
- 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 Protégé should run after the restore.

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

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

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

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

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

8. Select the persistent volume claims that you want to restore.
 - 🔗 Tip** Select the **Disks** check box to restore all persistent volume claims.
9. Use the **Keep original volumes** switch if you want to keep the original persistent volumes. If you disable the switch, the original volumes will be overwritten by the restored ones.

10. Click **Restore**.

Restoring resource objects

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

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

Procedure

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

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

2. In the Detail view, select the preferred restore point, and then click **Restore**. The Application Restore dialog box opens.
3. Select **Restore Resource Objects**, and then click **Next**.
4. From the Restore from drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among which you can select:
 - **Automatic**: This option ensures the fastest and most cost-effective restore.
 - **Backup (Snapshot)**
 - **Backup (Target)**
 - **Copy**
 - **Archive—(daily, weekly, monthly, yearly)**
5. From the Target cluster drop-down menu, select the cluster to which you want to restore resource objects. The original cluster of the application is preselected.
6. *Optional*. In the Post-restore script field, enter the path to the script that HYCU Protégé should run after the restore.

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

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

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

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

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

Chapter 6

Protecting instances

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

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

- [“Planning instance protection” below](#)
- [“Backing up instances” on page 82](#)
- [“Restoring instances” on page 86](#)
- [“Restoring individual files or folders” on page 118](#)

Planning instance protection

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

- [“Preparing your data protection environment” below](#)
- [“Configuring instance backup options” on page 78](#)

Preparing your data protection environment

Prerequisites

- *For AWS:*
 - To protect instances in Virtual Private Clouds (VPC) without public IPs or in subnets without public IPs, you must create the following VPC endpoints:

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

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

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

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

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

Limitations

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

Considerations

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

[page 192](#).

- Backup data (including copies of backup data and data archives) that HYCU Protégé creates is crash-consistent, but it may not always be application-consistent. If pre-snapshot scripts are not provided, the application consistency of backup data is limited to the following:
 - Applications that store their data on a single disk.
 - AWS instances and applications that comply with the prerequisites for creating a Windows Volume Shadow Copy Service (VSS) snapshot. For more information about Windows VSS snapshot prerequisites, see [“Backing up instances” on page 82](#).
 - Google Cloud instances and applications that comply with the restrictions for creating a Windows Volume Shadow Copy Service (VSS) snapshot. For details, see Google Cloud documentation.
- *For Google Cloud:* HYCU Protégé uses an external IP address to access Google Cloud APIs if Private Google Access is disabled on subnets. If your data protection environment requires the use of an internal IP address, make sure Private Google Access is enabled on subnets. For details, see Google Cloud documentation.

Recommendation

Only if you delete an instance from Google Cloud. If an instance that you delete from Google Cloud still has at least one valid restore point available in HYCU Protégé, 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 Protégé 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 170](#).

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.
Allowing the restore of individual files	<p>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.</p> <p>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 cloud, and by doing so, instruct HYCU Protégé to allow it automatically. For details, see “Allowing the restore of files by tagging the instance in cloud” on page 82.</p>
Specifying the temporary instance location and subnet	You can specify the region, the zone, and the subnet where you want HYCU Protégé to create a temporary instance during the backup. By default, the temporary instance is created in the source 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 44](#).
 - 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 are protecting Google Cloud instances and you plan to specify a different subnet for the temporary instance.* If you plan to use pre-snapshot and post-snapshot scripts, or back up instances for which the restore of individual files is allowed, VPC Network Peering must be configured. For

details on how to configure VPC Network Peering, see Google Cloud documentation.


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 Protégé uses for running the scripts.

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


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

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:

I want to...	Instructions
Run the pre-snapshot and post-snapshot scripts.	<p>On the Pre/post scripts tab, do the following:</p> <ul style="list-style-type: none"> • In the Pre-snapshot script field, enter the script that HYCU Protégé runs before it creates a snapshot of the instance. The following are examples of the scripts: <ul style="list-style-type: none"> ◦ Linux: <pre>bash /home/<UserName>/freeze_db.sh</pre> ◦ Windows: <pre>%USERPROFILE%\quiesce_db.bat</pre> • In the Post-snapshot script field, enter the script that HYCU Protégé runs after it creates a snapshot of the instance. The following are examples of the scripts: <ul style="list-style-type: none"> ◦ Linux: <pre>bash /home/<UserName>/thaw_db.sh</pre> ◦ Windows: <pre>%USERPROFILE%\resume_db.bat</pre>
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 for the temporary instance.	<p>On the Temporary instance configuration tab, select the following:</p> <ol style="list-style-type: none"> 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 cloud

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

Procedure

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

Name	Value
<code>hycu-enable-flr</code>	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 Protégé 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 44](#).


Backing up instances

With HYCU Protégé, you can back up your instance data securely and efficiently.

Prerequisites

For AWS instances running Windows for which you want to ensure application consistency of backup data:

- You must create an IAM role for VSS-enabled snapshots and attach it to the instance. For details on how to create an IAM role for VSS-enabled snapshots, see AWS documentation.
- All attached disks must be online.

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

Prerequisites when planning to restore individual files or folders

- The restore of individual files or folders must be enabled for the instance. For instructions on how to enable the restore of individual files or folders, see [“Configuring instance backup options” on page 78](#).
- *For AWS instances:*
 - The security group that the instance belongs to must have an inbound rule for the following ports:
 - WinRM: TCP port 5985 for HTTP and TCP port 5986 for HTTPS
 - Linux: TCP port 22 (or a different port if configured for SSH communication)

For instructions on how to configure and apply the network firewall rule, see AWS documentation.
 - The correct credential group must be assigned to the original instance, and the corresponding credentials must belong to a user account with sufficient privileges. For instructions on how to assign access credentials, see [“Enabling access to data” on page 44](#).
- On the instances that you plan to protect, WinRM must be configured to use a TCP port (by default, 5986) for HTTPS or HTTP communication. The firewall must be configured to enable inbound network traffic through this port.
- *For Google Cloud instances:* If you want to use custom credentials for the restore, the correct credential group must be assigned to the original instance, and the corresponding credentials must belong to a user account with sufficient privileges. For instructions on how to assign access credentials, see [“Enabling access to data” on page 44](#).
- *For Google Cloud instances running Windows:*
 - In the Google Cloud Console, there must be 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 must be configured to use a TCP port (by default, 5986) for HTTPS or HTTP communication. The firewall must be configured to enable inbound network traffic through this port.

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

- *For instances running 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 Windows:* On the instances that you plan to protect, WinRM must be configured to use a TCP port (by default, 5986) for HTTPS or HTTP communication. The firewall must be configured to enable inbound network traffic through this port.
- *For instances running Windows, and for instances running 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 44](#).
- *For AWS instances:* The security group that the instance belongs to must have an inbound rule for the following ports:
 - WinRM: TCP port 5986 (or a different port if configured for SSH communication)
 - Linux: TCP port 22 (or a different port if configured for SSH communication)

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

- *For Google Cloud instances:* In the Google Cloud Console, there must be 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:

- Linux: TCP port 22 (or a different port if configured for SSH communication)
- 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.


Limitations

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


Consideration

When backing up an instance with multiple disks, HYCU Protégé 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**.


Procedure

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

To narrow down the list of displayed instances, use the filtering options as described in “[Filtering and sorting data in panels](#)” on page 161.

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

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

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

You can also perform a manual backup of individual instances at any time. For details, see [“Performing manual backups” on page 169](#).

Restoring instances

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

Prerequisites

Only if you plan to specify post-restore scripts.

- Access to the instance file system must be enabled. For instructions, see [“Enabling access to data” on page 44](#).
- 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 Protégé uses for running the scripts.

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

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

- *The instance has a credential group assigned:* The user account specified in the credential group.
- *For Google Cloud instances running Windows:* The user account that is assigned to the instance in HYCU Protégé 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 89
Move Instance	Enables you to move an instance by restoring it to a different cloud platform.	“Moving an instance” on page 99
Restore Disks	Enables you to restore disks and attach them to the same instance.	“Restoring disks” on page 108
Clone Disks	Enables you to restore disks by creating their clones and attaching them to the same or a different instance, or by creating their clones in the same or a different source and zone and leaving them unattached.	“Cloning disks” on page 109
Move Disks	Enables you to move disks by restoring them and attaching them to an instance on a different cloud platform, or by restoring them to a different cloud platform and leaving them unattached.	“Moving disks” on page 112

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 by creating clones of the instances.	“Restoring multiple instances in a single session” on page 114
Restore Disks	Enables you to restore multiple disks on multiple instances at once.	“Restoring multiple disks in a single session” on page 116
Restore from JSON	Enables you to upload an existing restore specification to HYCU Protégé and use it to restore multiple instances or disks.	“Restoring multiple instances or disks from a JSON file” on page 118

Accessing the Instances panel

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

Restoring an instance


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

Consideration


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


Procedure


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

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

2. In the Detail view that appears at the bottom of the screen, select the preferred restore point.

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

 **Note** All disks of the instance are preselected for the restore. With such selection, the entire instance is restored. The boot disk is restored even if you do not select it.
8. *Optional*. In the Post-restore script field, enter the path to the script or a command that HYCU Protégé should run on the instance after the restore.

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

Cloning an instance

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

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

- [“Cloning an AWS instance” on the next page](#)
- [“Cloning a Google Cloud instance” on page 94](#)


Cloning an AWS instance


Limitations

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


Procedure

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


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


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

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

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


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



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

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

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

If you selected Provisioned IOPS SSD or General Purpose SSD, enter the IOPS number.
 - iv. *Only if you want to add labels to the restored disk.*

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

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

- v. Click **Save**.


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

- VPC ID
- Subnet ID



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

Modifying network settings


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

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



Option	Description
None	The network interface does not use a public IP address. This option is preselected if the network interface of the original instance did not use a public IP address.



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

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

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

- e. Click **Add** or **Save**.

- Click  **Delete** next to the network interface that you want to delete. Keep in mind that you cannot restore the instance without a network interface.
14. *Only if you want to add tags to the restored instance.*
 - a. Click **Advanced**.
 - b. For each custom metadata tag that you want to add, click  **Manage**.
 - c. Enter a key and a value, and then click **Add**.

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

Cloning a Google Cloud instance

Limitation

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


Considerations



Only if you plan to replicate disks.


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



Procedure



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

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

- view.
2. In the Detail view that appears at the bottom of the screen, select the preferred restore point.
 3. Click  **Restore Instance**. The Restore Options dialog box opens.
 4. Select **Clone Options**, and then click **Next**.
 5. Select **Clone Instance**, and then click **Next**.
 6. In the New instance name field, specify a new name for the instance.
 7. From the Restore from drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among which you can select:
 - **Automatic**: This option ensures the fastest and most cost-effective restore.
 - **Backup (Snapshot)**
 - **Backup (Target)**
 - **Copy**
 - **Archive—(daily, weekly, monthly, yearly)**
 8. *Optional*. In the Post-restore script field, enter the path to the script or a command that HYCU Protégé should run on the restored instance after the restore.
-  **Note** You can enter any command that the command-line interface of your instance supports.
9. From the Destination source drop-down menu, select the source to which you want to restore the instance. The original source of the instance is preselected. You can choose from sources that belong to the currently selected protection set and that your user account can access.
 10. From the Destination region and Destination zone drop-down menus, select the Google Cloud region and zone to which you want to restore the instance. The original region and zone of the instance are preselected.
 11. Under Disk name, do the following:
 - a. Select the instance disks that you want to restore.

 **Note** All disks of the instance are preselected for the restore. With such selection, the entire instance is restored. The boot disk is restored even if you do not select it.
 - b. Edit the disks as required. For each selected disk, do the following:


- i. Click  **Edit Disk**.
 - ii. *Only if you do not want HYCU Protégé to automatically generate a name for the restored disk device or disk.* Do the following:
 - I. In the New device name field, enter a name for the restored disk device.
 - II. In the New disk name field, enter a name for the restored disk.
 - iii. If you want to change the disk type, from the Disk type drop-down menu, select one of the available disk types for the restored disk (Standard persistent disk, Balanced persistent disks, or SSD persistent disk). By default, the original disk type is selected.
 - iv. If you want to replicate data between two zones in the region of the instance, make sure the **Replicate this disk within region** check box is selected, and then, from the Target zone drop-down menu, select to which zone you want to replicate data. If the selected disk was regional at backup time, the two zones across which the disk is replicated are shown, otherwise, a list of all zones in the region of the instance is shown.
 - v. *Only if you want to add labels to the restored disk.*
 - I. Click **Advanced**.
 - II. Click  **Manage**. The Custom Metadata dialog box opens.
 - III. Enter a key and a value, and then click **Add** for each label that you want to add.

 **Note** If the selected disk already has one or more labels added, they are listed under Labels. If you want to delete any of the added labels, click  **Delete** next to it.
 - vi. Click **Save**.
12. Under Network interfaces, review the list of networks that the original instance was configured in at the time of backup. The list shows the following for each such network:
- 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 public and private IP address types. By default, the public IP address configuration of the original instance is kept.

Modifying network settings

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

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


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

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

options:


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

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

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

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

a. Click **Advanced**.

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

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

📄 Note If the selected instance already has one or more labels, network tags, and/or custom metadata tags added, they are listed

under Labels, Network tags, or Custom metadata. If you want to delete any of the added labels, network tags, and/or custom metadata tags, click **✕ Delete** next to it.

14. Click **Restore**.

Moving an instance

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

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

- [“Moving an instance to AWS” below](#)
- [“Moving an instance to Google Cloud” on page 103](#)


Moving an instance to AWS


Limitation

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

Procedure


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

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


2. In the Detail view that appears at the bottom of the screen, select the preferred restore point.
3. Click  **Restore Instance**. The Restore Options dialog box opens.
4. Select **Move Options**, and then click **Next**.
5. Select **Move Instance**, and then click **Next**.
6. From the Restore from drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among which you can select:
 - **Automatic**: This option ensures the fastest and most cost-effective restore.

- **Backup (Snapshot)**
- **Backup (Target)**
- **Copy**
- **Archive—(daily, weekly, monthly, yearly)**

7. From the Destination source drop-down menu, select the source to which you want to restore the instance. You can choose from sources that belong to the currently selected protection set and that your user account can access.
8. Click **Next**.
9. In the New instance name field, enter the name for the instance.
10. *Only if the original operating system image was not found.* From the Image drop-down menu, select the operating system image you want to use.
To use a custom image, select **Use custom image** and enter the image AMI ID.
11. *Optional.* In the Post-restore script field, enter the path to the script or a command that HYCU Protégé should run on the restored instance after the restore.


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


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

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


16. Under Disk name, do the following:



- a. Select the disks that you want to restore.

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

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

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

If you selected Provisioned IOPS SSD or General Purpose SSD, enter the IOPS number.
 - iv. *Only if you want to add labels to the restored disk.*
 - I. Click **Advanced**.
 - II. Click  **Manage**. The Custom Metadata dialog box opens.
 - III. Enter a key and a value, and then click **Add** for each label that you want to add.


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



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

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

Modifying network settings

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


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

Option	Description
None	<p>The network interface does not use a public IP address.</p> <p>This option is preselected if the network interface of the original instance did not use a public IP address.</p>
Auto-assign	<p>The network interface uses an automatically allocated public IP address.</p> <p>This option is preselected if the network interface of the original instance used a public IP address.</p> <p> Note Auto-assign will not work if the Auto-assign public IPv4 address on a subnet option is set to No or if more than one network interface is specified.</p>
Elastic IP (Reserved)	<p>The network interface uses an elastic public IP address that was reserved in Amazon EC2 in advance.</p>
Elastic IP (New)	<p>The network interface uses a new elastic public IP address.</p> <p> Note Allocation of the IP address in Amazon EC2 is performed at the very beginning of the</p>


Option	Description
	restore. If the allocation fails, the restore task is terminated without being logged.

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

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

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

18. *Only if you want to add tags to the restored instance.*

- Click **Advanced**.
- For each custom metadata tag that you want to add, click  **Manage**.
- Enter a key and a value, and then click **Add**.

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

19. Click **Restore**.


Moving an instance to Google Cloud


Limitation


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

Procedure

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


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

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


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



11. From the Destination region and Destination zone drop-down menus, select the Google Cloud region and zone to which you want to restore the instance.
12. In the vCPU cores field, enter the number of virtual CPUs for the restored instance multiplied by the number of cores per virtual CPU.



13. In the Memory field, set the amount of memory (in GiB) for the restored instance.
14. From the Instance type drop-down menu, select the instance type for the restored instance.

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

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

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

- b. Edit the disks as required. For each selected disk, do the following:
 - i. Click  **Edit Disk**.
 - ii. *Only if you do not want HYCU Protégé to automatically generate a name for the restored disk device or disk. Do the following:*
 - I. In the New device name field, enter a name for the restored disk device.
 - II. In the New disk name field, enter a name for the restored disk.
 - iii. If you want to change the disk type, from the Disk type drop-down menu, select one of the available disk types for the restored disk (Standard persistent disk, Balanced persistent disks, or SSD persistent disk). By default, the original disk type is selected.
 - iv. *Only if you want to add labels to the restored disk.*
 - I. Click **Advanced**.
 - II. Click  **Manage**. The Custom Metadata dialog box opens.
 - III. Enter a key and a value, and then click **Add** for each label that you want to add.


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

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

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

Modifying network settings

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


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

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


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


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

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

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

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

- Click **Advanced**.
- For each label, network tag, or custom metadata tag that you want to add, click  **Manage**.
- Enter a key and a value, and then click **Add**.

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


18. Click **Restore**.


Restoring disks


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

Procedure

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


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

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

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

7. From the Restore from drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among which you can select:
 - **Automatic**: This option ensures the fastest and most cost-effective restore.
 - **Backup (Snapshot)**
 - **Backup (Target)**
 - **Copy**
 - **Archive—(daily, weekly, monthly, yearly)**

8. *Optional.* In the Post-restore script field, enter the path to the script or a command that HYCU Protégé should run after the restore on the instance to which the restored disks are attached.

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

9. Click **Restore**.

Cloning disks

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

Limitations

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


Considerations


- For details on how the restored disks are named, see [“Objects created by HYCU Protégé” on page 215](#).
- *Only if you are restoring Google Cloud instance disks and you plan to replicate disks.*
 - The boot disk cannot be replicated.
 - Standard persistent disks smaller than 200 GiB cannot be replicated.
 - Regional disks can be replicated only across two zones in the same region. One of these zones must be the same as the zone of the destination instance.
 - If the region or zone of the destination instance changes, all regional


disks are automatically converted to zonal disks. In this case, the procedure of replicating the disks must be performed again.

Procedure


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

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


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


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

11. Edit the disks as required. For each selected disk, do the following:

- a. Click  **Edit Disk**.
- b. *Only if you do not want HYCU Protégé to automatically generate a name for the restored disk device or disk. Do the following:*
 - i. In the New device name field, enter a name for the restored disk device.
 - ii. In the New disk name field, enter a name for the restored disk.
- c. If you want to change the disk type, from the Disk type drop-down menu, select one of the available disk types for the restored disk. By default, the original disk type is selected. The following disk types are available:
 - For Google Cloud: Standard persistent disk, Balanced persistent disks, and SSD persistent disk.
 - For AWS: General Purpose SSD, Previous Generation Volume, and Provisioned IOPS SSD.

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

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

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

- f. Click **Save**.
12. Click **Restore**.

Moving disks


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


Limitation

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


Procedure


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

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


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


10. Select the region and the zone that contain the instance to which you want to attach the restored disks.
11. From the Destination instance drop-down menu, select the instance to which you want to attach the restored disks. If you do not want to attach the disks to an instance, select **None (Leave unattached)**.
12. *Optional.* In the Post-restore script field, enter the path to the script or a command that HYCU Protégé should run after the restore on the instance to which the restored disks are attached.

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

13. Edit the disks as required. For each selected disk, do the following:
 - a. Click  **Edit Disk**.
 - b. *Only if you do not want HYCU Protégé to automatically generate a name for the restored disk device or disk.* Do the following:
 - i. In the New device name field, enter a name for the restored disk device.
 - ii. In the New disk name field, enter a name for the restored disk.
 - c. If you want to change the disk type, from the Disk type drop-down menu, select one of the available disk types for the restored disk. By default, the original disk type is selected. The following disk types are available:
 - For Google Cloud: Standard persistent disk, Balanced persistent disks, and SSD persistent disk.
 - For AWS: General Purpose SSD, Previous Generation Volume, and Provisioned IOPS SSD.

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

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

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

| next to it.

- e. Click **Save**.
14. 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. You can also download the restore specification and use it the next time you want to restore multiple instances to speed up the restore process.


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**.


Procedure

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

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

8. Enable the **Overwrite existing** switch to overwrite the existing instances. By default, this option is disabled and the restore of the instance fails if an instance with the same name exists in the destination zone.
9. *Only if you want to start the restore immediately.* Click **Continue to Summary**, and then do the following:
 - To download the restore summary, click **Download Restore Summary as JSON**.
 - To start the restore, click **Start restore**.
10. *Only if you want to apply custom settings to an instance, or edit or download the restore specification.*
 - a. Click **Advanced Settings**.
 - b. *Only if you want to apply custom settings to an instance.* Do the following:
 - i. Under Instance options, from the Instance drop-down menu, select the instance to which you want to apply the custom settings.
 - ii. If you want to rename the instance, select **Rename Instance**, and then enter the new name for the instance and click **Confirm**.
 - iii. *Only if you are restoring AWS instances.* From the Image drop-down menu, select the operating system image you want to use.
To use a custom image, select **Use custom image** and enter the image AMI ID.
 - iv. From the Destination source drop-down menu, select the source to which you want to restore the instance.
 - v. From the Destination zone drop-down menu, select the zone to which you want to restore the instance.
 - vi. *Only if you want to rename a restored disk.* Under Disk options, do the following:
 - I. From the Disk drop-down menu, select the disk that you want to rename.
 - II. Select **Rename disk**, and then enter a new name for the disk and click **Confirm**.
 - vii. Click **Save Restore Settings** to apply the custom settings to the instance.
 - c. *Only if you want to edit the restore specification.* Do the following:

- i. Click **Edit or download restore JSON**. The restore specification generated by HYCU Protégé for all selected instances is displayed, and you can edit it as required.
- ii. Click Restore from .json to start the restore.

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

- d. *Only if you want to download the restore specification.* Do the following:
 - i. Click **Edit or download restore JSON**. The restore specification generated by HYCU Protégé for all selected instances is displayed.
 - ii. Click Download JSON to download the 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. You can also download the restore specification and use it the next time you want to restore multiple disks to speed up the restore process.


Limitation

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


Accessing the Instances panel

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

Procedure

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


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

 **Note** If you want to edit the restore specification by using the REST API Explorer, you can use the URL that is displayed under Request URL.
 - d. *Only if you want to download the restore specification.* Do the following:
 - i. Click **Edit or download restore JSON**. The restore specification generated by HYCU Protégé for all selected instances is displayed.
 - ii. Click Download JSON to download the restore specification.

Restoring multiple instances or disks from a JSON file

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

Procedure

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


Restoring individual files or folders

You can restore one or more individual files or folders to an instance or to a 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” below](#)
- [“Restoring files or folders to a target” on page 124](#)

Accessing the Instances panel

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

Restoring files or folders to an instance

You can restore one or more individual files or folders to the 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 uses one of the supported file systems. For details, see the *HYCU Protégé Compatibility Matrix*.

- *For AWS instances:*

The security group that the instance belongs to must have an inbound rule for the following ports:

- WinRM: TCP port 5985 for HTTP and TCP port 5986 for HTTPS
- Linux: TCP port 22 (or a different port if configured for SSH communication)

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

- *For Google Cloud instances:* In the Google Cloud Console, there must be 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:


- Linux: TCP port 22 (or a different port if configured for SSH communication)
- 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 Amazon EC2 or 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 Protégé. For instructions on how to manually assign credential groups, see [“Enabling access to data” on page 44](#).
 - *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 Protégé 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 HYCU Protégé” on page 215](#).
- *For Google Cloud instances running Linux:* Depending on whether the instance has a credential group assigned in HYCU Protégé, 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 Detail view appears at the bottom of the screen.


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

2. In the Detail view, select the preferred restore point, and then click  **Restore Files**.
3. Select **Restore to original instance**, and then click **Next**.
4. From the Restore from drop-down menu, select which tier you want to use for the restore, and then click **Next**.

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. 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 Original 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 HYCU Protégé” on page 215.

- **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.


- Use the **Restore ACL** switch if you want to restore the original access control list. If enabled, HYCU Protégé preserves original ACLs. If disabled, HYCU Protégé 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

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

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

- In the Detail view, select the desired restore point, and then click  **Restore Files**.
- Select **Restore to different instance**, and then click **Next**.
- In the Restore to Different Instance dialog box, do the following:

- a. From the Restore from drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among which you can select:
 - **Automatic:** This option ensures the fastest and most cost-effective restore.
 - **Backup (Snapshot)**
 - **Backup (Target)**
 - **Copy**
 - **Archive—(daily, weekly, monthly, yearly)**
 - b. From the Zone drop-down menu, select the zone to which the instance to which you want to restore data belongs.
 - c. From the Instance drop-down menu, select the instance to which you want to restore data.
 - d. Click **Next**.
5. In the Choose Files and Folders dialog box, from the list of available files and folders, select the ones that you want to restore, 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. Use the **Restore ACL** switch if you want to restore the original access control list. If enabled, HYCU Protégé preserves original ACLs. If disabled, HYCU Protégé 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 source of the original instance. For information on how to add manually created targets, see [“Adding a bucket to HYCU Protégé as a target” on page 29](#).

Consideration

For details on how the restored individual files or folders are named, see [“Objects created by HYCU Protégé” on page 215](#).

Procedure


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

 **Note** The Detail view appears only if you click an instance. Selecting the check box before the name of the instance does not open the Detail view.
2. In the Detail view, select the preferred restore point, and then click  **Restore Files**.
3. Select **Restore to target**, and then click **Next**.
4. From the Restore from drop-down menu, select which tier you want to use for the restore, and then click **Next**.

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. 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 drop-down menu, select a target to which you want to restore data.
7. Click **Restore**.

Chapter 7

Protecting buckets

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

Prerequisites

For Google Cloud:

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

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 Protégé as a target, see [“Adding a bucket to HYCU Protégé as a target” on page 29](#).

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 192](#).
- *For Google Cloud:* HYCU Protégé uses an external IP address to access Google Cloud APIs if Private Google Access is disabled on subnets. If your data protection environment requires the use of an internal IP address, make

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

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

- [“Configuring bucket backup options” below](#)
- [“Backing up buckets” on page 130](#)
- [“Restoring buckets” on page 131](#)

Configuring bucket backup options

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

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

Prerequisites

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

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

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

Limitations

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

- Only Python scripts are supported.
- *For AWS:* The pre-backup and post-backup scripts must be located in the same account and in the same region as the bucket.
- *For Google Cloud:* Only the `googleapiclient` Python library can be used for making Google Cloud API calls.

Considerations

Only if you are specifying the location or the subnet for a temporary instance. If not specified otherwise, the temporary instance will be created:


- *For AWS:* In the same region as the bucket (for example, US-EAST-1).
- *For Google Cloud:* In the following region (based on the location type of the bucket):
 - *The region:* In the same region as the bucket (for example, US-CENTRAL1).
 - *The dual-region:*

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


- *The multi-region:*

Multi-region name	Temporary instance region
ASIA	ASIA-EAST1
EU	EUROPE-WEST1
US	US-CENTRAL1

Accessing the Buckets panel

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

Procedure

1. In the Buckets panel, select the 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:
 - *Only if specifying the pre-backup and post-backup scripts.* On the Pre/post scripts tab, specify the scripts to perform necessary actions before and/or after the backup of the bucket is performed:
 - In the Pre-backup script field, enter the path to the script that HYCU Protégé 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 Protégé will run immediately after it performs the backup of the bucket.

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

- *For AWS:* `s3://bucket-name/script.py parameter1 parameter2 ...`
- *For Google Cloud:* `gs://bucket-name/script.py parameter1 parameter2 ...`

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

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


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

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


- *Only if specifying the temporary instance location and subnet.* On the Temporary instance configuration tab, provide the following

information:

- a. From the Region drop-down menu, select the preferred region.

 **Note** It is recommended that you select the same region as the one where the bucket resides. Otherwise, you will be charged for outbound data transfer. For details, see Amazon S3 or Google Cloud pricing.

- b. From the Subnet drop-down menu, select the preferred subnet. By default, the temporary instance is created in the default subnet of the preferred region and zone.

 **Important** A policy cannot be assigned to a bucket on which HYCU Protégé could not detect the subnet.

4. Click **Save**.

Backing up buckets

With HYCU Protégé, you can back up your Amazon S3 and Google Cloud Storage bucket data securely and efficiently.

Prerequisite

For Google Cloud: If you plan to back up buckets for which a Shared VPC subnet is specified in configuration, your user account or the HYCU Managed Service Account (HMSA) must be granted the following permissions in the Shared VPC host project: `compute.firewalls.list`, `compute.networks.list`, `compute.networks.get`, `compute.subnetworks.list`, `compute.subnetworks.use`, and `compute.subnetworks.get`.


Limitation

For AWS: Backing up bucket objects that are stored in the S3 Glacier Flexible Retrieval and S3 Glacier Deep Archive storage classes is not supported.


Consideration

The information on the bucket size becomes available in the Detail view after you assign a policy to the bucket. Keep in mind that this size is always rounded up to the full unit, the minimum being 1 GiB.


Accessing the Buckets panel

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

Procedure

1. In the Buckets panel, select the buckets that you want to back up. You can update the bucket list by clicking  **Refresh**.

 **Tip** To narrow down the list of displayed buckets, you can use the filtering options as described in “[Filtering and sorting data in panels](#)” on [page 161](#).

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

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

If required, you can also perform a manual backup of any bucket at any time. For details, see “[Performing manual backups](#)” on [page 169](#).

Restoring buckets

HYCU Protégé enables you to restore one or more individual files or folders inside an Amazon S3 bucket or a Google Cloud Storage bucket to the original or a different bucket.

Prerequisite

For Google Cloud: If you plan to restore buckets for which a Shared VPC subnet is specified in configuration, your user account or the HYCU Managed Service Account (HMSA) must be granted the following permissions in the Shared VPC host project: `compute.firewalls.list`, `compute.networks.list`, `compute.networks.get`, `compute.subnetworks.list`, `compute.subnetworks.use`, and `compute.subnetworks.get`.

Limitation

Only if you plan to restore the original access control list. The Restore ACL option is not available if you are restoring files to a bucket residing on a different cloud


platform.


Consideration



For details on how the restored individual files or folders are named, see [“Objects created by HYCU Protégé” on page 215](#).

Procedure

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

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

2. In the Detail view, select the preferred restore point, and then click  **Restore Files**. The File Restore Options dialog box opens.



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


3. From the Restore from drop-down menu, select which tier you want to use for the restore. Your restore point can contain one or more tiers among which you can select:

- **Automatic:** This option ensures the fastest and most cost-effective restore.
- **Backup (Target)**
- **Copy**
- **Archive—(daily, weekly, monthly, yearly)**



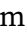
4. Click **Next**. The Choose Files and Folders dialog box opens.



5. From the list of available files and folders, select the ones that you want to restore, and then click **Next**.


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

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

6. Depending on where you want to restore data, select the preferred restore option, click **Next**, and then follow the instructions:

Restore option	Instructions
Restore to original bucket	<p>a. Select the location on the bucket where you want to restore the files or the folders, and then provide the required information:</p> <ul style="list-style-type: none"> • 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. <p>b. Use the Restore ACL switch if you want to restore the original access control list. If enabled, HYCU Protégé preserves original ACLs. If disabled, HYCU Protégé applies inherited ACLs on the restored files (according to the ACL permissions at the bucket or source level).</p> <p>c. <i>Only if you want to add custom metadata tags to the restored bucket objects.</i> Click Advanced, and then, in the Advanced section that opens, do the following:</p> <ol style="list-style-type: none"> i. Click  Manage. The Custom Metadata dialog box opens. ii. Enter a key and a value, and then click Add for each custom metadata tag that you want to add. iii. Click Save. <p> Note If you want to delete any of the added custom metadata tags, click  Delete next to it.</p>
Restore to different bucket	<p>a. From the Source drop-down menu, select the source that contains the bucket to which you want to restore data.</p>

Restore option	Instructions
	<p data-bbox="571 331 1254 409">  Note You can select only among the sources inside the selected protection set. </p> <p data-bbox="520 445 1294 562">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.</p> <p data-bbox="520 584 1267 701">c. Select the location on the bucket where you want to restore the files or folders, and provide the required information:</p> <ul style="list-style-type: none"> <li data-bbox="576 723 847 757">• 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). <li data-bbox="576 1003 863 1037">• 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. <p data-bbox="520 1245 1299 1451">d. Use the Restore ACL switch if you want to restore the original access control list. If enabled, HYCU Protégé preserves original ACLs. If disabled, HYCU Protégé applies inherited ACLs on the restored files (according to the ACL permissions at the bucket or source level).</p> <p data-bbox="520 1473 1302 1590">e. <i>Only if you want to add custom metadata tags to the restored bucket objects.</i> Click Advanced, and then, in the Advanced section that opens, do the following:</p> <ol style="list-style-type: none"> <li data-bbox="571 1612 1299 1691">i. Click  Manage. The Custom Metadata dialog box opens. <li data-bbox="571 1713 1315 1787">ii. Enter a key and a value, and then click Add for each custom metadata tag that you want to add. <li data-bbox="555 1809 756 1843">iii. Click Save.

Restore option	Instructions
	 Note If you want to delete any of the added custom metadata tags, click × Delete next to it.

7. Click **Restore**.

Chapter 8

Performing daily tasks

To ensure your data protection environment is in the optimal state in terms of security, reliability, and efficiency, HYCU Protégé 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 192](#).


I want to...	Instructions
Get an at-a-glance overview of the data protection environment topology and state, identify eventual bottlenecks, and inspect different areas of the data protection environment.	“Using the HYCU Protégé dashboard” on the next page
See the visual representation of my data protection platform.	“Exploring R-Graph” on page 183
View the details of my protected entities.	“Viewing entity details” on page 138
View policy information, edit a policy, or delete a policy.	“Managing policies” on page 143
View target information, activate or deactivate a target, and edit or remove a target.	“Managing targets” on page 145
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 149
View all events that occurred in my data protection environment.	“Viewing events” on page 150

I want to...	Instructions
Configure HYCU Protégé to send notifications when new events occur in my data protection environment.	“Configuring event notifications” on page 152
Obtain HYCU Protégé reports on different aspects of the data protection environment.	“Using HYCU Protégé reports” on page 155
Narrow down the list of displayed items by applying filters and sort the items in panels.	“Filtering and sorting data in panels” on page 161
Back up my data manually.	“Performing manual backups” on page 169
Mark a restore point as expired.	“Expiring backups manually” on page 170
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 172
View subscription information	“Viewing subscription information” on page 173

Using the HYCU Protégé dashboard


The HYCU Protégé dashboard enables you to monitor your data protection environment, observe the relevant data protection activity, and quickly identify the areas that need your attention. You can use this dashboard as a starting point for your everyday tasks because it enables you to easily access the area of interest by simply clicking the corresponding widget.

Accessing the Dashboard panel

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


The following table describes what kind of information you can find within each widget.

Widget	Description
R-Graph	The R-Graph is a visual representation of your data protection environment. For details on the R-Graph, see “Exploring R-Graph” on page 183 .
Policies	Percentage of compliant policies. A policy is considered compliant if all entities to which a policy is assigned are compliant with the policy settings. For details on policies, see “Defining your backup strategy” on page 31 .
Targets	Number of targets in the protection set, and the information about how much space is available for storing the backup data. For details on targets, see “Setting up targets” on page 27 .
Backups	Number of backups performed per day in the protection set, and the backup task success rate for the last seven days in percentages. For details on backups, see “Defining your backup strategy” on page 31 .
Events	Total number of events in the protection set and the number of events according to their severity level (Success, Warning, Failed) in the last 48 hours. For details on events, see “Viewing events” on page 150 .

 **Tip** By clicking icons that denote different statuses within a widget, you are automatically taken to the corresponding panel with the data already filtered accordingly.

Viewing entity details

You can view the details about each discovered entity in the Detail view of the SaaS, Applications, Instances, or Buckets panel.

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


The following details are available:



Entity information	Description
Summary	Shows detailed information about the selected entity.



Restore point



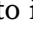
Shows the following information for the restore point:

- Creation date and time.
- Available tiers from which you can restore data:
 - *For SaaS applications, GKE applications, and instances:*
 - **SNAP** or **S**: Snapshot. Displayed if a snapshot of the instance using persistent disks exists. Snapshots allow faster completion of restore tasks.
For SaaS applications, the snapshot tier represents a backup to a staging target or remote storage.
 - **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.




	<p>This depends on the disks included in the corresponding backup.</p> <p>Visual labels of the tiers may be specially marked to denote different statuses. For more information, see “Tier statuses” on page 143.</p> <ul style="list-style-type: none"> ○ <i>For SAP HANA applications:</i> <ul style="list-style-type: none"> ▪ FULL: Full backup. ▪ INCR: Incremental backup. ○ <i>For buckets:</i> <ul style="list-style-type: none"> ▪ 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	<p>Shows the compliance status of the backup (and the resulting restore point):</p> <ul style="list-style-type: none"> • The  icon: The backup is compliant (the RPO setting in the policy assigned to the SaaS application, Google Cloud application, instance, or bucket was met). • The  icon: The backup is not compliant (the RPO setting in the policy assigned to the SaaS application, Google Cloud application, instance, or bucket was not met).




	<ul style="list-style-type: none"> The  icon: The backup compliance status is undefined (the backup is still running). <p>By pausing on a compliance status icon, additional information about the backup is available. You can see backup frequency, the elapsed time since the last successful backup, and the expiration time for each available tier.</p>
Backup status	Shows the backup status of your entity. For more information, see “Viewing the backup status of entities” below.
Restore status	Shows a progress bar indicating the progress of the restore for your entity. <p> Tip If you double-click a progress bar, you are directed to the Tasks panel where you can check details about the related task.</p>

 **Tip** To minimize the Detail view, click  **Minimize** or press the Spacebar. To return the Detail view to its original size, click  **Maximize** or press the Spacebar.

Viewing the backup status of entities

The backup status of your entities determines whether it is possible to restore them.

Backup status	Restore a SaaS app, a GKE app, an instance, or a disk?	Restore files?	Restore a SAP HANA app?	Restore a bucket?
 (Done)	✓	✓ ^a	✓	✓
 (Done with warnings)	✓	✓ ^a	✓	✓
 (Done with errors)	✓ ^b	⚠ ^c	✓ ^d	✓ ^e

Backup status	Restore a SaaS app, a GKE app, an instance, or a disk?	Restore files?	Restore a SAP HANA app?	Restore a bucket?
 (Failed)	x	x	x	x
 (Aborted)	x	x	x	x
 (Expired / Inaccessible on Source / Deleted from Source)	x	x	x	x

^a All disks were backed up successfully, but the disk catalog creation task might have failed. In this case, you will not be able to restore individual files or folders.

^b This backup status may indicate one of the following:


- Not all entities were backed up successfully. Therefore, the entity can be restored only partially. If backing up a boot disk of an instance failed, you may not be able to start the instance after the restore.
- Creating a copy of backup data or a data archive failed. However, the entity can still be fully restored from the backup.
- The backup is not application-consistent.
- *Applicable only if you are using the pre-backup and post-backup scripts.* The script or some actions specified by the script were not executed.

^c This backup status may indicate one of the following:

- Not all disks were backed up successfully and the disk catalog creation task might have failed. In this case, you will not be able to restore individual files or folders.
- Not all disks were backed up successfully, therefore only the files belonging to the disks that were successfully backed up can be restored.















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

^e *Applicable only if you are using the pre-backup and post-backup scripts.* The script or some actions specified by the script were not executed.

 **Note** By pausing on a backup status icon, additional information about the restore point is shown. You can see the backup duration and ID.

Tier statuses


Tier labels may be visually marked to represent backup statuses of individual tiers. These statuses define whether it is possible to restore an entity. The following is an example of possible marks:

Tier status	Restore an entity?
 or  (Done)	✓
 or  (Done with warnings 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 entities” on page 141.
 or  (Failed)	×
 or  (Aborted)	×
 or  (Expired)	×
 or  (Inaccessible on source)	×
 or  (Deleted from the source)	×

Managing policies

You can view policy information, edit policy properties, or delete a policy if you do not want to use it for protecting data anymore.




Accessing the Policies panel

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




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

Viewing policy information

You can view information about each policy in the list of policies in the Policies panel.

Property name	Description
Name	Policy name.
Compliance	Compliance status of the policy: <ul style="list-style-type: none"> • 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 entity, or this is the exclude policy.
SaaS Count	Number of the SaaS applications that have the policy assigned to them.
Instance Count	Number of the instances that have the policy assigned to them.
Application Count	Number of the applications that have the policy assigned to them.
Bucket Count	Number of the buckets that have the policy assigned to them.
Description	Description of the policy.

To open the Detail view where you can find more details about the policy, click the preferred policy.


 **Tip** To minimize the Detail view, click  **Minimize** or press the Spacebar. To return the Detail view to its original size, click  **Maximize** or press the Spacebar.

Creating a policy

See [“Creating custom policies” on page 32](#).

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 32](#).
3. Click **Save**.

Deleting a policy

Limitation

You cannot delete the exclude policy.


Procedure

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

Managing targets

You can view target information, edit a target, deactivate or activate a target, or remove a target if you do not want to use it for storing backup data anymore.


Accessing the Targets panel





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

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




Viewing target information

You can view information about each target in the list of targets in the Targets panel. This allows you to have an overview of the general status of the targets. The following information is available for each target:

Property name	Description
Name	<p>Target name (globally unique).</p> <p>A target that has Object Lock (WORM) enabled is represented by the  icon in the list of targets.</p> <p>For information on how automatically created targets are named, see “Objects created by HYCU Protégé” on page 215.</p>
Target Type	Cloud-specific target type (Amazon S3 or Google Cloud Storage).
Location	Name of the Google Cloud Storage or Amazon S3 geographical region in which the target resides.
Storage Class	<p>Storage classes define the storage availability and pricing model. The default object storage class is displayed for the target.</p> <p>The available options for Google Cloud Storage are: Standard, Nearline, Coldline, and Archive.</p> <p>The available options for Amazon S3 are: S3 Standard, S3 Intelligent-Tiering, S3 Standard-IA, S3 One Zone-IA, S3 Glacier Instant Retrieval, S3 Glacier Flexible Retrieval, or S3 Glacier Deep Archive.</p>
State	<p>Status of the target:</p> <ul style="list-style-type: none"> • Active: You can use the target for backing up data, creating data archives, and restoring data. • Inactive: The target has been deactivated within HYCU Protégé. Until it is activated, you can use it only for restoring data. • Inaccessible on source: HYCU Protégé cannot access the target.


Property name	Description
	<ul style="list-style-type: none"> Deleted from source: The target no longer exists in Google Cloud Storage or Amazon S3. <p>For instructions on how to change the status of active or inactive targets, see “Deactivating and activating targets” on the next page.</p>
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 Protégé. The amount represents a soft limit, therefore actual usage may exceed it.
Health	<p>Health status of the target:</p> <ul style="list-style-type: none"> The  icon: Indicates one of the following: <ul style="list-style-type: none"> 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 or AWS. The  icon: Indicates one of the following: <ul style="list-style-type: none"> 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. Active lifecycle rules are configured for the target.
Utilization	Ratio (expressed in percentage) between the target storage space occupied by backup data and the configured size limit.
Tags	<p>The tag shows if the target:</p> <ul style="list-style-type: none"> Was created automatically by HYCU Protégé (Automatic). Is a staging target for a SaaS application (Staging).

To open the Detail view where you can find more details about the target, click the preferred target.

 **Tip** To minimize the Detail view, click  **Minimize** or press the Spacebar. To return the Detail view to its original size, click  **Maximize** or press the Spacebar.

Editing targets

Procedure

1. In the Targets panel, select the target that you want to edit, and then click  **Edit**. The Edit Target dialog box appears.
2. Edit the selected target as required.
3. Click **Save**.

Deactivating and activating targets

Deactivation of a target makes the target unavailable for backup operations in HYCU Protégé. The target remains registered with HYCU Protégé 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 Protégé.



Prerequisite

For target deactivation: The target must not be 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 if you are deactivating a target.* Click **Yes** to confirm that you want to deactivate the selected target.

Removing targets

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


Prerequisites

- The target must not contain any backup data.
- The target must not be 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 Protégé unless they have been deleted from Google Cloud or AWS.

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.

Checking task statuses


In the Tasks panel, you can do the following:


- Check the overall status of the tasks in your data protection environment.
- Check the status of tasks that are currently running.
- Check the status of completed and stopped tasks.
- Check more details about a specific task.


The information is presented in the Detail view that appears at the bottom of the screen after you select the task.

 **Tip** To minimize the Detail view, click  **Minimize** or press the Spacebar. To return the Detail view to its original size, click  **Maximize** or press the Spacebar.


- Generate a report about a specific task.

To generate the report, select a task, and then click  **View Task Report**.

To copy the report to the clipboard, in the View Task Report dialog box that opens, click  **Copy to Clipboard**.

- Cancel any currently running task by selecting it, and then clicking  **Abort Task**.

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).
Subtask	The list of subordinate tasks.
Started	The task's start date and time.
Finished	The task's finish date and time.




Viewing events

In the Events panel, you can do the following:


- View all events that occurred in your data protection environment.
- Check more details about a specific event in the Detail view that appears at the bottom of the screen after you select the event.

 **Tip** If you click the related task link in the Detail view, you are directed to the Tasks panel where you can view more details about the related task.




- List the events that match the specified filter.
- Configure HYCU Protégé to send notifications when new events occur in your data protection environment. For details, see [“Configuring event notifications” on page 152](#).

 **Tip** To minimize the Detail view, click  **Minimize** or press the Spacebar. To return the Detail view to its original size, click  **Maximize** or press the Spacebar.

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	<p>Severity level of the event:</p> <ul style="list-style-type: none"> •  (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	<p>HYCU Protégé functional area to which the event belongs:</p> <ul style="list-style-type: none"> • Administration: Protection environment changes, such as updated configurations, added/removed sources, or added/removed targets. • Archive: Creation or deletion of archives. • Backup: Events that take place during backup and notifications about skipped backup tasks. • Backup_Window: Events that take place when a backup misses a time period defined for the backup window. • Configuration: Events related to setting backup options for entities. • Credentials: Events related to instance credentials management. • Export: Event list export. • Notification: Possible failures or system malfunctions. • Policies: Creation, updates or removal of policies. • Protege: Events related to HYCU Protégé imports and exports. • Reporting: Events related to report management and generation.



	<ul style="list-style-type: none"> • Restore: Events that take place during restore. • IAM: Added or removed users, updates of user roles, and status changes. • System: Events not related to any other category. Events of this type usually take place independently of your interaction with HYCU Protégé. • Targets: Events related to target management.
Timestamp	Event creation date and time.

Configuring event notifications

You can configure HYCU Protégé to send notifications when new events occur in your data protection environment. This allows you to monitor and manage your data protection environment more efficiently, and to immediately respond to the events if required. You can set up emails or webhooks as a notification channel.

Important Make sure to configure event notifications for each protection set separately.

Accessing the Notifications dialog box


To access the Notifications dialog box, click  **Events** in the navigation pane, and then click  **Notifications** in the toolbar.

Depending on which notification channel you want to use, see one of the following sections:

- [“Creating email notifications” below](#)
- [“Creating webhook notifications” on the next page](#)

Creating email notifications



Procedure

1. In the Notifications dialog box, click the **Email** tab, and then click  **New**.
2. In the Subject field, enter a subject for the email notification.
3. From the Category drop-down menu, select one or more categories. To include all categories, click **Select All**. For a description of categories, see

[“Viewing events” on page 150.](#)


4. From the Status drop-down menu, select one or more statuses. To include all statuses, click **Select All**. For a description of statuses, see [“Viewing events” on page 150.](#)
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 event categories. To include all categories, click **Select All**. For a description of categories, see [“Viewing events” on page 150.](#)
4. From the Status drop-down menu, select one or more statuses. To include all statuses, click **Select All**. For a description of statuses, see [“Viewing events” on page 150.](#)
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 Protégé. 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 Protégé sends to the specified URL, see [“Configuring event notifications” on page 152](#).

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 Protégé
- HTTP request body sent by HYCU Protégé
- HTTP response code sent by the webhook endpoint and received by HYCU Protégé

HTTP request headers

The request headers are sent in the following format:


```
content-type = application/json
x-hycu-signature = base64(hmac(body, secret, 'sha256'))
```

📄 Note The `x-hycu-signature` request header is sent only if the webhook secret is specified.

HTTP request body

The request body is sent in the following format:

```
{
  "severity": "<severity-value>",
  "created": "<created-value>",
  "details": "<details-value>",
  "category": "<category-value>",
  "message": "<message-value>",
  "user": "<user-value>",
  "taskId": "<taskId-value>"
}
```


 **Note** Null values are ignored.

HTTP response code

Your webhook URL should return a response with HTTP status code 204.


Using HYCU Protégé reports

HYCU Protégé 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” on the next page](#), you can continue as follows:

- View reports. For details, see [“Viewing reports” on page 158](#).
- Generate reports. For details, see [“Generating reports” on page 158](#).
- Schedule reports. For details, see [“Scheduling reports” on page 159](#).

 **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 160](#).

Accessing the Reports panel


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

Getting started with reporting

You can take advantage of predefined reports or create additional reports to better understand your data protection environment, identify potential problems, and improve performance.

For a list of predefined reports, see [“Predefined reports”](#) below. 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 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-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-data-on-targets-per-vm	Amount of protected data on targets for each protected instance.
protected-vm-disk-capacity-per-policy	Amount of protected instance disk capacity for each policy.
total-protected-data-on-targets-trend	Total amount of protected data on targets through time.
total-vm-disk-capacity-trend	Total amount of instance disk capacity through time.
transferred-data-per-vm-for-previous-month	Amount of transferred data for each protected instance (per backup tier) for the previous month.
unprotected-vms	List of unprotected instances.
vm-compliance-status	List of instances, their compliance

Name	Description
	statuses, assigned policies, and the corresponding policy tiers.

Tip To minimize the Detail view, click **Minimize** or press the Spacebar. To return the Detail view to its original size, click **Maximize** or press the Spacebar.

Creating reports

If none of the predefined reports meets your reporting requirements, you can create a new report and tailor it to your needs.



Depending on whether you want to create a new report from scratch or edit an existing report and save it as a new report, do the following:


I want to...	Procedure
Create a new report from scratch.	<ol style="list-style-type: none"> 1. Click New. The New Report dialog box opens. 2. Enter a report name and, optionally, its description. 3. Select the type of report (a table or a chart). 4. Select the aggregation value that you want to use to perform a calculation on a set of collected data. 5. Specify the time range for the report. The Time Range drop-down menu allows you to: <ul style="list-style-type: none"> • Select one of the predefined time ranges. • Define a custom from-to time range. Click Custom to define the custom time range using a date and time picker. 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.	<ol style="list-style-type: none"> 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 Preview Report dialog box

I want to...	Procedure
	<p>opens.</p> <ol style="list-style-type: none"> 2. Enter a new name for the report, and then make the required modifications. 3. Click Save as to save the edited report as a new report, or Save to save the changes to the existing report.

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 preferred report, and click  Preview .
View a saved report version.	<ol style="list-style-type: none"> 1. From the list of reports, select the preferred report. 2. In the Detail view that appears at the bottom of the screen, select the preferred report version, and then click  View. <p>For instructions on how to generate report versions manually or automatically, see “Generating reports” below or “Scheduling reports” on the next page.</p>


In the dialog box that opens, besides viewing the report data, you can also download and export the report in the PDF, PNG, or CSV format. To do so, click  **Download**, and then select one of the available formats.


Generating reports



When you generate a report, you save a copy of the current version of the selected report (a report version) for future reference.

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 157.

2. In the Detail view that appears at the bottom of the screen, click  **Generate**. The Generate Report Version dialog box opens.
3. *Optional.* Enter a description for the report version.
4. Click **Generate**.

 **Tip** You can save a version of the selected report also by clicking  **Preview** followed by **Generate**.

The generated report version is added to the list of report versions in the Detail view that appears at the bottom of the screen when you select a corresponding report.


You can later do the following:


- View the saved report versions. For details, see “[Viewing reports](#)” on the [previous page](#).
- Delete the saved report versions that you do not need anymore. To do so, select the preferred report version, and then click  **Delete**.

Scheduling reports

You can use scheduling to generate report versions automatically at a particular time each day, week, or month. You can view these report versions in the web browser or schedule them by email.



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 157.

2. In the Schedule date field, specify the date and the time of day when you want the report generation to begin.

3. From the Interval drop-down menu, select how often you want the report versions to be generated (daily, weekly, or monthly).
4. Use the **Send** switch if you want to schedule the automatic delivery of the reports to email recipients, and then do the following:
 - a. From the Report format drop-down menu, select a file format for your report (PDF, PNG, or CSV).
 - b. In the Email address field, enter one or more email recipients that should receive the reports. If you are entering more than one email address, make sure to press the Spacebar after entering each one.
5. Click **Schedule**.

 **Tip** The reports that are generated automatically are marked by the  icon in the Scheduled column of the Reports panel.

You can later do the following:


- Edit scheduling options of any of the scheduled reports. To do so, select the report, click  **Scheduler**, make the required modification, and then click **Schedule**.
- Unschedule any of the reports if you do not want them to be generated automatically anymore. To do so, select the report, click  **Scheduler**, and then click **Unschedule**.

Exporting and importing reports

HYCU Protégé enables you to share user-created reports among different HYCU Protégé subscriptions by exporting the reports to a JSON file and then importing the reports from the JSON file.

Exporting reports


Procedure


From the list of all reports, select the one that you want to export, and then click  **Export**.

The selected report will be exported to a JSON file and saved to the download location on your system.

Importing reports

Procedure

1. Click  **Import**. The Import Report dialog box opens.
2. Browse your file system for a report that you want to import.
3. Enter a name for the report and, optionally, its description.



 **Note** If the JSON file name and description are already defined in the file itself, the Name and Description fields will be populated automatically. You can, however, use another name and description.

4. Click **Import**.

A new report will be added to the list of the reports.

Filtering and sorting data in panels

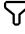
HYCU Protégé enables you to filter data in the panels so you can easily find what you need. Each panel contains different filtering options and it can display only the entries that meet the specified filter criteria. For example, filtering the data in the Instances panel helps you to focus only on the 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 SaaS panel” below](#)
- [“Filtering options in the Applications panel” on the next page](#)
- [“Filtering options in the Instances panel” on page 164](#)
- [“Filtering options in the Buckets panel” on page 165](#)
- [“Filtering options in the Policies panel” on page 166](#)
- [“Filtering options in the Targets panel” on page 166](#)
- [“Filtering options in the Tasks panel” on page 167](#)
- [“Filtering options in the Events panel” on page 168](#)
- [“Filtering options in the IAM panel” on page 169](#)

Filtering options in the SaaS panel

You can enter a SaaS 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
Source	From the drop-down menu, select the relevant HYCU Protégé SaaS source.
Policy	From the drop-down menu, select the policies that are assigned to the SaaS applications.
Compliance	Select one or more options to filter by the compliance status: <ul style="list-style-type: none"> • Success: The SaaS application is compliant. • Failure: The SaaS application is not compliant. • Undefined: The exclude policy is assigned to the SaaS

Filtering option	Action
	application or the SaaS application does not have a policy assigned.
Protection	Select one or more options to filter by the protection status: <ul style="list-style-type: none"> • Yes: The SaaS application is protected. • No: The SaaS application is not protected. • Deleted: The SaaS application is deleted.

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
Source	From the drop-down menu, select the relevant sources to filter the applications.
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: <ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • Yes: The application is protected. • No: The application is not protected. • Deleted: <ul style="list-style-type: none"> • <i>For SAP HANA applications:</i> 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.

Filtering option	Action
	<ul style="list-style-type: none"> • 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	<p>Select one or more options to filter by the application discovery status:</p> <ul style="list-style-type: none"> • 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
Source	From the drop-down menu, select the relevant Google Cloud projects or AWS accounts to filter the instances.
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 instance zones.
Compliance	<p>Select one or more options to filter by the compliance status:</p> <ul style="list-style-type: none"> • 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.

Filtering option	Action
Protection	<p>Select one or more options to filter by the protection status:</p> <ul style="list-style-type: none"> • Yes: The instance is protected. • No: The instance is not protected. • Deleted: The instance no longer exists, but at least one of its valid backups does.
Discovery	<p>Select one or more options to filter by the instance discovery status:</p> <ul style="list-style-type: none"> • 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 source 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:

Source	From the drop-down menu, select the relevant Google Cloud projects or AWS accounts to filter the buckets.
Policy	From the drop-down menu, select the policies that are assigned to the buckets.
Location	From the drop-down menu, select the geographical location of the buckets.
Compliance	<p>Select one or more options to filter by the compliance status:</p> <ul style="list-style-type: none"> • 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	<p>Select one or more options to filter by the protection status:</p> <ul style="list-style-type: none"> • 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	<p>Select one or more options to filter by the compliance status:</p> <ul style="list-style-type: none"> • Success: All entities to which the policy is assigned are compliant. • Failure: Not all entities to which the policy is assigned are compliant. • Undefined: The policy is not assigned to any entity, or an entity is under 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	<p>Select one or more options to filter the targets by the Amazon S3 or Google Cloud storage classes:</p> <ul style="list-style-type: none"> • S3 Standard • S3 Intelligent-Tiering • S3 Standard-IA • S3 One Zone-IA • S3 Glacier Instant Retrieval • S3 Glacier Flexible Retrieval

Filtering option	Action
	<ul style="list-style-type: none"> • S3 Glacier Deep Archive • Standard • Nearline • Coldline • Archive
Health	<p>Select one or more options to filter by the status of the target:</p> <ul style="list-style-type: none"> • 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
Source	From the drop-down menu, select the relevant Google Cloud projects, AWS accounts, or Protégé SaaS modules to filter the tasks.
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	<p>Select one or more options to filter by the status of the task:</p> <ul style="list-style-type: none"> • Ready • Running • Aborting • Aborted • Done • Failed

Filtering option	Action
	<ul style="list-style-type: none"> • 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 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
Source	From the drop-down menu, select the relevant Google Cloud projects, AWS accounts, or Protégé SaaS modules to filter the events.
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: <ul style="list-style-type: none"> • 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.

Filtering options in the IAM 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
Type	Select one or more options to filter by the type: <ul style="list-style-type: none"> • User • Service Account
Status	Select one or more option to filter by status: <ul style="list-style-type: none"> • Active • Deactivated


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.

Performing manual backups

HYCU Protégé backs up your data automatically after you assign a policy to the selected entities. However, you can also back up your data manually at any time, for example, for testing purposes or if an automatic backup fails.


Prerequisite



A policy other than the exclude policy must be assigned to the entity.

Consideration

When the assigned policy uses a backup window, manual backups may prevent the scheduled backup from starting within the defined time frame. This may result in data not being protected until the next backup window or the next manual backup.

Procedure

1. In the SaaS, Applications, Instances, or Buckets panel, select which entities you want to back up.
2. Click  **Backup** to perform the backup of the selected entities.
3. Click **Yes** to confirm that you want to start the manual backup.

 **Tip** In the navigation pane, click  **Tasks** to check the overall progress of the backup.


Expiring backups manually

HYCU Protégé expires backups automatically according to the retention period that is set for the backup data in the policy. However, if there is a restore point that you do not want to use for restoring data anymore, you can at any time expire it manually. You can do this also for restore points whose backup status is Failed or Aborted if you want to free storage space.

A restore point represents data that was backed up at a specified point in time. Your restore point can contain one or more tiers—Backup, Copy, Archive—that can be marked as expired also individually. Keep in mind that the Catalog tier cannot be marked as expired.

Depending on whether the selected restore point belongs to a SaaS application, a Google Cloud application, an instance, or a bucket, it can contain one or more tiers that you can mark as expired:

- For *SaaS applications, instances, and Google Kubernetes Engine applications*: Backup (Target), Backup (Snapshot), Copy, and/or Archive

 **Important** Only the Backup tier is available for GKE applications not using persistent volumes.

- For *SAP HANA applications*: Full or Incremental

ⓘ Important Only Full can be marked as expired if at least one successful full backup has been created after it.

- *For buckets:* Backup, Copy, and/or Archive

You can mark as expired one of the following:

- Entire restore point





Make sure that all tiers are marked for expiration.

- One or more tiers:

Make sure that only the tiers that you want to expire are marked for expiration.

ⓘ Important Marking a restore point or its tiers as expired cannot be undone. If you are marking an application restore point as expired, keep in mind that all previous backups are also marked for expiration.

Depending on whether you want expire backups for a SaaS application, a Google Cloud application, an instance, or a bucket, access one of the following panels:


- Accessing the SaaS panel
To access the SaaS panel, in the navigation pane, click  **SaaS**.
- Accessing the Applications panel
To access the Applications panel, in the navigation pane, click  **Applications**.
- Accessing the Instances panel
To access the Instances panel, in the navigation pane, click  **Instances**.
- Accessing the Buckets panel
To access the Buckets panel, in the navigation pane, click  **Buckets**.


Limitation

You cannot manually expire tiers on targets with Object Lock (WORM) enabled.

Procedure

1. In the relevant panel, click the entity for which you want to expire a backup. The Detail view appears at the bottom of the screen.

 **Note** The Detail view appears only if you click a backup entity. Selecting the check box before its name does not open the Detail view.

2. In the Detail view, select the restore point that you want to mark as expired.
3. Click  **Expire**.
4. *Only if marking an entity restore point as expired and its backup status is not Failed or Aborted.* Select the tiers that you want to mark as expired:
 - Backup (Snapshot): *Available only for SaaS applications, 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 Protégé removes the corresponding data from SaaS remote storage (snapshots), Google Compute Engine (snapshots), Google Cloud Storage (other tiers), Amazon EC2 (snapshots), or Amazon S3 (other tiers).


Exporting the contents of the panel

Data that you can view in a table in any of the panels can be exported to a file in JSON or CSV format.

Consideration

If you want to export only specific data, click  **Filters**, select your filter criteria based on what kind of data you want to export to a file, and then click **Apply Filters**. You can also use the Search field on the left side of the main panel to filter the data.

Procedure


1. Navigate to the panel whose data you want to export.
2. Click  **Export**, and then, from the drop-down menu, select one of the following options:

Option	Description
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.

Viewing subscription information

This section describes the HYCU Protégé subscription information that is provided in the HYCU Protégé web user interface. You can check the information about the current subscription which you can use with your user account (corresponding to each billing account that is linked from any project which you can access).

Accessing the Subscription Information dialog box

To access the Subscription Information dialog box, click  **<email address>** in the toolbar, and then select **Subscription Information**.

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

Subscriber	
First name	Information about the person who subscribed to HYCU Protégé.
Last name	
Company	
Notification email recipients	A list of recipients to whom notifications related to the selected HYCU Protégé subscription will be sent. If this field is empty, all important notifications

	related to the HYCU Protégé 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 or AWS.
Subscription plan	The plan that your HYCU Protégé 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 18.
Subscribed on	The date of subscribing to HYCU Protégé.
Version	Current HYCU Protégé version.
HYCU Account	
HYCU Account ID	Information about the billing account that is billed for the subscription cost.
Login URL	The login URL for the HYCU account.
Alias	An alias for your HYCU account that you can use to sign in to HYCU Protégé.

Chapter 9

Customizing HYCU Protégé

After you subscribe to HYCU Protégé, you can perform various tasks to customize HYCU Protégé for your data protection environment.

Consideration

Keep in mind that the role you have assigned determines what kind of actions you can perform. For details on roles, see [“Managing roles” on page 192](#).

If you have the Administrator role assigned, the scope of tasks you can perform depends on the user interface context you select. You can switch between the following two contexts:


- Subscription

In the subscription context, only the IAM panel and the dashboard are active. Use this context to perform administration tasks related to your subscription, such as adding identity providers, adding or removing users, or changing roles.

- Protection set

In the protection set context, you select the scope of data protection by selecting a specific protection set.

Switching the user interface context

1. On the toolbar, click  next to the name of the selected protection set or subscription.
2. From the drop-down menu, select the context.

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

Tasks

Task	Instructions
Add, edit, or remove sources.	“Managing sources” on the next page

Task	Instructions
Configure service discovery and explore your data protection environment.	“Discovering services” on page 182
Manage identity providers, add or remove users and add or remove roles.	“Managing identity and access” on page 187
Manage HYCU Protégé protection sets.	“Managing protection sets” on page 193
Import service accounts.	“Importing service accounts” on page 198
Hide instances from HYCU Protégé.	“Excluding instances from synchronization by tagging the instance in AWS or Google Cloud” on page 199
Stop protecting individual sources.	“Stopping protection for individual sources” on page 199

Managing sources

HYCU Protégé provides data protection for the following sources:


- AWS accounts
- Google Cloud projects
- Protégé SaaS modules

You can add, edit, or remove the sources directly from HYCU Protégé. When adding sources, HYCU Protégé may request you to grant the required permissions or roles to HYCU Protégé.

For details on how to manage sources, see the following topics:

- [“Managing AWS accounts” on the next page](#)
- [“Managing Google Cloud projects” on page 178](#)
- [“Managing Protégé SaaS modules” on page 180](#)

Accessing the Sources dialog box

To access the Sources dialog box, click  **Administration**, and then select **Sources**.

Accessing the Sources dialog box from an empty R-Graph

In an empty R-Graph with no sources configured, click **Set Up Source**.

Managing AWS accounts

You can perform the following tasks related to AWS accounts:

Task	Instructions
Add an AWS account.	“Adding accounts” below
Edit an existing AWS account.	“Managing sources” on the previous page
Remove an AWS account that you no longer need.	“Managing sources” on the previous page

Adding accounts

If you have the required permissions granted, you can add accounts.

Procedure


1. In the Sources dialog box, select the **AWS** page and click **+** **New**.
2. Enter the account ID, and optionally a display name.
Click **Add**.
3. Click **Create IAM Role**. The AWS Management Console opens.

ⓘ Important You must be logged on to AWS Management Console with the account that you are adding to HYCU Protégé. If you are already logged in to AWS Management Console with a different account when you create the IAM roles, the creation fails.

4. In the AWS Management Console, on the Quick create stack page, confirm the capabilities required by HYCU Protégé by clicking **I acknowledge that AWS CloudFormation might create IAM resources with custom names** and then click **Create stack**.
5. Return to the HYCU Protégé web user interface and click **Save**.
The account is added to the list of sources.

Editing accounts

Procedure

1. In the Sources dialog box, from the list of account IDs, select the one that you want to edit, and then click  **Edit**.
2. Edit the display name and click **Save**.


Removing accounts

You can at any time remove sources that you no longer need.

Consideration

Removing the account from HYCU Protégé does not delete any IAM resources that were created in the AWS account.

Procedure

1. In the Sources dialog box, from the list of account IDs, select the one that you want to remove from HYCU Protégé, and then click  **Remove**.
2. Click **Yes** to confirm that you want to remove the selected account.

Managing Google Cloud projects

You can perform the following tasks related to Google Cloud projects:

Task	Instructions
Add a Google Cloud project and enable the HYCU Managed Service Account for it.	“Adding projects” on the next page
Remove a Google Cloud project that you no longer need.	“Removing projects” on the next page

Enabling the HYCU Managed Service Account

The HYCU Managed Service Account (HMSA) is a special type of account that is designed specifically for HYCU Protégé to run data protection operations. It provides business continuity of your data protection environment by enforcing a single service account that cannot be deleted accidentally, and at the same time it also delivers enhanced security by uniquely identifying the service and

using key rotation to limit risks associated with potential service account key leaks.

You enable the HMSA for a project by following the HYCU Managed Service Account configuration wizard when adding a Google Cloud project as a source.



Prerequisite

You must have the Administrator role assigned.

Adding projects

If you have the required permissions granted, you can add Google Cloud projects.

Procedure

1. In the Sources dialog box, select the **Google Cloud** page and click  **New**.
2. Enter the project ID.
Click **Add**.
3. The HMSA email is displayed. Copy it to the clipboard using the  **Copy to Clipboard** button. You need the email address to assign permissions to HMSA.
Click **Grant Access** to open the HYCU Managed Service Account configuration wizard.
4. The HYCU Managed Service Account configuration wizard guides you through all the required steps of enabling the HMSA for the project.
5. Return to the HYCU Protégé web user interface and click **Save**.
The project is added to the list of sources.


Removing projects

You can at any time remove sources that you no longer need.

Consideration

Removing the project from HYCU Protégé does not delete any IAM resources that were created in the Google Cloud project.

Procedure

1. In the Sources dialog box, on the Google Cloud tab, select the Google Cloud project that you want to remove from HYCU Protégé, and then click 

Delete.

2. Click **Yes** to confirm that you want to remove the selected project.


Managing Protégé SaaS modules

You can perform the following tasks related to Protégé SaaS modules:


Task	Instructions
Add a Protégé SaaS module as a source.	“Adding Protégé SaaS modules” below
Edit an existing Protégé SaaS module.	“Editing Protégé SaaS modules” on the next page
Remove a Protégé SaaS module that you no longer need.	“Removing Protégé SaaS modules” on the next page

Adding Protégé SaaS modules

If you have the required permissions granted, you can add Protégé SaaS modules.

 **Tip** You can access the SaaS tab in the source configuration dialog from HYCU Protégé marketplace directly. Click **Marketplace** in the toolbar, and then click **Configure** in the preferred SaaS module

Procedure

1. In the Sources dialog box, select the **SaaS** page and click  **New**.
2. Select the Protégé SaaS module.
3. Enter the display name.
4. Select the protection set.
5. *Only if your SaaS module supports backups to a staging target.* From the Staging target drop-down menu, select the staging target. Only valid staging targets are listed, see [“Staging targets for SaaS backup” on page 28](#).

If you select the **Automatically selected** option, HYCU Protégé creates a staging target. Automatically created targets can be selected only if the Protégé SaaS module supports this option.


6. Depending on the selected deployment, select the authentication type and enter the required data, such as organizational names, user names, API

tokens, and so on.

7. Click **Save**. The SaaS application is added to the list of sources.

Editing Protégé SaaS modules


Procedure

1. In the Sources dialog box, from the list of HYCU Protégé SaaS modules, select the one that you want to edit, and then click  **Edit**.
2. Edit the display name and, depending to the Protégé SaaS module, the authentication type and other required data, such as organizational names, user names, API tokens, and so on.
3. Click **Save**.


Removing Protégé SaaS modules

You can at any time remove modules that you no longer need.

Prerequisites

- All the policies must be unassigned from all the SaaS applications in the module. To unassign the policies from the SaaS applications, in the SaaS panel, select the applications, and then click  **Assign Policy**. Click **Unassign**, and then click **Yes** to confirm that you want to unassign the policies from the selected SaaS applications.
- No restore points must be present for any of the SaaS applications in the module. If any of the SaaS applications in the module still have valid restore points, you must expire them manually and wait for the next retention maintenance task to finish before removing the module. For details on how to expire restore points, see [“Expiring backups manually” on page 170](#).
- No tasks with the Ready status or a progress bar indicating the Running status must be present for the module.

Procedure

1. In the Sources dialog box, from the list of Protégé SaaS module, select the one that you want to remove from HYCU Protégé, and then click  **Delete**.
2. Click **Yes** to confirm that you want to remove the selected module.

Discovering services

As part of SaaS native data protection, HYCU Protégé offers the means to discover all of your subscribed SaaS services and map them via the R-Graph. This enables you to understand the scope of unprotected SaaS data and leverage HYCU Protégé to have it backed up.

HYCU Protégé discovers sources and SaaS applications automatically as soon as you provide the necessary credentials for accessing the source platform. Automatic discovery of SaaS services is done through supported identity provider services.

Tasks

To successfully discover services, you must complete the following tasks:

Task	Instructions
Add, edit, or delete an identity provider to enable service discovery.	“Configuring service discovery” below
Explore R-Graph.	“Exploring R-Graph” on the next page

Configuring service discovery


You need to configure at least one supported identity provider and provide the necessary credentials to enable HYCU Protégé to automatically discover SaaS services.

Prerequisite


HYCU Protégé requires the following permissions to perform a service discovery:

- **Azure Active Directory:** Use a service principal which has the `Microsoft Graph/Application.Read.All` application resource permissions.
- **Okta:**
 - Use a custom role with permissions to view applications and their details.
 - Use a custom resource set that is constrained to all applications.

Accessing the Discover Services dialog box



To access the Discover Services dialog box, click  **Administration**, and then select **Discover Services**.

Procedure

1. In the Discovery Services dialog box, click  **New**.
2. From the Identity Provider drop-down menu, select the identity provider, and then follow the instructions:

Identity provider	Instructions
Azure Active Directory	<ol style="list-style-type: none"> a. In the Name field, enter a display name for the identity provider. b. In the Client ID field, enter the application ID that is generated by the identity provider. c. In the Client Secret field, enter the application secret that is associated with the client ID and generated by the identity provider. d. In the Tenant ID, enter the tenant ID.
Okta	<ol style="list-style-type: none"> a. In the Name field, enter a display name for the identity provider. b. In the Base Url field, enter the base URL of your provider. c. In the API Key field, enter the API key.

3. Click **Save**. The identity provider is added to the list of configured identity providers.

You can later edit identity provider settings (click  **Edit** and make the required modifications) or delete an identity provider that you do not need anymore (click  **Delete**).

Exploring R-Graph

The R-Graph is a visual representation of your data protection environment, displaying the topology, data protection and compliance statuses of different data sources—cloud workloads, applications and databases, and SaaS applications.

Prerequisite

You need to configure a supported identity provider through which HYCU Protégé performs service discovery. See [“Configuring service discovery”](#)


[on the previous page.](#)

Consideration

If no identity provider is configured, only the protection sets and sources that you already configured in HYCU Protégé are shown. If no sources are configured, an empty R-Graph is shown.

Navigating the R-Graph

Use the following actions to navigate the R-Graph:

- Expand the R-Graph over the entire display pane: Click the  arrow in the bottom right corner of the graph.
- Zoom in and zoom out: Scroll the mouse wheel or double-click an empty area of the graph.
- Move around: When zoomed in, click in the area between nodes and drag the graph to display the part of the graph you are interested in.
- Display additional information about the node: Pause the pointer on the node.

Perform additional actions on nodes:

- Switch to a protection set: Double-click the protection set node.
- Create a protection set and add a source to it: Double-click the generic protection set placeholder to open the Sources > New dialog box.
- Show entities of a particular source: Double-click the source node to open the Detail view of a panel and filter entities by this source.

R-Graph elements and structure

Your data protection environment is represented as a tree graph. Each node represents an element of the data protection environment—protection sets, services, and sources. Connections between nodes represent relationships between these elements—sources are grouped under services and services are grouped under protection sets. Intuitive icons help you to quickly glance the status of data protection for each node and overlay icons show their compliance status.

The following layers of objects represent the data protection environment:

- Root element: HYCU subscription. Visible in Subscription context only.
- ① Protection sets. The default protection set, protection sets that you create, and generic protection set placeholders that group newly discovered modules.
- ② Services. Various types of services or platforms that can have one or more sources. Examples of service types are: “Amazon EC2”, “Salesforce”, “Atlassian Jira”, or “Google Cloud Storage”.
- ③ Sources. AWS accounts, Google Cloud projects, or HYCU Protégé modules for SaaS services.

The following figure shows an example of a zoomed-in part of the R-Graph with three layers of objects, their protection status icons (P), and a tooltip (T) with protection details:

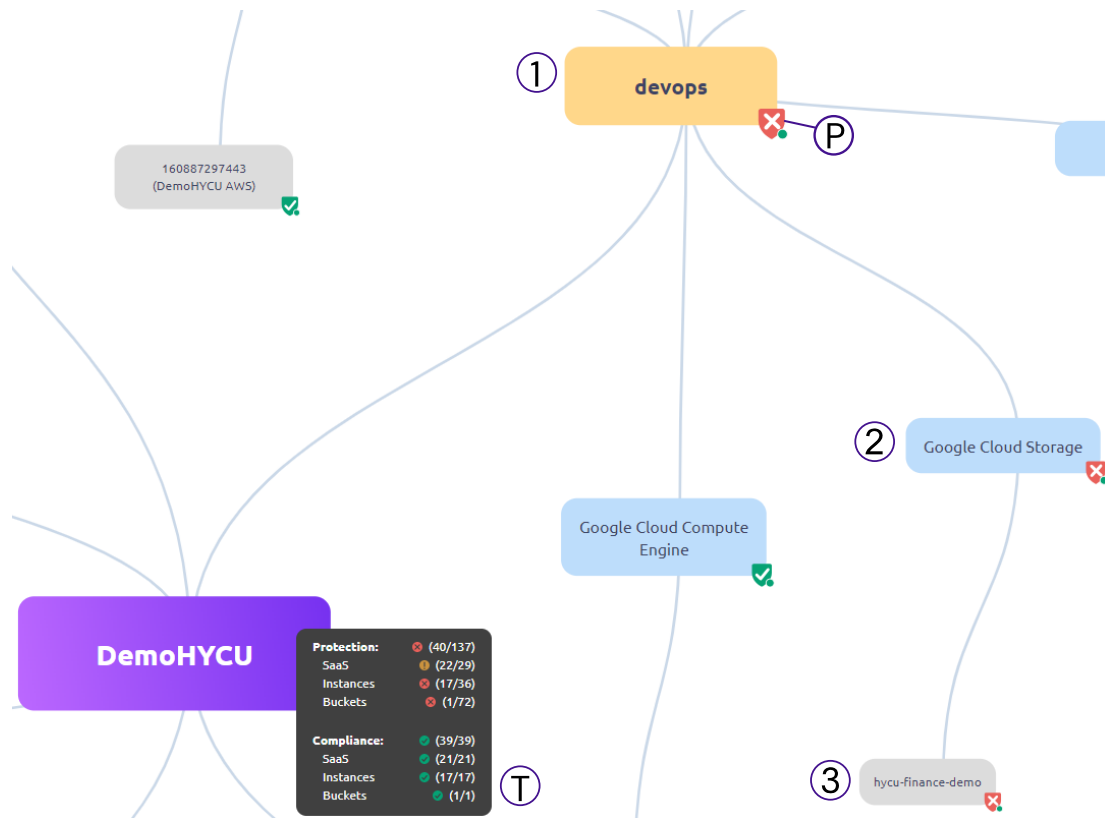


Figure 9-1: R-Graph elements





During the initial scan, HYCU Protégé matches the discovered SaaS services with existing protection sets and maps them accordingly. If no source for a particular Protégé SaaS module is added to any protection set, HYCU Protégé shows the module as part of a generic protection set that would be automatically created when you add the first source, using a name based on the

industry or domain under which the service is classified. If you do not want to create such a protection set, you can add the source to an existing protection set or create a different protection set.

Node protection status





The node protection status is indicated by status icons, compliance overlay icons, and tooltips with more details.

Protection indicators

Status icon	Description
 Green	80% or more included entities have the protection status Green.
 Yellow	Between 60% and 80% of included entities have the protection status Green.
 Red	Less than 60% of included entities have the protection status Green.
 Gray	There are no entities in the node.

Compliance indicators

Node compliance is indicated with overlay icons:

Overlay icon	Description
 Green	80% or more included entities have the compliance status Green.
 Yellow	Between 60% and 80% of included entities have the compliance status Green.
 Red	Less than 60% of included entities have the compliance status Green.
 Gray	There are no entities in the node.

Detailed information about a node

Pause the pointer on the node to display additional information. Each node has the following properties:

Property	Description
Protection	Shows the protection status of the node and individual groups of entities in the node.
Compliance	Shows the compliance status of the node and individual groups of entities in the node.
RPO	<i>Available only for SaaS applications.</i> Shows the RPO as defined by the SaaS application.
Retention	<i>Available only for SaaS applications.</i> Shows the retention period as defined by the SaaS application.

Protection and compliance calculation and inheritance

The protection and compliance status of a node is calculated based on the status of child entities:

- An entity is protected if it has at least one valid restore point available and the entity has a policy assigned.
- An entity is compliant if the RPO set in the assigned policy is met.
- Only entities with an assigned policy are included in compliance calculation.
- Entities with the Exclude policy assigned are excluded from the protection calculation.

The status of a source, service, or protection set node is based on the status of all entities that are included in the node.

Services discovered through identity providers

The protection and compliance of services discovered through identity providers are shown as follows:

- An entity is protected if the RPO of the SaaS application is defined, otherwise it is marked as unprotected.
- A service is not compliant until it is added as a source and is protected by HYCU Protégé.

Managing identity and access

You can use the Identity and access management (IAM) panel to manage identity providers, users, and user roles in HYCU Protégé.

The scope of tasks you can perform depends on your assigned roles and the selected user interface context:

- **Subscription:**

Task	Instructions
Add, edit, or remove identity providers from HYCU Protégé.	“Managing identity providers” below
Add, deactivate, or remove users.	“Managing users” on page 190
Add or remove user roles.	“Managing roles” on page 192

- **Protection set:**

Task	Instructions
Add users.	“Managing users” on page 190
Assign or unassign user roles.	“Managing roles” on page 192

Accessing the IAM panel

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

Managing identity providers

You can integrate HYCU with identity providers that support the OpenID Connect authentication protocol, such as Google, Microsoft, and Okta, to give users the possibility to securely sign in to HYCU Protégé by using these identity providers, without the need to maintain dedicated credentials for HYCU Protégé.

Prerequisites

Only when adding identity providers that support the OpenID Connect authentication protocol. HYCU Protégé must be registered as a web application within the identity provider that you plan to add to HYCU Protégé. When registering HYCU Protégé, make sure the following is done:

- *Only if you are using Microsoft as an identity provider.* In Azure, HYCU Protégé must be given access permissions to the following Azure API: Microsoft Graph with delegated permissions for User .Read.
- *Only if you are using Okta as an identity provider.* In Okta, you must select **Authorization Code** under Client acting on behalf of a user as the grant type.


For instructions on how to register an application, see the respective identity provider documentation.

Accessing the Identity Providers dialog box

To access the Identity Providers dialog box, in the Subscription context, in the IAM panel, click  **Identity Providers**.


Adding an identity provider to HYCU Protégé

Procedure



1. In the Identity Providers dialog box, click  **New**.
2. Enter a name for the identity provider. The name that you specify can contain only lowercase letters and hyphens, must begin and end with a lowercase letter, and cannot be longer than 63 characters.
3. From the Type drop-down menu, select one of the following types of identity providers, and then follow the instructions:

Identity provider type	Instructions
Google	<ol style="list-style-type: none"> a. In the Client ID field, enter the application ID that is generated by the identity provider. b. In the Client secret field, enter the application secret that is associated with the client ID and generated by the identity provider.

Identity provider type	Instructions
Microsoft	a. In the Client ID field, enter the application ID that is generated by the identity provider.
Okta	b. In the Client secret field, enter the application secret that is associated with the client ID and generated by the identity provider.
OIDC	
Cognito	c. In the Issuer field, enter the URL of the issuer of the identity provider.

4. Click  **Copy to Clipboard** to copy the redirect URL that you need to input when you create the application integration with HYCU Protégé.
5. Click **Save**.
6. Configure your identity provider and enter the redirect URL that you copied. For details on the required format, see the respective identity provider documentation.

You can later do the following:

- Edit information about any of the existing identity providers by clicking  **Edit** and making the required modifications.
- Delete any of the existing identity providers by clicking  **Delete**.


Managing users


The HYCU Protégé user management system provides security mechanisms to help prevent unauthorized users from accessing protected data. Only users that are given specific rights have access to the data protection environment. These users can be authenticated either by HYCU or any of the supported identity providers. For details on identity providers, see [“Managing identity providers” on page 188](#).

Consideration


The scope of tasks you can perform depends on the selected UI context. In the Protection set context, you can only add users but cannot deactivate or remove them.

Adding a user

1. In the IAM panel, click  **New User**. The New User dialog box opens.
2. Enter the email address of the user that you want to add.
3. *Optional, if the user will log on using an identity provider.* Select **Generate password** to automatically generate a password. The user must change the generated password during the first log on.

 **Important** If the user has no identity provider configured and you do not generate a password, the user will not be able to log on to HYCU Protégé.

4. *Only if you are adding a user in the Subscription context.* Select one of the following options:
 - **Assign to subscription**
Assign the user to the subscription.
 - **Assign to protection set**
From the list of protection sets, select the one to which you assign the user.


 **Tip** You can search for a protection set by entering its name in the Protection set search field and then pressing **Enter**. By selecting the Name check box, you select all protection sets at once.
5. From the Role drop-down menu, select the role for the user.
You can select more than one role if needed. For more information about user roles, see [“HYCU Protégé roles” on the next page](#).
6. Click **Save**.

Deactivating a user

Consideration

When you deactivate a user, the user can no longer perform any actions. However, the inactive account is preserved in cloud, including all of the data that the user has backed up.

Procedure

1. In the IAM panel, from the list of available users, select the user that you want to deactivate.
2. Click  **Deactivate**. The Deactivate User dialog box opens.
3. Click Deactivate to confirm the deactivation of the user.


Deleting a user


Considerations

- Deleting a user from HYCU Protégé does not remove it from cloud.
- You cannot delete yourself from HYCU Protégé.
- Any upcoming data protection tasks related to the user that you delete will be automatically assigned to you.

Procedure

1. In the IAM panel, from the list of available users, select the one that you want to delete.

 **Tip** You can also search for a user by entering their name in the Search field.

2. Click  **Remove**. The Remove Account dialog box opens.
3. Click **Remove** to confirm that you want the selected user to be deleted from HYCU Protégé.

Managing roles

A role determines the scope of actions that can be performed in the HYCU Protégé data protection environment by a specific user or service account. This means that access to data and information within the data protection environment is limited based on the assigned role. As an administrator, you can manage these roles and define what actions can be performed by each user or service account.

Considerations

- Each user that signs in to HYCU Protégé or each configured service account has by default the Administrator role assigned.
- At least one user with the Administrator role assigned must exist in the data protection environment for each subscription, at the subscription level.
- User roles are inherited from the subscription level to all protection sets under one subscription. User roles set in a protection set are local to that protection set.

HYCU Protégé roles

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

Role	Allowed actions
Administrator	Perform all actions in the data protection environment.
Backup Operator	Define backup strategies, back up SaaS applications, applications, instances, and buckets, and acquire the same information as Viewer.
Restore Operator	Restore SaaS applications, applications, instances, and buckets, and acquire the same information as Viewer.
Viewer	Acquire information about SaaS applications, applications, instances, buckets, policies, targets, tasks, events, reports, service accounts, and protection sets in the data protection environment.


Assigning or unassigning roles

Consideration

If you plan to remove your own Administrator role, keep in mind the following:

- At least one user with the Administrator role assigned must exist in the data protection environment for each subscription.
- You will not be able to change your role back to Administrator yourself.

Procedure

1. In the IAM panel, from the list of available users, select the user for whom you want to change the roles and then click  **Edit**.
2. In the Edit Role dialog box, from the drop-down list, select the roles that you want to assign or unassign. You can select or deselect roles individually or you can click **Select all** to select all roles at once.
3. Click **Save** to save the selected roles.

Managing protection sets


By default, a predefined protection set is created automatically (named default-protection-set). All the Google Cloud projects that are linked to the billing account of your HYCU Protégé subscription are included in the default protection set, and AWS accounts and SaaS modules are added to it when they are added to sources. You can adjust the default setup to better suit your needs

by creating additional protection sets and distributing your sources among them.

You can perform the following tasks related to protection sets:

Task	Instructions
Create a protection set and include preferred sources in it.	“Creating protection sets” below
Edit an existing protection set.	“Editing protection sets” on the next page
Add an AWS account or a Google Cloud project to a protection set by using a label.	“Adding Google Cloud projects to a protection set by using a label” on page 196
Remove an AWS account or a Google Cloud project from a protection set by using a label.	“Removing Google Cloud projects from a protection set by using a label” on page 197
Delete a protection set that you no longer need.	“Deleting protection sets” on page 197

Accessing the Protection Sets dialog box

To access the Protection Sets dialog box, click  **Administration** , 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 sources.


Considerations

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

- Policies will be automatically unassigned from the entities in the source.
- If you move an AWS account or a Google Cloud project, the credential groups that were manually assigned to the instances in the account or the project will be automatically unassigned from those instances.

Procedure

1. In the Protection Sets dialog box, click  **New**. The New Protection Set dialog box opens.
2. Enter a name for your protection set and, optionally, its description.
3. From the list of available sources, select one or more sources that you want to include in the protection set.

 **Tip** You can search for a source by entering its name in the search field and then pressing **Enter**. By selecting the Source check box, you select all sources at once.

4. Click **Save**.

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

Editing protection sets

If you have the required permissions granted, you can change the name of a protection set, add sources to the protection set, or remove sources from the protection set.

When you remove a source from the protection set other than the default one, it is automatically moved to the default protection set. If you want to completely remove the source from HYCU Protégé and stop protecting its resources, you must remove the source from the default protection set.

As an alternative to adding or removing sources by using the HYCU Protégé web user interface, you can also add or remove Google Cloud projects from protection sets by using a label. For details, see the following sections:


- [“Adding Google Cloud projects to a protection set by using a label” on the next page](#)
- [“Removing Google Cloud projects from a protection set by using a label” on page 197](#)

Consideration

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

- Policies will be automatically unassigned from the entities in the source.
- If you move an AWS account or a Google Cloud project, the credential groups that were manually assigned to the instances in the account or the project will be automatically unassigned from those instances.

Procedure

1. In the Protection Sets dialog box, from the list of protection sets, select the one that you want to edit, and then click  **Edit**.
2. Edit the name of the protection set and its description.
3. *Only if you want to add sources to the protection set.* From the list of sources, select one or more sources that you want to add to the protection set. The sources that already belong to the protection set are preselected.
4. *Only if you want to remove sources from the protection set.* From the list of sources, deselect one or more sources that you want to remove from the protection set. The sources that belong to the protection set are preselected.
5. Click **Save**.
6. *Only if you want to add or remove sources from the protection set.* Click **Yes** to confirm that you want to add or remove the sources from the protection set.

Adding Google Cloud projects to a protection set by using a label

As an alternative to adding a project to a protection set by using the HYCU Protégé 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 Protégé.

Procedure

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

Key	Value
<code>hycu-protection-set</code>	<p><code><ProtectionSetName></code></p> <p>In this case, <code><ProtectionSetName></code> is the name of the protection set to which you want to add the project.</p>

For detailed instructions on how to create and manage labels, see Google Cloud documentation.

Removing Google Cloud projects from a protection set by using a label

As an alternative to removing a project from a protection set by using the HYCU Protégé web user interface, you can also remove a project from a protection set by attaching the `hycu-protection-set` label to the project in Google Cloud.

Consideration

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

Procedure

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

Key	Value
<code>hycu-project-exclude</code>	<code>true</code>

After you add the label to the project, it is no longer included in the protection set and HYCU Protégé no longer retrieves its information from Google Cloud.

For detailed instructions on how to create and manage labels, see Google Cloud documentation.

Deleting protection sets

You can at any time delete protection sets that you no longer need.


Prerequisites

- The protection set that you want to delete is empty with no included sources.
- 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 Protégé cannot be deleted.

Procedure

1. In the Protection Sets dialog box, from the list of protection sets, select the one that you want to delete from HYCU Protégé, and then click  **Delete**.
2. 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 Google Cloud target. For details, see [“Adding a bucket to HYCU Protégé as a target” on page 29](#) and [“Managing targets” on page 145](#).


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 Protégé subscription.

Accessing the Service Accounts dialog box


To access the Service Accounts dialog box, click  **Administration**, and then select **Service Accounts**.

Procedure

1. In the Service Accounts dialog box, click  **Import**. The Service Accounts > Import dialog box opens.
2. Click **Browse**.
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 Protégé by selecting it, and then clicking  **Delete**. Keep in mind that deleting the service account from HYCU Protégé does not remove it from Google Cloud.

Stopping protection for individual sources

This section provides instructions that you must follow to stop protecting individual sources in HYCU Protégé.

 **Note** If you want to stop using HYCU Protégé completely, see [“Unsubscribing from HYCU Protégé” on page 208](#).

Procedure

1. In HYCU Protégé, unassign policies from all protected entities in the source. For instructions, see [“Stopping service charges” on page 208](#).
2. In HYCU Protégé, manually mark restore points of all entities in the source as expired. For instructions, see [“Expiring backups manually” on page 170](#).
3. Remove the source from any protection set. For instructions, see [“Editing protection sets” on page 195](#).

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

Excluding instances from synchronization by tagging the instance in AWS or Google Cloud

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

Procedure

1. Depending on your cloud platform, do the following:

Cloud platform	Instructions				
AWS	<p>a. In the AWS Management Console, choose the AWS account to which the instances that you want to leave unprotected belong.</p> <p>b. Within the AWS account, choose an instance and add it the <code>hycu-instance-sync</code> tag in Amazon EC2. Use the following key/value pair:</p> <table border="1"> <thead> <tr> <th>Key</th> <th>Value</th> </tr> </thead> <tbody> <tr> <td><code>hycu-instance-sync</code></td> <td><code>false</code></td> </tr> </tbody> </table> <p>Custom tags can be added from the Amazon EC2 console. For instructions, see AWS documentation.</p>	Key	Value	<code>hycu-instance-sync</code>	<code>false</code>
Key	Value				
<code>hycu-instance-sync</code>	<code>false</code>				
Google Cloud	<p>a. In the Google Cloud Console, choose the Google Cloud project to which the instances that you want to leave unprotected belong.</p> <p>b. Within the project, choose an instance and add it the <code>hycu-instance-sync</code> custom metadata tag in Google Compute Engine. Use the following key/value pair:</p> <table border="1"> <thead> <tr> <th>Key</th> <th>Value</th> </tr> </thead> <tbody> <tr> <td><code>hycu-instance-sync</code></td> <td><code>false</code></td> </tr> </tbody> </table> <p>Custom metadata tags can be added from the Google Cloud Console, the <code>gcloud</code> command line, or by using the Google Cloud API. For instructions, see Google Cloud documentation.</p>	Key	Value	<code>hycu-instance-sync</code>	<code>false</code>
Key	Value				
<code>hycu-instance-sync</code>	<code>false</code>				

2. Repeat step 2 for each additional instance that you want to make invisible to HYCU Protégé.
3. Sign in to the HYCU Protégé web user interface.

4. Select the protection set that includes the same AWS account or Google Cloud project as you selected in step 1 of the procedure. For instructions on selecting protection sets in HYCU Protégé, see [“Selecting a HYCU Protégé protection set” on page 26](#).
5. In the navigation pane, click  **Instances**.
6. 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 10

Troubleshooting

If you encounter a problem while using HYCU Protégé, 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 Protégé?
 - 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 150](#) and [“Checking task statuses” on page 149](#).
 - 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)

Known problems and solutions

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

Missing Google Cloud projects

Problem

When configuring protection sets in HYCU Protégé, not all of your Google Cloud projects are listed in the Protection Sets dialog box. Switching to a billing account of another HYCU Protégé 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 Protégé subscription. For instructions, see Google Cloud documentation.
- Subscribe to HYCU Protégé again and select the billing account to which your missing projects are linked.

Inability to set up manually created Google Cloud targets

Problem

When you try to add a manually created target, HYCU Protégé 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 Protégé” on page 22](#).

Assigning a policy to a Google Cloud instance fails

Problem

After adding the `hycu-policy` custom metadata tag to an instance in Google Compute Engine, no policy is assigned to the instance in HYCU Protégé.

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 Protégé web user interface, go to the Events panel and search for the following error message:

```
Failed to assign a policy
```

2. Click the message entry, check the Message details section for the root cause of the problem, and act accordingly.

Snapshot creation fails for instances in a specific Google Cloud project

Problem

When a backup task for any instance in a specific Google Cloud project is started, the snapshot creation task fails and reports an error.

Solution

In Google Compute Engine, grant your Google Account the Compute Admin (`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 Protégé” on page 22](#).

Task progress indicator remains at 0% during the backup of a Google Cloud instance

Problem

You experience one of the following symptoms:

- When you start a backup task, its child task for creating disk catalog never makes any progress.
- After you start a backup or restore task, the task gets started, but it never makes any progress.

Solution

Check if the Google Cloud project that the instance belongs to has the Cloud Pub/Sub API enabled. If it does not, enable the API for the project through the Google Cloud Console.

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 disk no longer exists, or the credential group that is assigned to the original instance in HYCU Protégé 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 Protégé.

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 Protégé cannot connect to the original instance because no credential group is assigned to the instance in HYCU Protégé 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 [“Manually enabling access to data” on page 45](#).

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 Protégé, only the currently selected protection set is available in the Protection set UI context. After your web user interface session ends, you are unable to sign in again.

Solution

Contact HYCU 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 Windows, automatic assignment of a credential group to the instance fails. HYCU Protégé 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 44](#).

Problem with sorting data in the Events panel

Problem

In the HYCU Protégé web user interface, sorting data in the Events panel by the Message parameter does not function properly.

Solution

There is no solution available for this problem. If data disappears while being sorted, sign out of the web user interface and sign in again to repopulate the table with data.

Chapter 11

Unsubscribing from HYCU Protégé


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










Unsubscribing from HYCU Protégé includes the following tasks:

Task	Instructions
1. Stop being charged for using HYCU Protégé.	“Stopping service charges” below
2. Prevent HYCU Protégé to access your account.	“Preventing account access” on page 210
3. <i>Optional.</i> Remove the HYCU Managed Service Account permissions.	“Removing the HYCU Managed Service Account permissions” on page 212
4. Cancel your HYCU Protégé subscription in the cloud.	“Canceling your HYCU Protégé subscription” on page 213

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 Protégé, unassign policies from all protected entities: <ul style="list-style-type: none">To unassign policies from SaaS applications:<ol style="list-style-type: none">In the navigation pane, click  SaaS.

	<ol style="list-style-type: none"> 2. Select all SaaS applications with assigned policies, and then click  Assign Policy. 3. Click Unassign, and then click Yes to confirm that you want to unassign the policies from the selected SaaS applications. <ul style="list-style-type: none"> • To unassign policies from applications: <ol style="list-style-type: none"> 1. In the navigation pane, click  Applications. 2. Select all applications with assigned policies, and then click  Assign Policy. 3. Click Unassign, and then click Yes to confirm that you want to unassign the policies from the selected applications. • To unassign policies from instances: <ol style="list-style-type: none"> 1. In the navigation pane, click  Instances. 2. Select all instances with assigned policies, and then click  Assign Policy. 3. Click Unassign, and then click Yes to confirm that you want to unassign the policies from the selected instances. • To unassign policies from buckets: <ol style="list-style-type: none"> 1. In the navigation pane, click  Buckets. 2. Select all buckets with assigned policies, and then click  Assign Policy. 3. Click Unassign, and then click Yes to confirm that you want to unassign the policies from the selected buckets. <p> Important If multiple protection sets are available in your data protection environment, make sure to follow these steps for each protection set separately.</p>
2. Stop charges for backup data storage.	<ol style="list-style-type: none"> 1. Manually mark restore points of all entities as expired. For instructions, see “Expiring backups manually” on page 170. <p> Important If multiple protection sets are available in your data protection environment,</p>

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 Protégé from AWS and 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 HYCU Protégé” on page 215](#). For instructions on how to delete targets and remove backup data from targets, see the respective cloud documentation.

4. Remove all snapshots created by HYCU Protégé from AWS and Google Cloud.

For the snapshot naming conventions, see [“Objects created by HYCU Protégé” on page 215](#). For instructions on how to remove snapshots, see the respective cloud documentation.

Preventing account access

As part of unsubscribing from HYCU Protégé, you must prevent HYCU Protégé to access your account.

Cloud platform	Instructions
AWS	“Preventing access to an AWS account” on the next page


Cloud platform	Instructions
Google Cloud	“Preventing access to a Google Cloud account” below

Preventing access to an AWS account

When you added an account as a source to HYCU Protégé, you assigned HYCU Protégé IAM roles to your AWS account. After you stop using the solution, you must remove the roles.

Procedure

1. Open a web browser, go to the [Sign in page](#) of the AWS Management Console and sign in.
2. Open the AWS CloudFormation console and in the navigation pane, choose **Stacks**.
3. In the list of stacks, select `CreateHycuRole` and delete it. When prompted, confirm the deletion.

 **Note** If multiple sources are available in your data protection environment, make sure to follow these steps for each source.

For details on removing AWS stacks, see AWS documentation.

Preventing access to a Google Cloud account

When you subscribed to HYCU Protégé, 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 Protégé**, 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 Protégé 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: 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 Protégé, you will have to grant the HMSA the following roles in Google Cloud on each project that you plan to protect:

- Compute Admin, Service Account User, and Storage Admin
- *Required only if protecting GKE applications.* Kubernetes Engine Admin

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

Canceling your HYCU Protégé subscription

As part of unsubscribing from HYCU Protégé, you must cancel your HYCU Protégé subscription.

Cloud platform	Instructions
AWS	“Cancelling the HYCU Protégé subscription in the AWS Marketplace” below
Google Cloud	“Cancelling the HYCU Protégé subscription in the Google Cloud Marketplace” on the next page

Cancelling the HYCU Protégé subscription in the AWS Marketplace

Prerequisite

Your user account has the `AWSMarketplaceManageSubscriptions` predefined role assigned.

Procedure

1. Open a web browser and go to the [HYCU | AWS Market](#) webpage.
2. Search for "HYCU Protégé for AWS" to find your subscription.
3. On the Manage Subscription page, cancel the subscription. For details on how to cancel an AWS subscription, see AWS documentation.

After you cancel your HYCU Protégé 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 Protégé, subscribe from the same account.

Cancelling the HYCU Protégé subscription in the Google Cloud Marketplace

Prerequisites

- You are signed in to Google with a Google Account that is granted the Billing Account Administrator (`roles/billing.admin`) role on the billing account of the HYCU Protégé subscription.
- Your currently selected project in the Google Cloud Console is linked to the billing account of the HYCU Protégé 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 Protégé 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 Protégé, resubscribe to HYCU Protégé and specify the billing account of the canceled subscription.

Appendix A

Objects created by HYCU Protégé

During data protection tasks, HYCU Protégé creates temporary and persistent HYCU objects in your AWS accounts or 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:
 - *For AWS:*
HYCU-*<Instance>*-snapshot
 - *For Google Cloud:*
hycu-snap-*<TaskUUID>*-*<Disk>*
- Automatically created target:
hycu-*<CloudStorageRegion>*-*<UUID>*
- Target folder with a backup, a backup copy, or a data archive:
hycu/backups/*<Source>*/*<Region/Zone>*/*<Instance>*/disks/*<Disk>*/*<StorageClass>*
- Target folder with a disk catalog:
hycu/backups/*<Source>*/*<Region/Zone>**<Instance>*
/tasks/*<TaskUUID>*/*<Disk>*

Names or location path templates of persistent HYCU 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):
 - *For AWS:*
`<OriginalFileName>
[.<OriginalFileExtension>].<TimeStamp>.restored`
 - *For Google Cloud:*
`<OriginalFileName>.hycu.restored
[.<OriginalFileExtension>]`
- Target folder with restored files or folders:
 - *For AWS:*
`hycu/restores/<Source>/<Instance>/<TaskUUID>/<Path>`
 - *For Google Cloud:*
`hycu/restores/<Source>/<Zone>/<Instance>/
<TaskUUID>/<Disk>/<Volume>/<Path>`
- Restored file:
`<FileName>.<FileExtension>.<TimeStamp>.restored`
- *For Google Cloud:* External IP address resource automatically allocated by HYCU Protégé during cloning:
`hycu-static-external-<UUID>`
- *For Google Cloud:* Internal IP address resource automatically allocated by HYCU Protégé during cloning:
`hycu-static-internal-<UUID>`
- Cloned disk, attached to an instance:
 - *For AWS:*
`<OriginalDiskName>`
 - *For Google Cloud:*
`hycu-disk-<TaskUUID>-<UUID>-<Disk>`
- Cloned or moved disk, unattached:
 - *For AWS:*
`hycu-export-<Disk>`

- *For Google Cloud:*

hycu-disk-<TaskUUID>-<UUID>-<Disk>

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

- Temporary disk:

- *For AWS:*

hycu-temporary-<SnapshotID>

- *For Google Cloud:*

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

Appendix B

Bulk restore specifications

Based on the bulk restore options you specify when restoring multiple instances or disks belonging to multiple instances, HYCU Protégé generates a bulk restore specification.

Elements of a bulk restore specification

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

Syntax

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

```

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

```

Basic elements

- `bulkRestoreType`: "VMS" | "DISKS"
The bulk restore type, VMS for instances and DISKS for disks.
- `overwriteExisting`: false | true
If set to true, instances in the destination region and zone with the same name as the source instances or disks attached to instances, with the same name as the source disks are overwritten during restore.
Default: false
- `items` []
An array of items to restore, each item element contains a source and a destination.

Items to restore

Each item consists of a source and a destination record:

- `source`
The source record contains the path and an array listing the disks.
- `destination`
The destination record contains the path, an array listing the disks, an array listing the network interfaces, and tags and labels.

Source and destination elements

- `path`
The path in the format
projects/
<project>/zones/<targetZone>/instance/<instanceName>.
- `disks`
An array containing the disks to be restored. Disks can either contain the disk name for the source disks or a record with the following elements in the case of destination disks:

- `sourceType`

The source type of the original disk: AWS for AWS, or GC for Google Cloud.
- `diskName`

The name of the original disk.
- `newDiskName`

The name of the restored disk, including the specified postfix. If no postfix is specified, the new name equals the original disk name.
- `newDeviceName`

The name of the restored disk device, including the specified postfix. If no postfix is specified, the new name equals the original device name.
- `diskType`

One of the available disk types for the restored disk.

For AWS, this is I02 or I01 for provisioned IOPS SSD disks, GP3 or GP2 for general purpose SSD disks, SC1 for cold HDD disks, and ST1 for throughput optimized HDD disks.

For Google Cloud, this is BALANCED for balanced persistent disks, EXTREME for extreme persistent disks, HDD for standard persistent disks, or SSD for SSD persistent disks.

By default, the original disk type is used.
- `region`

Specifies the region. You can use it to define a different destination region for the disk.
- `replicaZones`

An array listing the replica zones.

For regional disks, by default only the same replica zone as in the path is added. You need to add the second zone.
- `networkInterfaces`

An array containing network interfaces. Each interface is a record with the following elements:

 - `sourceType`

The source type of the network interface: AWS for AWS, or GC for Google Cloud.
 - `path`

The path to the network device.

- `externalIpType`
The external IP type for the network interface.
For AWS, the `AUTO_ASSIGN` option is always selected.
For Google Cloud, you can select among the following options: `NONE`, `Ephemeral`, `StaticReserved`, `StaticNew`.
- `externalIp`
The external IP value, if supported by the external IP type.
- `internalIpType`
The internal IP address type for the network interface.
For AWS, the `AUTO_ASSIGN` option is always selected.
For Google Cloud, you can select among the following options: `EphemeralAutomatic`, `EphemeralCustom`, `StaticReserved`, `StaticNew`.
- `internalIp`
The internal IP value, if supported by the internal IP type.

The interfaces are selected in the following order:

1. A legacy network with same name.
 2. Shared subnetworks that are accessible in destination projects.
 3. A subnetwork with the name "default".
 4. The first subnetwork in the specified region (sorted by name alphabetically).
- Labels, metadata, and tags
 - `metadata`
Custom metadata tags, consisting of a key and a value pair.
 - `labels`
Labels for the restored disk, consisting of a key and a value pair.
 - `tags`
An array of tags (strings).
 - *For AWS*: Operating system image AMI ID
 - `imageId`
AMI ID of the custom operating system image.

Appendix C

Least-privilege permissions used by HYCU Protégé

To access your data protection environment and perform different tasks such as discovering entities, backing up data, and restoring data, HYCU Protégé does the following:


- *For AWS:* Creates an IAM role for your AWS account with a predefined set of permissions.
- *For Google Cloud:* Uses the permissions that you granted to the Google Account, the Google Service Account, or the HMSA in Google Cloud.

However, if you need to create a custom role with the least-privilege permissions needed to access your data protection environment, you can use the HYCU Protégé role template that contains a predefined set of these permissions. Depending on your cloud platform, see one of the following sections:

Cloud platform	Instructions
AWS	“Using a role template for AWS” below
Google Cloud	“Using a role template for Google Cloud” on page 226

Using a role template for AWS

Prerequisite

You must have the HYCU account ID of your subscription. To get the HYCU account ID, click  **<EmailAddress>** in the toolbar, and then click **Subscription Information** to open the Subscription Information dialog box. The account ID is listed under the HYCU Account section.

Consideration

Make sure that the account for which you are creating the role is not already added as a source in HYCU Protégé, otherwise the creation of the least-privileges role will fail and the role with default permissions will stay in place. If you already added the account as a source, delete its role or the AWS CloudFormation stack with which you created the original role before you start the process or use a different account.

Procedure

To add the role template to your AWS account, perform the following:

1. Open the following URL in your browser:

https://us-east-2.console.aws.amazon.com/cloudformation/home?region=us-east-2#/stacks/quickcreate?templateUrl=https%3A%2F%2Fhycu-resources.s3.amazonaws.com%2Fcloudformation%2F08082022-HycuRoleTemplate-AWSLeastPermissions.json&stackName=HycuStack¶m_ExternalId=<HycuAccountId>

In this URL, *<HycuAccountId>* at the end of the URL is the account ID of your subscription.

ⓘ Important You must be signed in to the AWS Management Console with the account for which you are creating roles. If you are already signed in to the AWS Management Console with a different account when you create the IAM roles, the creation fails.

2. In the AWS Management Console, on the Quick create stack page, confirm the capabilities required by HYCU Protégé by clicking **I acknowledge that AWS CloudFormation might create IAM resources**, and then click **Create stack**.

AWS permissions required by HYCU Protégé

The following is a list of AWS permissions required by HYCU Protégé:

Service	Permissions
S3	ListAllMyBuckets ListBucket GetBucketLocation

	<p>GetBucketLogging GetBucketObjectLockConfiguration GetBucketTagging GetBucketVersioning GetObject GetObjectTagging DeleteJobTagging DeleteObjectTagging DeleteObjectVersionTagging DeleteStorageLensConfigurationTagging PutBucketTagging PutJobTagging PutObjectTagging PutObjectVersionTagging PutStorageLensConfigurationTagging ReplicateTags CreateBucket PutObject</p>
STS	<p>AssumeRole</p>
SQS	<p>GetQueueUrl ListQueues ReceiveMessage CreateQueue DeleteMessage DeleteQueue SendMessage</p>
IAM	<p>GetAccountSummary PassRole</p>
EC2	<p>DescribeAddresses DescribeAvailabilityZones DescribeInstances DescribeInstanceStatus DescribeInstanceTypes DescribeRegions DescribeSecurityGroups</p>

	DescribeSnapshots DescribeSubnets DescribeVolumes GetConsoleOutput CreateTags AllocateAddress AssociateAddress AttachVolume CopyFpgaImage CopyImage CopySnapshot CreateNetworkInterface CreateSnapshot CreateSnapshots CreateVolume DeleteSnapshot DeleteVolume DeregisterImage DetachVolume ImportImage ImportInstance ImportKeyPair ImportSnapshot ImportVolume RegisterImage RunInstances StartInstances StopInstances TerminateInstances
Elastic Block	Store CompleteSnapshot StartSnapshot GetSnapshotBlock ListChangedBlocks ListSnapshotBlocks PutSnapshotBlock
SNS	ListSubscriptions ListSubscriptionsByTopic

	ListTopics GetSubscriptionAttributes GetTopicAttributes ListTagsForResource TagResource UntagResource ConfirmSubscription CreateTopic DeleteTopic Publish SetSubscriptionAttributes SetTopicAttributes Subscribe Unsubscribe
S3 Object Lambda	ListBucket ListBucketMultipartUploads ListBucketVersions ListMultipartUploadParts GetObject GetObjectRetention PutObject PutObjectLegalHold PutObjectRetention RestoreObject WriteGetObjectResponse

Using a role template for Google Cloud

Prerequisite

Your account must have the `iam.roles.create` permission. If you are a project or organization owner, you have this permission by default. If you are not an owner, you must have either the Organization Role Administrator or the IAM Role Administrator role assigned.

Procedure

1. Download the HYCU Protégé 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 a role and grant it the permissions required by HYCU Protégé. To do so, run the following command:

```
gcloud iam roles create <RoleID> --project=<ProjectID> --file=<RoleDefinitionFilePath>
```

In this command, *<RoleID>* is the name of the role (for example *hycuRole*), *<ProjectID>* is the name of your project, and *<RoleDefinitionFilePath>* is the path to the location of the downloaded template that contains the custom role definition.

For details on creating and managing custom roles, see Google Cloud documentation.

Google Cloud permissions required by HYCU Protégé

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

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

<p>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</p>
--

	<p>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</p>
<p>Google Kubernetes Engine</p>	<p>container.clusterRoleBindings.list container.clusterRoles.list container.configMaps.list container.controllerRevisions.list container.cronJobs.list container.customResourceDefinitions.list container.daemonSets.list container.deployments.list container.endpoints.list container.jobs.list container.limitRanges.list container.networkPolicies.list container.podTemplates.list container.replicationControllers.list container.resourceQuotas.list container.roleBindings.list container.roles.list container.secrets.list container.statefulSets.list container.thirdPartyObjects.list resourcemanager.projects.get</p>
<p>Google Cloud Storage</p>	<p>storage.buckets.create storage.buckets.createTagBinding storage.buckets.delete storage.buckets.get</p>

<code>storage.buckets.getIamPolicy</code>
<code>storage.buckets.list</code>
<code>storage.buckets.listTagBindings</code>
<code>storage.buckets.setIamPolicy</code>
<code>storage.buckets.update</code>
<code>storage.objects.create</code>
<code>storage.objects.delete</code>
<code>storage.objects.get</code>
<code>storage.objects.getIamPolicy</code>
<code>storage.objects.list</code>
<code>storage.objects.setIamPolicy</code>
<code>storage.objects.update</code>

Appendix D

Deploying a HYCU backup controller

If you are employing the HYCU Protégé solution in HYCU for Enterprise Clouds, you can use the HYCU Protégé web user interface to deploy a HYCU backup controller to AWS or Google Cloud. This enables you to restore your data in the event of a disaster in your HYCU for Enterprise Clouds data protection environment.

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

Depending on the cloud platform to which you want to deploy the HYCU backup controller, see one of the following sections:

Cloud platform	Instructions
AWS	“Deploying a HYCU backup controller to AWS” below
Google Cloud	“Deploying a HYCU backup controller to Google Cloud” on page 235

Deploying a HYCU backup controller to AWS

To deploy a HYCU backup controller to AWS, follow the procedure described in this section. After you deploy the HYCU backup controller, you must also configure a port in AWS to be able to access the HYCU web user interface. For details, see [“Accessing the HYCU web user interface” on page 235](#).

Prerequisites

- You must own the HYCU and HYCU Protégé licenses. For details on how to obtain these licenses, see HYCU documentation.
- You must have the Administrator role assigned.

Considerations

- The recommended requirements for the HYCU backup controller are 4 vCPU cores and 8 GiB of memory.
- Each HYCU backup controller is by default deployed with the system disk size of 10 GiB and the data disk size of 32 GiB.


Accessing the HYCU Controller Deployment dialog box

To access the HYCU Controller Deployment dialog box, click 

Administration, and then select **HYCU Controller Deployment**.


Procedure

1. In the HYCU Controller Deployment dialog box, select AWS.
2. From the Source drop-down menu, select the account to which you want to deploy the HYCU backup controller.
3. From the Region drop-down menu, select the geographic region for the HYCU backup controller.

 **Important** Make sure that at least one virtual network is configured in the selected region.

4. From the Zone drop-down menu, select the zone for the HYCU backup controller.
5. Click **Next**.
6. In the Instance name field, enter a name for the HYCU backup controller.
7. In the vCPU cores field, enter the number of virtual CPUs to be assigned to the HYCU backup controller multiplied by the number of cores per virtual CPU. The value that you specify must be a whole number and cannot be higher than 128.
8. In the Memory field, enter the amount of memory (in GiB) to be assigned to the HYCU backup controller. The value that you specify must be a whole number and cannot be higher than 24576.
9. From the Available versions drop-down menu, select the preferred version for the HYCU backup controller. By default, the latest version is selected.

10. From the Instance type drop-down menu, select the instance type.


 **Note** The list of available instance types is based on the number of virtual CPU cores and the amount of memory that you specified. If no instance type corresponds to the specified values, the list is empty and you need to adjust the values in the vCPU and Memory fields.

11. Under Network interfaces, you can view the network interface that will be added to the HYCU backup controller. By default, this is the first network interface from the region and zone that you selected for the HYCU backup controller.


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


Modifying network settings

To modify a network interface:


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


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

 **Note** Auto-assign will not work if the Auto-assign public IPv4 address on a subnet option is

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

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

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

- e. Click **Add**.
- Click  **Delete** next to the network interface that you want to delete. Keep in mind that you cannot deploy the HYCU backup controller without a network interface.

12. Click **Deploy**.

Accessing the HYCU web user interface

After you deploy the HYCU backup controller, you must configure a port in AWS to be able to access the HYCU web user interface.

Procedure

Add rules to the security groups to allow inbound traffic. Specify the following settings:

- Source port ranges: 0.0.0.0/0 (to allow any source port)
- Destination port ranges: 8443

For instructions, see AWS documentation.

You can access the HYCU web user interface by entering the following URL:

```
https://<HYCUBackupControllerPublicIPAddress>:8443
```

On the logon page, enter your logon name and password. You can use the default user name and password for initial access:

User name: **admin**

Password: **admin**

ⓘ Important For security purposes, it is highly recommended that you change the default password.

Deploying a HYCU backup controller to Google Cloud

To deploy a HYCU backup controller to Google Cloud, follow the procedure described in this section. After you deploy the HYCU backup controller, you must also configure a port in Google Cloud to be able to access the HYCU web user interface. For details, see [“Accessing the HYCU web user interface” on page 239](#).

Prerequisites

- You must own the HYCU and HYCU Protégé licenses. For details on how to obtain these licenses, see HYCU documentation.

- You must have the Administrator role assigned.
- The Compute Engine default service account must be enabled for the project to which you plan to deploy the HYCU backup controller.

Considerations

- The recommended requirements for the HYCU backup controller are 4 vCPU cores and 8 GiB of memory.
- Each HYCU backup controller is by default deployed with the system disk size of 10 GiB and the data disk size of 32 GiB.


Accessing the HYCU Controller Deployment dialog box

To access the HYCU Controller Deployment dialog box, click 


Administration, and then select **HYCU Controller Deployment**.

Procedure

1. In the HYCU Controller Deployment dialog box, select Google Cloud, and then click **Next**.
2. From the Source drop-down menu, select the project to which you want to deploy the HYCU backup controller.
3. From the Region drop-down menu, select the geographic region for the HYCU backup controller.

 **Important** Make sure that at least one virtual network is configured in the selected region.
4. From the Zone drop-down menu, select the zone for the HYCU backup controller.
5. Click **Next**.
6. In the Instance name field, enter a name for the HYCU backup controller.
7. In the vCPU cores field, enter the number of virtual CPUs to be assigned to the HYCU backup controller multiplied by the number of cores per virtual CPU. The value that you specify must be a whole number and cannot be higher than 1024.
8. In the Memory field, enter the amount of memory (in GiB) to be assigned to the HYCU backup controller. The value that you specify must be a whole number and cannot be higher than 4096.
9. From the Available versions drop-down menu, select the preferred version for the HYCU backup controller. By default, the latest version is selected.

10. From the Instance type drop-down menu, select the instance type.


 **Note** The list of available instance types is based on the number of virtual CPU cores and the amount of memory that you specified. If no instance type corresponds to the specified values, the list is empty and you need to adjust the values in the vCPU and Memory fields.

11. Under Network interfaces, you can view the network interface that will be added to the HYCU backup controller. By default, this is the first network interface from the region and zone that you selected for the HYCU backup controller.


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

Modifying network settings

To modify a network interface:

- Click **Add network interface** to add a network interface or click  **Edit** next to the network interface that you want to edit, and then follow these steps:

- a. *Only if you are adding a network interface.* From the Destination Networks drop-down menu, select the destination network.

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

- b. In the Public address type field, select the public IP address for the network interface. You can select among the following options:


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

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

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

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

- Click  **Delete** next to the network interface that you want to delete. Keep in mind that you cannot deploy the HYCU backup controller without a network interface.

12. Click **Deploy**.

Accessing the HYCU web user interface

After you deploy the HYCU backup controller, you must configure a port in Google Cloud to be able to access the HYCU web user interface.

Procedure

Create an inbound security rule to allow traffic. Specify the following settings:

- Source port ranges: 0.0.0.0/0 (to allow any source port)
- Destination port ranges: 8443

For instructions, see Google Cloud documentation.


You can access the HYCU web user interface by entering the following URL:

```
https://<HYCUBackupControllerPublicIPAddress>:8443
```

On the logon page, enter your logon name and password. You can use the default user name and password for initial access:

User name: **admin**

Password: **admin**

 **Important** For security purposes, it is highly recommended that you change the default password.

Provide feedback

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

info@hycu.com

We will be glad to hear from you!

