

USER GUIDE

HYCU Data Protection as a Service for GCP

Service update date: April 2, 2020
Document edition: Second



Legal notices

Copyright notice

© 2020 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.

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

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

Microsoft and Windows are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or 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 for GCP	9
Key features and benefits	10
Data protection environment overview	11
Data protection with HYCU for GCP	12
Extent of data protection	13
2 Service pricing	15
Backup and recovery pricing	17
Pricing tiers	17
HYCU for GCP subscription plans	17
Pay-as-You-Go subscription plan	18
Calculating your data protection price	18
Price calculation examples	19
Token-based subscription plans	20
Estimating your HYCU token requirements	20
Token consumption calculation examples	21
3 Subscribing to and signing in to HYCU for GCP	23
Subscribing to the service	23
Signing in to HYCU for GCP	25
Service updates	27
4 Establishing a data protection environment	28
Setting the billing account viewer	31
Selecting HYCU for GCP protection sets	31
Selecting a different protection set	32
Defining your backup strategy	33
Choosing between policy types	33
Creating backup window specifications	38
Creating data archives	40

Setting default backup policies	42
Assigning backup policies from the Google Compute Engine service	43
Setting up buckets	45
Setting up user-created buckets	45
Enabling access to objects inside instances	46
Configuring and assigning credential groups manually	47
5 Protecting instances	51
Configuring instance backup options	52
Backing up instances	54
Methods of invoking backups	54
Methods of assigning backup policies	55
Assigning backup policies manually	55
Restoring instances	58
Restoring instances or their disks	58
Cloning instances	59
Restoring individual files or folders	62
6 Performing common tasks	67
Using the HYCU for GCP dashboard	68
Filtering and sorting data in panels	70
Filtering data in panels	70
Filtering options in the Instances panel	70
Filtering options in the Policies panel	71
Filtering options in the Buckets panel	72
Filtering options in the Tasks panel	72
Filtering options in the Events panel	73
Sorting data in panels	73
Checking task statuses	74
Viewing events	74
Viewing detailed instance information	75
Viewing instance summary and restore point information	76

Restore point backup statuses	78
Statuses of restore point entities	79
Using HYCU for GCP reports	80
Getting started with reporting	80
Viewing reports	81
Generating reports	81
Scheduling reports	82
Performing manual backups	82
Manually marking restore points as expired	83
Managing buckets	84
Viewing bucket information	84
Editing buckets	86
Deactivating and activating buckets	86
Removing buckets	87
Managing backup policies	87
Viewing backup policy information	88
Creating backup policies	88
Editing backup policies	88
Deleting backup policies	89
Viewing subscription information	89
7 Administering	91
Managing users	91
Managing protection sets	92
Configuring protection sets	93
Editing protection sets	93
Deleting protection sets	94
Excluding projects from any protection set	95
Configuring protection set service accounts	95
Importing service accounts	97
Assigning imported service accounts to protection sets	98

Configuring event notifications	98
Configuring email-based notifications	99
Configuring webhook-based notifications	100
Excluding instances from synchronization	101
Stopping protection for individual projects	102
8 Troubleshooting	103
General troubleshooting guidelines	103
Problems and solutions	104
Missing Google Cloud Platform projects	104
Inability to set up user-created buckets	104
Backup policy assignment failures	105
Snapshot creation failures	105
Task progress indicator stuck at 0% forever	105
Restore of individual files or folders ending with errors or failing	106
Restore of individual files or folders failing	106
Inability to change the protection set or to sign in	107
Instance backup option reconfiguration failure	107
Getting assistance	107
Customer support	108
Getting additional information and latest updates	108
Before contacting HYCU Customer Support	108
9 Ceasing to use HYCU for GCP	110
Stopping service charges	110
Unassigning backup policies	111
Marking all restore points in a protection set as expired	111
Removing backup data from buckets	112
Removing snapshots	113
Preventing account access	113
Canceling HYCU for GCP subscription	114
A Objects created by the service	115

Glossary	117
HYCU Customer Support and information	121
Customer Support	121
Company resources on the web	121
General information	121
Feedback	121

Chapter 1

About HYCU for GCP

HYCU Data Protection as a Service for GCP (HYCU for GCP) is the first purpose-built backup and recovery solution for the Google Cloud Platform service suite. It is delivered in the form of software-as-a-service (SaaS), a managed service implemented on the basis of the Google Cloud Storage and Google Compute Engine services. It can protect instances running in the Google Compute Engine service. HYCU for GCP is agentless, simple to use, and cost-effective. You can subscribe to HYCU for GCP from the Google Cloud Platform Marketplace.




Figure 1-1: Solution overview

Key features and benefits

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

- **Protection against data loss**

Delivers native data protection for instances (virtual machines) in your projects in the Google Cloud Platform service suite, and ensures easy recoverability.

- **Express setup**

You can enable data protection for your instances in a few minutes after the service is activated for you—no prior actions for solution deployment are required.

- **Predefined and custom policies**

Simplifies implementation of data protection by providing predefined backup 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).

- **Choice of sources and targets**

Selection of backup sources (Google Cloud Platform projects) and backup targets (Google Cloud Storage buckets) is the administrator's choice. Backup data from multiple projects can also be stored in buckets of a dedicated project.

- **Centralized data protection management and monitoring**

You can join Google Cloud Platform projects related to the same service subscription into protection sets to establish centralized data protection management and monitoring for these projects.

- **Optimization and leveraging of the platform value**

Native Google Compute Engine snapshots can be used for both backup and recovery. Google Cloud Storage buckets are created and reused automatically, thus avoiding data protection management overhead.

- **Low impact on applications, backup windows**

Agentless architecture heavily reduces backup load on production instances. In addition, through proper definition of backup windows, you can completely avoid the impact of backup activity on your production environment during peak hours.

- **Application-consistent backup data**

Support for pre-snapshot commands facilitates application data consistency of your backup images.

- **Optimized consumption of storage space**

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

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

A dashboard in the web user interface helps you identify potential problems and bottlenecks to improve the performance of your data protection environment. Predefined reports serve a similar purpose.

- **Restore of individual files or folders**

A possibility to restore one or more files that have become corrupted or have been deleted for some reason and are missing on the instance is an alternative to restoring the whole instance.

- **Use of backup images for cloning**

You are given the possibility of using backup images to clone your instances. Clones can be created in the user-selected project and zone, and can have different network settings.

- **Data archives**

Data archives, isolated from other backup images, can be used for keeping your backup data for a longer time period for future reference. Each data archive can defined multiple archive tiers with different archiving cycles and different retention periods.

- **Integration with Google's billing system, progressive discount**

Cost of data protection is safely billed by Google through existing billing accounts, without requiring you to enter additional billing information. You get a discount depending on the amount of protected data.

Data protection environment overview

The data protection environment of HYCU for GCP is built on two groups of components:

- Google Cloud Platform service suite components. These are native parts of the Google Cloud Platform services. For explanation of the related terms, see the [Google Cloud Platform | Documentation](#) webpage.
- Components provided by HYCU for GCP. They implement data discovery, protection, and analysis.

Table 1-1: Terms related to the HYCU for GCP data protection environment

Term	Description
HYCU for GCP	A managed service implemented on the basis of the Google Cloud Storage and Google Compute Engine services. It provides protection for instances in the Google Compute Engine service.

Web user interface (WUI)	A web-based graphical console that provides access to HYCU for GCP. It enables activities such as: <ul style="list-style-type: none"> • Configuring and maintenance of your data protection environment. • Scheduling and starting backup tasks. • Monitoring task progress. • Browsing the event log.
Sign-in user account	A Google Account that is used to sign in to the HYCU for GCP web user interface.
Entity	An object that you can protect with a backup policy independently of others. In HYCU for GCP, such objects are instances in the Google Compute Engine service.
Source	A resource which includes entities that you can protect. HYCU for GCP treats Google Cloud Platform projects as sources.
Protection set	A group of Google Cloud Platform projects which share the same data protection setup in HYCU for GCP: credential groups, backup policies, backup window specifications, data archives. Such projects may also share an assigned service account.
Target	A storage location used for storing backup data. HYCU for GCP uses Google Cloud Storage buckets as targets.
Backup image	Primary backup data of an instance that can be used to restore the instance. It can be kept in one of the following forms: snapshot, or backup data in a bucket.

For a full glossary of terms that are used in the HYCU for GCP documentation, see [“Glossary” on page 117](#).

Data protection with HYCU for GCP

By using HYCU for GCP as your backup and recovery solution, you can be confident that your business data is backed up in a consistent state, reliably stored, and can be restored and accessed in an uncorrupted state.

After you subscribe to the service through the Google Cloud Platform Marketplace, and sign in to the HYCU for GCP web user interface, you can enable data protection for a Google Cloud Platform project in a few simple steps:

1. *Optional.* Configure a custom protection set, and switch the web user interface to this protection set.

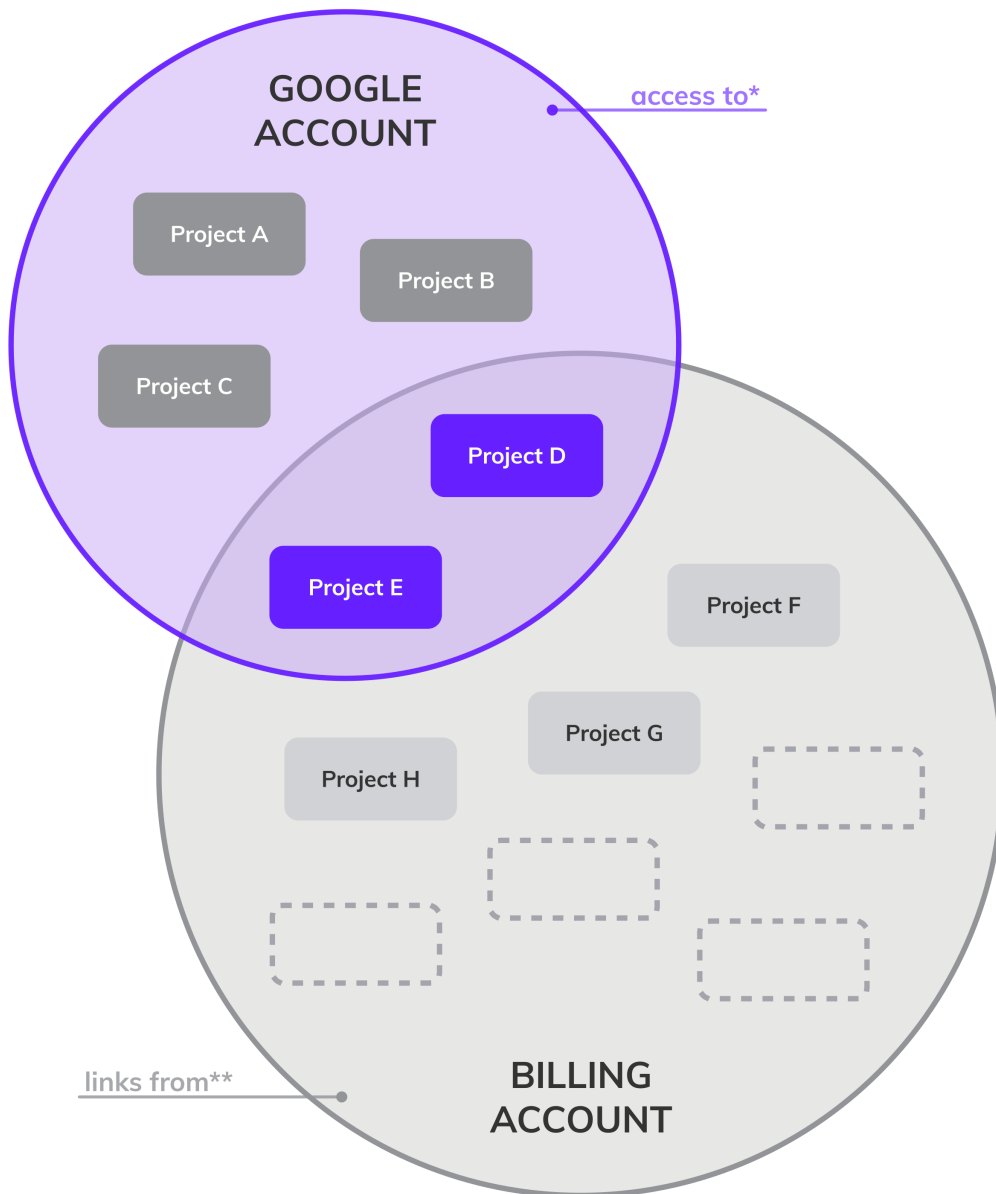
2. Choose appropriate predefined backup policies or create your own ones.
3. Select instances that belong to projects included in the protection set.
4. Assign the backup policies to the instances.

The service then automatically creates instance snapshots or appropriate buckets to store backup images, copied of backup images, or data archives into. When you complete the first backup, you can restore the data from the snapshot or a bucket if the original data becomes damaged or corrupted.

Extent of data protection

The billing account that you select for your HYCU for GCP subscription defines the scope of data protection—the set of Google Cloud Platform projects that are visible to and can be protected by HYCU for GCP within the same subscription. Such projects are referred to as scoped projects. Make sure that all Google Cloud Platform projects that you want to protect are linked to the selected billing account.

This figure depicts an example set-up where your Google Account has access to five projects in the Google Cloud Platform service suite. By using HYCU for GCP, you are able to protect only two of them (projects D and E); only these two projects are linked to the selected billing account and belong to the set of scoped projects for the HYCU for GCP subscription.



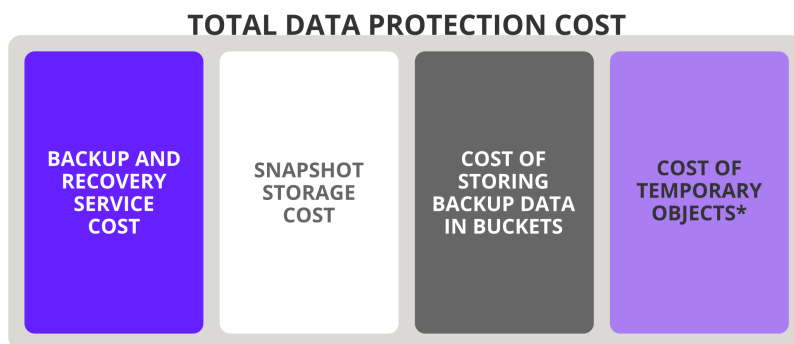
- * Google Account can access projects.
- ** Projects are linked to billing account.

Figure 1-2: Billing account limiting the data protection extent

Chapter 2

Service pricing

HYCU for GCP is a solution for the Google Cloud Platform service suite, and utilizes the Google Cloud Platform resources for its operational needs. When you enable data protection for a Google Cloud Platform project, HYCU for GCP allocates those resources within the same project. Thus, while using HYCU for GCP, you are charged for both the backup and recovery service and the allocated resources it requires.



*This cost exists only with specific configurations and scenarios.

Figure 2-1: Shares of total data protection cost

Total cost for protecting your data in Google Cloud Platform service suite consists of the following shares:

- Backup and recovery service cost

This is a cost generated by using HYCU for GCP to protect your data. It is billed by Google in scope of your Google Account.

This cost depends on the following factors:

- Measured actual usage of the service: frequency of backups, capacity of protected disks in protected instances, and time period for which the data is protected. Data quantity is measured in one gibibyte increments, and protection time is measured in one hour increments. Capacity of protected disks in the protected instances is considered for measurement, not the amount of data that is actually backed up.
- The HYCU for GCP subscription plan: the Pay-as-You-Go plan (also referred to as the Basic plan) or one of the token-based plans.

For more information, see [“Backup and recovery pricing” on the next page](#).

- Snapshot storage cost

This cost is generated by persistent disk snapshots that occupy storage space in the Google Compute Engine resources. If backup images are kept as data in buckets rather than snapshots, this cost is minimal. The cost is charged according to the Google Compute Engine pricing policy. For details, see the [Disks and images pricing | Compute Engine Documentation | Google Cloud](#) webpage.

Within HYCU for GCP, the following factors affect this cost:

- The backup target type
- Frequency of backups
- Snapshot size
- Backup retention period

- Cost of storing backup data in buckets

This cost is generated by backup data that occupies storage space in the Google Cloud Storage buckets. The cost is charged according to the Google Cloud Storage pricing policy. For details, see the [Cloud Storage pricing | Cloud Storage | Google Cloud](#) webpage.

Within HYCU for GCP, the following factors affect this cost:

- The backup target type
- Frequency of backups
- Size of backup data
- Use of backup image copies
- Backup retention period
- Use of data archives, configured archive tiers and their retention periods
- Enabled restore of individual files or folders

- Cost of temporary objects that are used by HYCU for GCP

This cost is generated by temporary objects that are created in your Google Cloud Platform project for the duration of the following processes and tasks:

Process or task	Temporary objects
Instance rediscovery after assigning a credential group	Instance with one disk
Instance rediscovery after selecting the Enable restore of individual files option	Instance with one disk
Backup of instances with the Enable restore of individual files option selected	Instance with multiple disks

Restore of instances or entire instance disks	Instance with one disk
Restore of individual files or folders	Instance with one disk, additional disks in your target instance

For the respective Google Compute Engine pricing policies, see the [All pricing | Compute Engine Documentation](#) | [Google Cloud](#) webpage.

Backup and recovery pricing

The pricing model used for HYCU for GCP is similar to the one used for the Google Cloud Platform service suite. This section explains the pricing model and briefly describes the HYCU for GCP subscription plans.

Pricing tiers

HYCU for GCP pricing model defines four pricing tiers, depending on how frequently your data is backed up. Each pricing tier defines:

- *With the Pay-as-You-Go subscription plan.* The fee that you are charged for protecting a particular instance.
- *With token-based subscription plans.* The number of HYCU tokens that you consume for protecting a particular instance.

In the following table, an RPO period is the recovery point objective setting (value of the Backup every option) in the assigned backup policy. For more information on the recovery point objective, see [“Defining your backup strategy” on page 33](#).

Table 2-1: HYCU for GCP pricing tiers

Pricing tier	RPO period
platinum	1-3 hours
gold	4-11 hours
silver	12-23 hours
bronze	24 hours or more

Protected instances are classified into pricing tiers depending on the assigned backup policy. For actual price tags of different pricing tiers, see the [HYCU | Marketplace - Google Cloud Platform](#) webpage.

HYCU for GCP subscription plans

HYCU for GCP offers several subscription plans to better fit your pricing policy to your data protection needs. The subscription plans are divided into two groups based on how you are charged for using the service: the Pay-as-You-Go subscription plan and token-based

subscription plans. You can purchase the Pay-as-You-Go subscription plan on the Google Cloud Platform Marketplace. For token-based subscription plans, contact HYCU sales at info@hycu.com.

- The Pay-as-You-Go subscription plan
Select this plan to be charged for your actual use of HYCU for GCP. For more information, see [“Pay-as-You-Go subscription plan” below](#).
- One of the token-based subscription plans
Select one of these plans to be charged a fixed subscription fee for a specific number of purchased HYCU tokens which corresponds to your planned use of HYCU for GCP. For more information, see [“Token-based subscription plans” on page 20](#).

Plans are billed on a monthly basis and for one day minimum. You can change your subscription plan any time. Plan upgrades take effect immediately while plan cancellations and downgrades take effect at the end of the subscription cycle. To upgrade or downgrade your subscription plan, contact HYCU sales at info@hycu.com.

Pay-as-You-Go subscription plan

This section describes the Pay-as-You-Go subscription plan. It also includes examples of how the subscription fee is calculated in the scope of this plan.

With this subscription plan, you pay for as much data as you protect with HYCU for GCP.


Calculating you data protection price

The price of protecting an instance for a particular time period depends on the following factors:

- Data quantity and time
This factor measures the following product of storage capacity and time period:
(capacity of protected instance disks) × (protection period)
- Pricing tier of protected data
This factor depends on the frequency of backups. For details, see [“Pricing tiers” on the previous page](#).

To calculate the data protection price, use the following formula:

$$\text{(capacity of protected instance disks)} \times \text{(protection period)} \times \text{(pricing tier-dependent price of one pricing unit)} = \text{(undiscounted data protection price)}$$

 **Note** If you use HYCU for GCP to protect a significant amount of data, you automatically qualify for a volume use discount. The discount rate increases with the quantity of your protected data. For the actual criteria defining different discount levels and their corresponding discount rates, see the [HYCU | Marketplace - Google Cloud](#)

[Platform](#) webpage.

Price calculation examples

The following examples illustrate how a data protection price is calculated for various use cases with different instance disk sizes, different RPOs, and different protection periods.

Note The actual measurement unit for reporting instance disk capacity is gibibyte (GiB), although the Google Cloud Platform Console presents it as gigabytes (GB).

Example

Suppose you have the following data protection requirement in your Google Cloud Platform project:

Instance count	1
Persistent disk capacity	100 GiB
Time period for instance protection	2 months
RPO setting in the effective backup policy	2 hours (= platinum tier)

For this requirement, the data protection price calculation is as follows:

$$(100 \text{ GiB} \times 2 \text{ M}) \times \langle \text{PlatinumTierPricePerUnit} \rangle / \text{GiB-M} = \langle \text{Total} \rangle$$

Example

Suppose you have the following data protection requirement in your Google Cloud Platform project:

Instance count	10
Persistent disk capacity per instance	2 TiB (= 2,048 GiB)
Time period for instance protection	6 months
RPO setting in the effective backup policy	4 hours (= gold tier)

For this requirement, the data protection price calculation is as follows:

$$10 \times 2,048 \text{ GiB} = 20,480 \text{ GiB}$$

$$(20,480 \text{ GiB} \times 6 \text{ M}) \times \langle \text{GoldTierPricePerUnit} \rangle / \text{GiB-M} = \langle \text{Total} \rangle$$

Example

Suppose you have the following data protection requirement in your Google Cloud Platform project:

Instance count	24
Persistent disk capacity per instance	5 TiB (= 5,120 GiB)
Time period for instance protection	2 weeks (= 14/30 month ¹)
RPO setting in the effective backup policy	48 hours (= bronze tier)

¹ In the pricing unit, one month represents 30 days.

For this requirement, the data protection price calculation is as follows:


$$24 \times 5,120 \text{ GiB} = 122,880 \text{ GiB}$$

$$(122,880 \text{ GiB} \times 14/30 \text{ M}) \times \langle \text{BronzeTierPricePerUnit} \rangle / \text{GiB-M} = \langle \text{Total} \rangle$$

Token-based subscription plans

This section describes token-based subscription plans. It also includes examples of how the subscription fee is calculated in scope of a token-based subscription plan.

With token-based subscription plans, you pay a fixed fee for the planned number of HYCU tokens that you then consume for data protection with HYCU for GCP within a month. This fee applies also in case when your actual use of HYCU for GCP is lower. If your token consumption exceeds the monthly quota defined in the subscription plan, you are charged an additional fee for the extra tokens that you consume until the end of the month (similarly to the Pay-as-You-Go subscription plan).

 **Note** You can purchase HYCU tokens at a lower fee by choosing a subscription plan that includes more tokens or by choosing a longer subscription period. The fee that you are charged for each consumed extra token also varies depending on the chosen subscription plan and subscription period.

Estimating your HYCU token requirements

With each subscription plan, you purchase a specific number of HYCU tokens. Tokens are consumed by protecting an instance for a particular time period. Depending on the pricing tier of the instance, tokens can be consumed in various quantities.

Pricing tier	Token consumption per gibibyte-hour ¹
platinum	4
gold	3
silver	2
bronze	1

¹ Data quantity refers to the amount of data on a protected instance disk.

The general formula for token consumption calculation is as follows:

$$(\text{capacity of protected instance disks} \times (\text{protection period}) \times (\text{token consumption per GiB-hour}) = (\text{consumed tokens}))$$


For example, if you protect two disks of 500 GiB with a backup policy that corresponds to the silver pricing tier, and the policy remains assigned to the instance for the duration of one month (31 days), the token consumption is calculated as follows:

$$31 \text{ days} = 744 \text{ hours}$$

$$(500 \text{ GiB} \times 2) \times 744 \text{ hours} \times 2/\text{GiB-hour} = 1,488,000$$

Token consumption calculation examples

The following examples illustrate how the count of consumed HYCU tokens is calculated for various use cases with different instance disk sizes, different recovery point objectives (RPOs), and different protection periods.

 **Note** The actual measurement unit for reporting instance disk capacity is gibibyte (GiB), although the Google Cloud Platform Console presents it as gigabytes (GB).

Example

Suppose you have the following data protection requirement in your Google Cloud Platform project:

Instance count	1
Persistent disk capacity	100 GiB
Time period for instance protection	60 days
RPO setting in the effective backup policy	2 hours (= platinum tier)

For this requirement, the calculation of the consumed HYCU tokens is as follows:

$$60 \text{ days} = 1,440 \text{ hours}$$

$$100 \text{ GiB} \times 1,440 \text{ hours} \times 4 / \text{GiB-hour} = 576,000$$

Example

Suppose you have the following data protection requirement in your Google Cloud Platform project:

Instance count	10
Persistent disk capacity per instance	2 TiB (= 2,048 GiB)

Time period for instance protection	1 month (31 days ²)
RPO setting in the effective backup policy	4 hours (= gold tier)

² The month in this example has 31 days.

For this requirement, the calculation of the consumed HYCU tokens is as follows:

$$10 \times 2,048 \text{ GiB} = 20,480 \text{ GiB}$$

$$31 \text{ days} = 744 \text{ hours}$$

$$20,480 \text{ GiB} \times 744 \text{ hours} \times 3 / \text{GiB-hour} = 45,711,360$$

Example

Suppose you have the following data protection requirement in your Google Cloud Platform project:

Instance count	24
Persistent disk capacity per instance	5 TiB (= 5,120 GiB)
Time period for instance protection	2 weeks
RPO setting in the effective backup policy	48 hours (= bronze tier)

For this requirement, the calculation of the consumed HYCU tokens is as follows:

$$24 \times 5,120 \text{ GiB} = 122,880 \text{ GiB}$$

$$2 \text{ weeks} = 336 \text{ hours}$$

$$122,880 \text{ GiB} \times 336 \text{ hours} \times 1 / \text{GiB-hour} = 41,287,680$$

Chapter 3

Subscribing to and signing in to HYCU for GCP

As a managed service that is built on the Google Compute Engine platform, HYCU for GCP does not require installation. You subscribe to HYCU for GCP online from the Google Cloud Platform Marketplace. HYCU then automatically performs the necessary provisioning and activates the service for you.

A single subscription to HYCU for GCP is needed to protect your environment, provided that all desired Google Cloud Platform projects are linked to the billing account selected during the subscription procedure. For more information on how your billing account selection limits the set of protected projects, see [“Extent of data protection” on page 13](#).

The process that leads to an active HYCU for GCP subscription which you can start using is as follows:

1. Subscribe to HYCU for GCP. For instructions, see [“Subscribing to the service” below](#).

This action is usually performed only once for an entire organization.

2. Sign in to the HYCU for GCP web user interface. For instructions, see [“Signing in to HYCU for GCP” on page 25](#).

This action is performed by anyone who wants to protect their data.

The HYCU for GCP web user interface provides information about your subscription. For details, see [“Viewing subscription information” on page 89](#).

Subscribing to the service

Prerequisites

- You have a Google Account.
- *Only if you are using multiple Google Cloud Platform billing accounts.* You have decided for a billing account that you want your data protection cost to be charged to.
- Your billing account meets the following requirements:
 - Your Google Account is granted the Billing Account Administrator (`roles/billing.admin`) role on the billing account. This role is required for purchasing solutions on the Google Cloud Platform Marketplace.

- The billing account has at least one linked Google Cloud Platform project that your Google Account has access to.
- You are signed in to Google, and your currently selected project in the Google Cloud Platform Console is linked to the billing account.

For information on the required procedures, see the Google and Google Cloud Platform documentation.


Considerations

- You cannot change the billing account for an active HYCU for GCP subscription. If you chose a wrong billing account and want to change it, contact HYCU Customer Support.
- The billing account viewer of a HYCU for GCP subscription is preset to the Google Account with which you subscribe to HYCU for GCP. Change of the billing account viewer is part of the process for establishing your data protection environment.
- If you violate the terms of use of HYCU for GCP, 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 Platform](#) webpage.
2. *Only if using Microsoft Edge or Internet Explorer.* Enable pop-ups for the *.cloud.google.com website.
3. Read through the solution description, and then click **SUBSCRIBE TO HYCU**.
4. On the New HYCU subscription page, in the Subscribe pane, check the displayed billing account information, take note of it, and click **Subscribe**.
5. In the Activate pane, click **Register with HYCU, Inc.**
6. On the HYCU Data Protection as a Service for GCP webpage, enter the required information.
7. Verify that the provided information is correct, and then click **Sign up with Google**.
8. In the Choose an account page, specify or select the email address of the same Google Account with which you were already signed in to Google. If needed, enter the corresponding password. Click **Next**.
9. In the Apps with access to your account webpage, review the listed actions for which HYCU for GCP must be granted permissions in scope of access to your Google Account. If you consent to grant the permissions, click **Allow**.

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

 **Note** You can revoke permissions for HYCU for GCP to access your Google Account any time. For instructions, see [“Preventing account access” on page 113](#).

10. On the HYCU Data Protection as a Service for GCP webpage, do one of the following:
 - Review the only billing account.
 - From the drop-down list, select the billing account that you chose for your subscription.
11. Under Projects, review the list of the linked Google Cloud Platform projects. If the billing account is correct, click **Submit**.
Provisioning of HYCU for GCP for your subscription starts.
12. Click **Close & Return to Google**.
13. On the New HYCU subscription page, click the left arrow icon to return to the solution's main page on the Google Cloud Platform Marketplace.
14. Click **Refresh to check status** to track the provisioning progress until the following message appears under the Your HYCU subscription heading:

You activated your HYCU service on <Date>

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

You are now ready to check the sign-in prerequisites and sign in to the web user interface. For instructions, see [“Signing in to HYCU for GCP” below](#).

After the service is activated for your subscription, your trial period starts. During this time, HYCU does not charge you for the backup and recovery service cost. Other shares of the total data protection cost are charged as usual. For more information, see [“Service pricing” on page 15](#). The trial period is limited to a certain time frame and to a certain fee that is initially credited to you by HYCU. For actual figures, see the [HYCU | Marketplace - Google Cloud Platform](#) webpage.

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

Signing in to HYCU for GCP

Prerequisites

- Your HYCU for GCP subscription is ready for use. For instructions on how to subscribe to HYCU for GCP, see [“Subscribing to the service” on page 23](#).
- You have a Google Account.
- Your Google Account has at least the Viewer (roles/viewer) role granted on at least

one Google Cloud Platform project that is linked to the billing account of a HYCU for GCP subscription. This enables you to sign in to the HYCU for GCP web user interface.

- The Google Cloud Platform projects whose instances you plan to protect are linked to the assigned billing account. This enables you to see the projects in the HYCU for GCP web user interface.

For more information on the visibility of projects in HYCU for GCP and the extent of data that you can protect, see [“Extent of data protection” on page 13](#).


- In the Google Compute Engine service, your Google Account has the following roles granted on the projects whose instances 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 buckets you plan to use for storing backup data (backup images, copies of backup images, or data archives).
- The Cloud Pub/Sub API is enabled on the Google Cloud Platform projects whose instances you plan to protect. For guidance on how to enable the API, see the [Quickstart: Using the Console \(UI\) | Cloud Pub/Sub Documentation | Google Cloud](#) webpage.
- You have a web browser installed that the web user interface of HYCU for GCP is compatible with. For a list of supported web browsers, see the *HYCU Data Protection as a Service for GCP Compatibility Matrix*.

For information on the required procedures related to Google and the Google Cloud Platform service suite, see the corresponding documentation from Google.


Procedure

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

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


 **Note** You can revoke permissions for HYCU for GCP to access your Google Account any time. For instructions, see [“Preventing account access” on page 113](#).


After a successful sign-in, the Dashboard panel of the HYCU for GCP web user interface shows information on the selected protection set.

 **Tip** If you sign in for the first time or you do not disable it before, a startup tutorial offers to assist you in getting familiar with a basic data protection setup. To start the tutorial, in the startup tutorial window, click **Begin**. Else, select the **Don't show this tutorial anymore** option or click **Close** as appropriate.

You can manually start the tutorial at any time. To do so, in the web user interface, click **?**, and then select **Startup Tutorial**.

You are now ready to establish your data protection environment and start protecting your data. For instructions, see [“Establishing a data protection environment” on page 28](#). Keep in mind that specific scenarios, for example, use of pre-snapshot and post-snapshot commands, restore of individual files or folders, and configuration of protection set service accounts require you to fulfill additional prerequisites.

 **Important** Your web user interface session expires after 15 minutes of inactivity. At that time, you are automatically signed out and your unsaved changes are lost.

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

Service updates

HYCU for GCP uses the software-as-a-service (SaaS) delivery model. The service is periodically updated to provide you with new features and enhancements. Before updating the service, HYCU sends out a pre-update notice to all customers with active subscription by email, stating the planned update date and time and listing major upcoming improvements. You are given the possibility of asking HYCU to postpone the service update roll-out,

A service update may induce specific configuration and operational changes in your data protection environment. For a list of such changes due to the latest service update, see the *HYCU Data Protection as a Service for GCP Release Notes*.

Chapter 4

Establishing a data protection environment

The first step to establishing a data protection environment is the selection of a Google Cloud Platform billing account. You already took this action while performing the procedure for subscribing to HYCU for GCP. For information on how the billing account affects the extent of data that you can protect, see [“Extent of data protection” on page 13](#).

The following flowchart depicts the first step towards an established data protection environment:

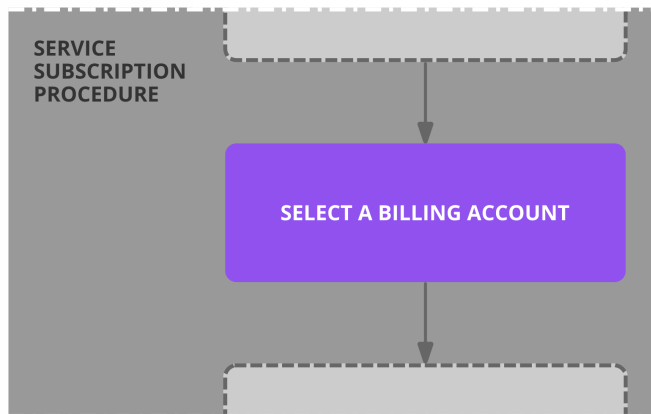


Figure 4-1: Preliminary part of establishing a data protection environment

When your HYCU for GCP subscription is activated, you must sign in to its web user interface and establish a data protection environment in which your data will be effectively protected. For instructions on how to sign in, see [“Signing in to HYCU for GCP” on page 25](#).

The sequence of actions is as follows:

1. *Strongly recommended for continuity of data protection and higher security.* Set a Google Cloud Platform service account as the billing account viewer of your HYCU for GCP subscription. For instructions, see [“Setting the billing account viewer” on page 31](#).

⚠ Caution The Google Cloud Platform identity that is set as the billing account viewer for a HYCU for GCP subscription should never lose permissions of the Billing Account Viewer (`roles/billing.viewer`) role on the billing account used by the

- subscription. If this occurs, you eventually lose access to HYCU for GCP.
2. *Only if you plan to use multiple protection sets.* Configure a protection set and select it. For instructions, see [“Configuring protection sets” on page 93](#) and [“Selecting HYCU for GCP protection sets” on page 31](#).
 3. *Only if you do not want to use automatic buckets.* Set up your own buckets. For instructions, see [“Setting up buckets” on page 45](#).
 4. Choose appropriate predefined policies or create custom policies. For instructions, see [“Choosing between policy types” on page 33](#).
 5. *Only with non-default configurations of SSH server or WinRM server within instances, with enabled OS Login, and for other special use-case scenarios.* Manually configure credential groups and assign them to the instances. For more information and instructions, see [“Enabling access to objects inside instances” on page 46](#).
 6. *Optional.* Configure instance backup options if you want to:
 - Exclude individual disks from a backup
 - Enable a restore of individual files or folders
 - Use pre-snapshot or post-snapshot commands

For instructions, see [“Configuring instance backup options” on page 52](#).

The following flowchart depicts the task flow up to the point when backup policies can be assigned:

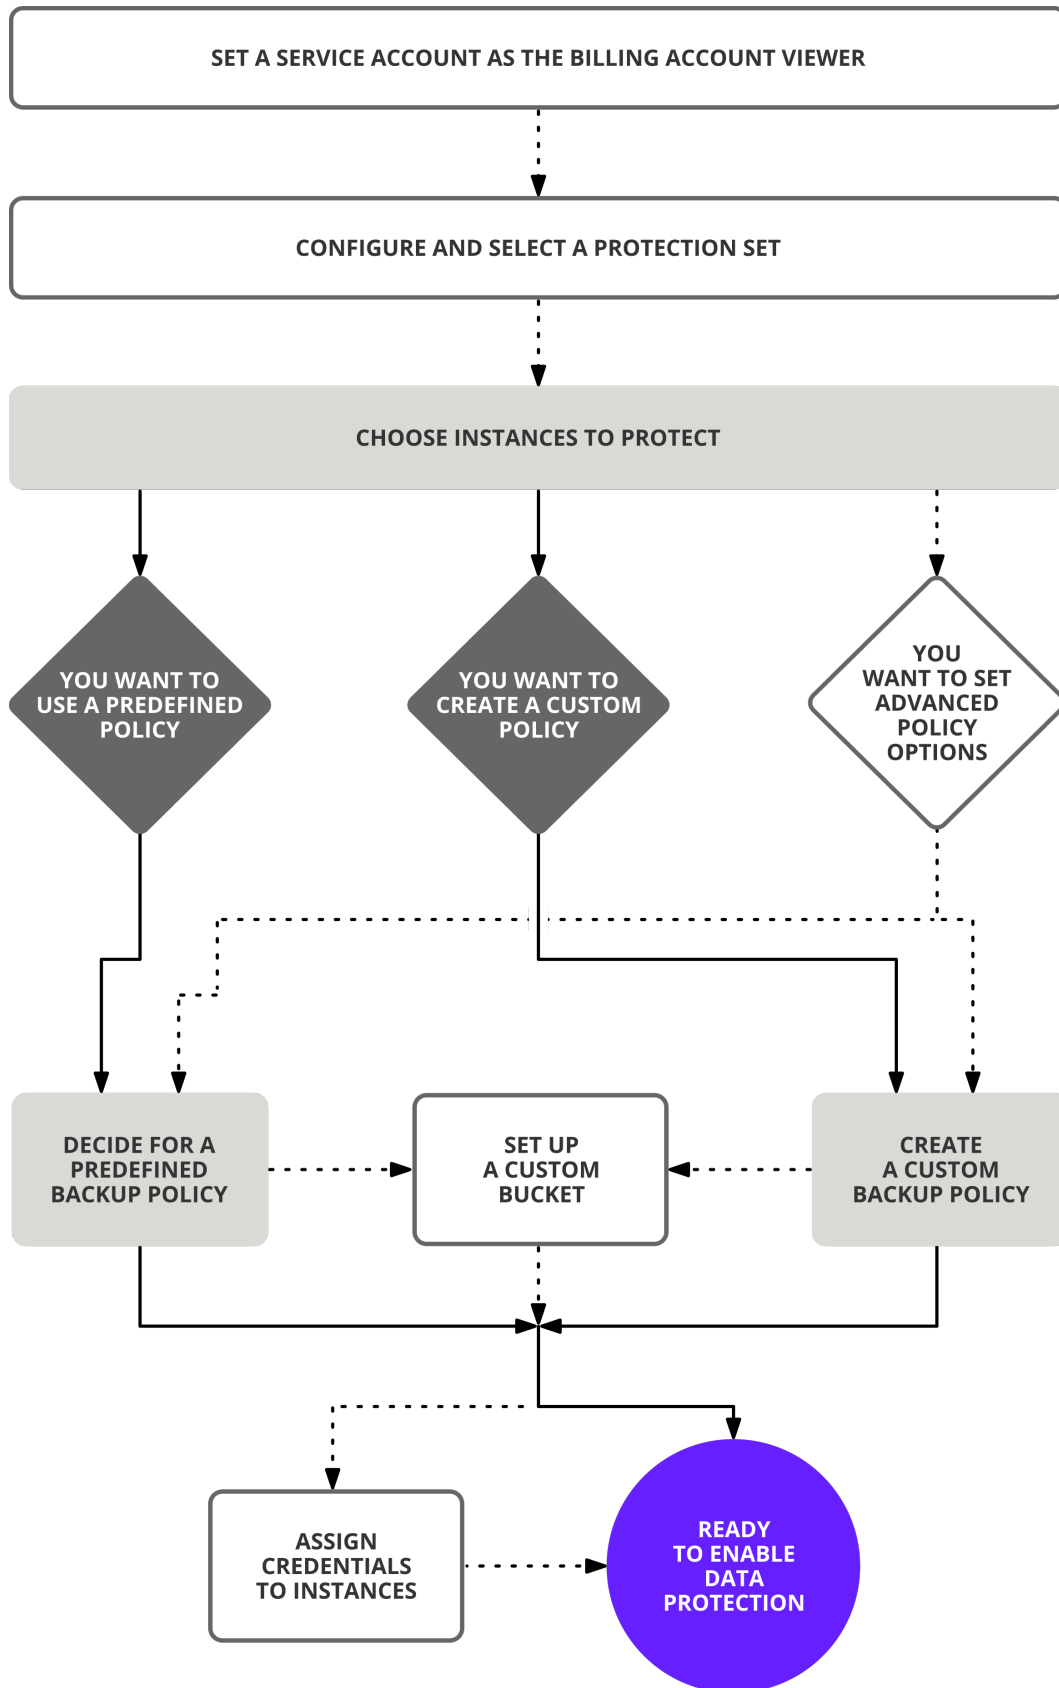



Figure 4-2: Establishing a data protection environment

After the data protection environment is established, data protection can be accomplished in several ways to fulfill your particular business needs. For instructions, see [“Backing up instances” on page 54](#).

Setting the billing account viewer

This section describes the steps that you must follow to properly set the billing account viewer for a HYCU for GCP subscription.


A billing account viewer is a Google Cloud Platform identity (a Google Account or a service account) that HYCU for GCP uses for determining which Google Cloud Platform projects are linked to the billing account.

 **Important** For continuity of data protection and higher security, we strongly recommend that you use a Google Cloud Platform service account rather than a Google Account as the billing account viewer.

Prerequisite

A Google Cloud Platform identity (a Google Account or a service account) that you plan to use as the billing account viewer for a HYCU for GCP subscription has at least the Billing Account Viewer (`roles/billing.viewer`) role granted on the billing account that is used by the subscription.

Procedure

1. *Only if you plan to use a Google Cloud Platform service account.* Check if the service account has been imported into HYCU for GCP and if it has not, import it. For instructions, see [“Importing service accounts” on page 97](#).
2. Click  in the toolbar, and then select **Subscription Information**.
3. In the Subscription Information dialog box, from the Subscription — Billing Account drop-down list, select the billing account name whose HYCU for GCP subscription you want to change the billing account viewer for. The list includes every billing account that is linked from any Google Cloud Platform project which you can access.
4. In the Billing account viewer text box, enter the email address of the chosen identity.
5. Click **Update**.


Selecting HYCU for GCP protection sets

Data protection for Google Cloud Platform projects is managed on the level of project groups called protection sets. Data protection can be enabled independently for each protection set that is configured in HYCU for GCP. This gives you the flexibility to meet specific requirements of different kinds of projects.

By default, all projects that are linked to the billing account of your HYCU for GCP subscription are joined in a single predefined protection set. For information on protection

set management, including configuration of protection sets, see [“Managing protection sets” on page 92](#).

Most actions of a web user interface session are performed in scope of the currently selected protection set.

 **Important** Regardless of your protection set configuration, you can see only the following Google Cloud Platform projects in the web user interface (both conditions apply):

- Projects linked to a billing account that was selected when subscribing to HYCU for GCP.

For the subscription procedure and the implications of choosing a billing account, see [“Subscribing to the service” on page 23](#) and [“Extent of data protection” on page 13](#).

- Projects that you can access with the sign-in user account.

For a list of general prerequisites for using the service, see [“Signing in to HYCU for GCP” on page 25](#).

Selecting a different protection set

You can switch between protection sets that are configured for either the same or a different HYCU for GCP subscription.


Prerequisite


The protection set that you want to select includes at least one Google Cloud Platform project which your sign-in user account can access.

Recommendation


Check HYCU for GCP subscriptions (billing accounts) that relate to your Google Cloud Platform projects. For instructions, see [“Viewing subscription information” on page 89](#).

Procedure

1. In the toolbar, click .
2. In the Protection Set Picker dialog box, from the **Subscription — Billing Account** drop-down list, select the name of the billing account that corresponds to the HYCU for GCP subscription for which the protection set is configured.


The list of configured protection sets within the selected subscription appears. The currently selected protection set (if part of the subscription) is marked with the  icon.

3. From the list of protection sets, select the desired protection set.

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

4. Click **Select**.

The web user interface switches the context to the selected HYCU for GCP protection set.

 **Note** The protection set that you selected last is remembered for the next time you sign in to the HYCU for GCP web user interface.


Defining your backup strategy

HYCU for GCP 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. Backup tasks can be scheduled to meet the specified RPO, run in the preferred time frame, and store backup data into data archives.

When defining your backup policy strategy, take into account the specific needs of your environment and consider the following:

- Recovery point objective (RPO)

A recovery point objective is the maximum tolerable data loss interval. It is a time period during which all incoming changes to the data are tolerated to be lost in case of a disaster.

 **Note** As a backup policy rule, the RPO affects the pricing tier of the instances to which the backup policy is assigned. For more information, see [“Pricing tiers” on page 17](#).

- Backup windows

You can avoid the impact of backup activity during peak hours in your production environment by specifying time frames when backups can be started.

- Data retention and use of data archives

HYCU for GCP can create copies of backup images with longer retention periods. In addition, it can create isolated data archives with multiple archive tiers each of which uses a different archiving cycle and a different retention period.

- Preferred backup policy

You can choose a specific backup policy to protect all instances in all your Google Cloud Platform projects that belong to a protection set. When a backup policy is set as the default backup policy, it is automatically assigned to unprotected instances.

- Policy assignment from the Google Compute Engine service

You can use instance custom metadata tags or instance labels (preferred) in the Google Compute Engine service to automate assignment of backup policies. To achieve this, you may need to properly configure the Labels option in backup policies in HYCU for GCP.

Choosing between policy types

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

- Applying a predefined backup policy

You can use any of the predefined backup policies (platinum, gold, silver, bronze, or exclude) to simplify the data protection implementation. For details, see [“Predefined backup policies” below](#).

- Creating a custom backup policy

If none of the predefined backup policies meets the needs of your backup plan, you can create a new backup policy. For details, see [“Custom backup policies” on the next page](#).


Predefined backup policies

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

HYCU for GCP comes with the following predefined backup policies:

Policy name	Description
platinum	Data is backed up every 2 hours, snapshots are kept for 1 day, copies of backup images are retained for 1 week.
gold	Data is backed up every 4 hours, snapshots are kept for 1 day, copies of backup images are retained for 1 week.
silver	Data is backed up every 12 hours, snapshots are kept for 1 day, copies of backup images are retained for 1 week.
bronze	Data is backed up every 24 hours, snapshots are kept for 2 days, copies of backup images are retained for 1 week.
exclude	Instances are not scheduled for backup and they cannot be manually backed up either.

Default names of the predefined backup policies match the pricing tiers of the instances protected by them. The exclude backup policy does not protect data, therefore it does not match any pricing tier.

 **Important** Adjusting the recovery point objective setting (the Backup every option) in a predefined backup policy may change the pricing tier for the instances protected by it.

Consider also the following:

- Predefined backup policies use automatic buckets for storing backup data. For information on bucket types in HYCU for GCP, see [“Setting up buckets” on page 45](#).
- No backup windows are assigned to predefined backup policies by default. For instructions on how to create and assign backup window specifications, see [“Creating](#)

[backup window specifications” on page 38.](#)

- No data archives are created by predefined backup policies by default. For instructions on how to create data archives, see [“Creating data archives” on page 40.](#)

Custom backup policies

If your data protection needs are not covered with any of the predefined backup policies, you can create a new backup policy. In this case, besides setting the desired backup frequency (recovery point objective—RPO), the retention period for backup data, and the backup image form (backup target type), you can also enable one or more additional policy options. These options are the following:

Policy option	Description
Backup Window	Allows you to choose a backup window specification. Each backup window specification can include one or more time frames during which new backup tasks are allowed to start.
Copy	Allows you to create a copy of the backup image and store it in a separate bucket.
Archiving	Allows you to preserve your data in an archive for future reference.
Labels	Allows you to set up automatic policy assignment based on the instance custom metadata tags or labels (preferred) that are added to the instance in the Google Compute Engine service.

Creating custom backup policies

You can create a custom backup policy that will meet all the needs of your data protection environment.

Prerequisites


- *Only if you plan to select Bucket for the Backup target type option in the backup policy and use a bucket of your choice.* You have created a bucket and set it up in HYCU for GCP. For instructions on how to set up buckets, see [“Setting up buckets” on page 45.](#)
- *Only if you plan to enable the Backup Window policy option.* A backup window specification exists for the selected HYCU for GCP protection set. For instructions on how to create backup window specifications, see [“Creating backup window specifications” on page 38.](#)
- *Only if you plan to enable the Archiving policy option.* A data archive exists for the selected HYCU for GCP protection set. For instructions on how to create data archives, see [“Creating data archives” on page 40.](#)
- *Only if you plan to enable the Labels policy option.* Make sure the following prerequisites are fulfilled:

- In HYCU for GCP, an appropriate service account is assigned to the protection set of that includes projects with the instances which you plan to protect.
- The labels that you plan to specify in HYCU for GCP are already added to the instances in the Google Compute Engine service in the form of either custom metadata tags or labels (preferred). For more information, see the [Storing and Retrieving Instance Metadata | Compute Engine Documentation | Google Cloud](#) and [Labeling Resources | Compute Engine Documentation | Google Cloud](#) webpages.

Considerations


- Storage of snapshots and storage of backup data in buckets are charged by Google according to their Google Compute Engine and Google Cloud Storage pricing policies. For more information, see the [Disks and images pricing | Compute Engine Documentation | Google Cloud](#) and [Cloud Storage pricing | Cloud Storage | Google Cloud](#) webpages.
- If you want that your data is stored as a snapshot and in a bucket, make sure to select the Snapshot backup target type and also enable the Copy option in the backup policy.
- *Only if you plan to enable the Labels policy option.* Consider the following:
 - When matched, the `hycu-policy` custom metadata tag of an instance takes precedence over other custom metadata tags or labels that might be added to the same instance in the Google Compute Engine service. In this case, HYCU for GCP assigns the corresponding backup policy based on this tag. For more information on the `hycu-policy` tag, see [“Assigning backup policies from the Google Compute Engine service” on page 43](#).
 - Labels that you specify in backup policies in HYCU for GCP must be unique within the selected protection set. Reusing a label (both its key and value) for another backup policy in the same protection set is not possible.
 - If an instance has several custom metadata tags or labels added in the Google Compute Engine service and individual labels are specified in different backup policies in HYCU for GCP, you do not have control over which of those policies gets assigned in case of a match. We therefore recommend that you avoid such configurations.
- *Only if you plan to store backup data in a bucket.* Backup and restore speed depends on the region of the chosen bucket and regions of the instances that you assign the backup policy. The optimum speed is achieved when the bucket and the instances reside in the same region.

Accessing the Policies panel

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

Procedure

1. In the Policies panel, click **+ New**. The New Policy dialog box appears.
2. Enter a name and, optionally, a description of your backup policy.
3. Enable the required policy options by clicking them:
 - **Backup** (*mandatory, already enabled by default*)
 - **Backup Window**
 - **Copy**
 - **Archiving**
 - **Labels**
4. In the Backup section, do the following:
 - a. In the Backup every boxes, set the recovery point objective (RPO; in months, weeks, days, hours, or minutes). This rule defines the period between two consecutive backups for the same instance.
 - b. In the Retention boxes, set a retention period (in months, weeks, or days) for the backup images.
 - c. Select one of the following backup target types:
 - **Snapshot:** Keeps snapshots of instance disks in the Google Compute Engine service.

 **Important** If snapshots created by HYCU for GCP are deleted—either from the Google Cloud Platform Console or the gcloud command line—you will not be able to restore from this target type. However, in this case, you can still restore your data from buckets if backup image copies or data archives exist on them.

For information on how the snapshots are named, see [“Objects created by the service” on page 115](#).

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

Example


If your instance resides in the `us-centra1-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-centra11` region.

- **Bucket:** Stores backup data to an automatically selected or user-created bucket in the Google Cloud Storage service.

From the Bucket drop-down list, select a bucket that you want to use for storing data.

If you select the **Automatically selected** option, HYCU for GCP creates a bucket in the region of the instance (or reuses a bucket previously created this way).

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

Enabled option	Procedure
Backup Window	<p>In the Backup Window section, from the Backup window specification drop-down list, 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	<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 a backup image. From the Bucket drop-down list, select a bucket that you want to use for storing data. <p>If you want your bucket to be selected automatically, make sure the Automatically selected option is selected. In this case, HYCU for GCP creates a bucket in the region of the instance (or reuses a bucket previously created this way). If you want to select a user-created bucket, make sure that this bucket is different from the one you selected for the backup.</p>
Archiving	<p>In the Archiving section, from the drop-down list, select a data archive.</p>
Labels	<p>In the Labels section, enter a label key and value, and then click Add. If required, repeat the action as appropriate.</p> <p> Note The backup policy is automatically assigned to all the instances that have at least one matching custom metadata tag or label added in the Google Compute Engine service during the next instance synchronization.</p>


6. Click **Save**.

The custom backup policy is created and added to the list of backup policies. For details on managing backup policies, see [“Managing backup policies” on page 87](#).


Creating backup window specifications

HYCU for GCP enables you to define time frames during which scheduled backup tasks are allowed to start. You can use time frame definitions to prevent your data protection environment from becoming overloaded. For example, you can schedule your backup tasks to run outside production hours to reduce the load during peak hours.


You can use backup window specifications with both predefined backup policies and custom backup policies.


 **Important** When defining a backup window specification, check if the recovery point objective (RPO) of each affected backup policy can be achieved with the specified time frames. If the RPO period is shorter than any time frame during which backups are not allowed to start, the backup tasks will be delayed and make the affected instances and backup policies non-compliant.

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**.
2. In the Backup Window dialog box, click **+ New**.
3. In the New dialog box, enter a name for your backup window specification and, optionally, a description.
4. From the Time Zone drop-down list, select the time zone for the backup window specification.
5. Select the days and hours during which you allow backups to run. Click and drag to quickly select a time frame that includes your preferred days and hours. Click a day label to select that entire day. Click an hour label to select that hourly period on all days of the week.

The selected time frames are displayed as entries in the Time Frames box. To delete a selected time frame, click  next to its entry.

6. Click **Save**.
7. In the Backup Window dialog box, click **Close**.

You can later edit any of the existing backup window specifications (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 specification, you can do the following:

- Specify a backup window specification when creating a new backup policy. For details, see [“Creating custom backup policies” on page 35](#).
- Assign a backup window specification to an existing backup policy. To do so, select the backup policy, click  **Edit**, and then make the required modifications.

Example

You have selected the bronze backup 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 6 AM, Saturdays and Sundays all day

TIMEZONE
(UTC-05:00) Eastern Time

SCHEDULE ?

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

TUE 18:00-06:00

WED 18:00-06:00

THU 18:00-06:00

FRI 18:00-06:00

SAT 00:00-24:00

SUN 00:00-24:00

TIME FRAMES

< 00:00 - 06:00 X 18:00 - 24:00 X 00:00 - 06:00 X 18:00 - 24:00 X 00:00 - 06:00 X 18:00 - 24:00 X 00:00 - 06:00 X 18:00 - 24:00 X 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 frame.

Creating data archives

HYCU for GCP enables you to create archives of your protected data and keep them for a longer period of time. By archiving protected data, the data is stored for future reference on a daily, weekly, monthly, or yearly basis. When included in a data archive, your data is isolated from other data protection activities and safely stored in a Google Cloud Storage bucket of your choice.

Prerequisite

Only if you plan to use a bucket of your choice for the data archive. You have created a bucket and added it to buckets of the selected protection set in HYCU for GCP.

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**.
2. In the Archiving dialog box, click **+ New**. The New dialog box appears.
3. Enter a name for your data archive and, optionally, a description.
4. Enable the required archiving options by clicking them:

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. For each enabled archiving option, do the following:
 - a. In the Start at text boxes, specify the hour and the minute when the archive task should start. From the Time zone drop-down list, specify the appropriate time zone.
 - b. *Not applicable to daily archives.* Provide information about the day when to archive data.
 - c. In the Retention box, set the retention period to be used.

 **Note** Make sure that the retention period is longer than the RPO set for the backup images. This is required to prevent the archive from expiring before a new backup is performed.

- d. From the Bucket drop-down list, select a bucket that you want to use for storing the data archive.



If you select the **Automatically selected** option, HYCU for GCP creates a bucket in the region of the instance (or reuses a bucket previously created this way).

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


If you select the **Automatically selected** option, a suitable storage class is automatically selected depending on the specified retention. If you manually select a storage class, make sure that your choice does not negatively affect your storage costs.

For description of storage classes in the Google Cloud Storage service, see the [Storage classes | Cloud Storage | Google Cloud](#) webpage.

6. 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 no longer need (click  **Delete**). Keep in mind that you cannot modify an archive bucket if an archiving task is in progress on that bucket.

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

- Specify a data archive when creating a new backup policy. For details, see [“Creating custom backup policies” on page 35](#).
- Include a data archive into an existing backup policy. To do so, select the policy, click  **Edit**, and then make the required modifications.

Setting default backup policies


If you consider one of the predefined or custom backup policies satisfies all your data protection needs, you can set that policy as the default backup policy in the selected protection set. HYCU for GCP then automatically assigns the policy to all newly identified instances. You can also assign the backup policy to each previously identified instance without an assigned policy.

 **Important** Automatic assignment of default backup policies is overridden by assignment of policies from the Google Compute Engine service. For more information, see [“Assigning backup policies from the Google Compute Engine service” on the next page](#).


Prerequisite

All relevant prerequisites that also apply for manual backup policy assignment are fulfilled. For details, see [“Assigning backup policies manually” on page 55](#).


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 default, and then click  **Set Default**.
2. In the Set Default Policy dialog box, do one of the following:
 - Click **Yes** to assign the default backup policy to all instances that do not have a backup policy assigned (that is, existing and newly identified ones).
 - Click **No** to assign the default backup policy only to newly identified instances.

HYCU for GCP marks the default policy with the  icon.

If you later decide not to use this backup policy as the default any more, click  **Clear Default**. This action does not unassign the backup policy from the instances to which it is assigned.

Assigning backup policies from the Google Compute Engine service

You can utilize the Google Compute Engine service itself to automate policy assignment in HYCU for GCP. When you add proper custom metadata tags or labels to instances in the Google Compute Engine service, HYCU for GCP can identify them during the next instance synchronization and automatically assigns corresponding backup policies.

This approach is especially useful within a complex data protection environment, and when your data protection approach requires use of numerous backup policies. It enhances the automatic policy assignment as follows:

- It offers you more detailed control over automatic policy assignment than the default policy provides.
- It enables mass assignment of multiple policies.
- You can use it for mass exclusion of instances from data protection—with the help of the predefined exclude policy,

To automate policy assignment from the Google Compute Engine service, you have two options:

- Use of your instance custom metadata tags or instance labels (preferred)
This option requires proper configuration of the Labels option in your backup policies.
- Use of the `hycu-policy` custom metadata tag
This option requires adding the tag to the instances that you plan to protect.

Prerequisites

- In HYCU for GCP, an appropriate service account is assigned to the protection set that includes projects with the instances which you plan to protect.
- All relevant prerequisites that also apply for manual backup policy assignment are fulfilled. For details, see [“Assigning backup policies manually” on page 55](#).

Consideration

Automated backup policy assignment from the Google Compute Engine service takes precedence over any other policy assignment from within HYCU for GCP (manual assignment, default policy). This means that for each identified instance with proper metadata tag or label, the tag or label defines which backup policy gets assigned to the instance, even for instances with already assigned policies.

Assigning backup policies based on your custom metadata tags or labels

To assign backup policies to instances from the Google Compute Engine service by using your instance custom metadata tags or instance labels (preferred), enable the Labels

option in your backup policies and specify appropriate keys and values. For instructions, see [“Creating custom backup policies” on page 35](#).

⚠ Important If an instance has a backup policy assigned by means of a corresponding metadata tag or label in the Google Compute Engine service, a subsequent policy unassignment or manual assignment of a different policy in HYCU for GCP causes the tag or label to be automatically deleted. Depending on the keys and values specified in the policy, multiple tags or labels may be deleted.

Assigning backup policies based on the hycu-policy custom metadata tag

Procedure

1. In the Google Cloud Platform Console, choose a Google Cloud Platform project whose instances you want to protect.
2. Within the project, choose an instance and add it the hycu-policy custom metadata tag. Use the following data:

Key	Value
hycu-policy	<PolicyName>

In the above table, <PolicyName> is the name of a backup policy that is configured for the protection set which includes the chosen Google Cloud Platform project in HYCU for GCP. If you specify exclude for the tag value, the instance will be assigned the exclude policy.

💡 Tip Custom metadata tags can also be added from the gcloud command line or by using the Google Cloud Platform API. For instructions, see the [Storing and Retrieving Instance Metadata | Compute Engine Documentation | Google Cloud](#) webpage.

3. Repeat step 2 for each additional instance to which you want to assign a policy this way.
4. Sign in to the HYCU for GCP web user interface.
5. Select the protection set that includes the same Google Cloud Platform project as you selected in step 1. For instructions on selecting protection sets in HYCU for GCP, see [“Selecting HYCU for GCP protection sets” on page 31](#).
6. In the navigation pane, click **Instances**.
7. Click **Synchronize**.

In the Instances panel, the names in the Policy column indicate a successful assignment of policies to instances.

When you modify the value of an existing hycu-policy custom metadata tag, the change is propagated in the same way as when the tag is added.

⚠ Important If an instance has a backup policy assigned by means of the

`hycu-policy` custom metadata tag, a subsequent policy unassignment or manual assignment of a different policy in HYCU for GCP deletes the tag from the instance.

Setting up buckets

Buckets are storage containers in the Google Cloud Storage service that store backup images, copies of backup images, and archives of protected data. For storing backup data, HYCU for GCP can use two bucket types, depending on how a backup policy is configured:

- User-created buckets

Buckets of this type are the buckets that you already created in the Google Cloud Storage service earlier. When configuring a backup policy, you can select a desired bucket from the currently selected HYCU for GCP protection set.

Only buckets of this type require your involvement to be set up.

- Automatic buckets

Buckets of this type are created automatically during a backup, and you can view them in the Buckets panel of the HYCU for GCP web user interface. They are created in the same Google Cloud Platform project and at the same location (the Google Cloud Storage region) as the backed up instances. The same automatic bucket is used for backup data of multiple instances where possible.

For information on how the automatic buckets are named, see [“Objects created by the service” on page 115](#).

Automatic buckets are visible to the Google Cloud Storage service in the same way as the buckets you create yourself. You can use them for storing your data, for example, for individual files and folders that you restore from backup images. When automatic buckets no longer contain backup images, they remain available in your project.

⚠ Caution Never delete any buckets used by HYCU for GCP. Deletion of a bucket that stores backup images results in data loss. Additionally, within buckets, ensure that the `hycu/backups/` folders are always kept intact.

Setting up user-created buckets

Setting up a user-created bucket means making the bucket available for backup and restore purposes within a protection set in HYCU for GCP.


Prerequisite

Your sign-in user account (or the protection set service account, if configured) has access to the bucket.

Consideration


You can set up the same bucket in multiple protection sets.

Accessing the Buckets panel

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

Procedure

1. In the Buckets panel, click **+ Add**. The Add Bucket dialog box appears.
2. In the Bucket text box, enter the name of the desired bucket.
3. For the Size option, specify the amount of bucket storage space that is allowed to be used by backup data created by HYCU for GCP. To do this, edit the value in the Size text box and select a unit of measurement from the drop-down list as required.

 **Note** The specified amount represents a soft limit, therefore actual usage may exceed it. To prevent the bucket's storage space from running out, we recommend that you monitor its usage in the Buckets panel, in the Health and Utilization columns and in the Details section.

The soft limit adds to the joint bucket capacity that is shown within the Buckets widget in the Dashboard panel.

4. Click **Save**.

The bucket is added to the list of buckets and can be used for configuring backup policies. For details on managing buckets, see ["Managing buckets" on page 84](#).

Enabling access to objects inside instances

When your data protection goals require handling individual objects in the file systems of your instances, HYCU for GCP must be able to access them. Enabled access to objects is a prerequisite for achieving the following goals:

- Enable restore of individual files or folders.
- Restore individual files or folders to the original instance.
- Run pre-snapshot or post-snapshot commands.

By default, HYCU for GCP automatically uses the following connection parameters to access the objects by using a network service protocol:


Guest operating system	Google Cloud Platform identity user name	Network service protocol	Port	Transport protocol
GNU/Linux	<UserName> where <UserName>@<DomainName> is the email address of the identity that is running the effective task in HYCU for GCP	SSH	22	Not applicable
Microsoft Windows	hycu	WinRM	5986	HTTPS

For instances running Microsoft Windows, HYCU for GCP also automatically configures an appropriate credential group, named `auto-<InstanceName>`, and assigns it to the instance. The default connection parameters are suitable for the majority of instance configurations. However, in the following cases you must manually enable access to the instances by assigning them proper credential groups in HYCU for GCP:

Guest operating system	Configuration details or use-case scenarios
any	<ul style="list-style-type: none"> You plan to restore individual files or folders by using a user account that you specify. You plan to use a specified user account for a restore, either to reuse an already existing user account or to comply with policies that impose restrictions on the utilized user names and passwords.
GNU/Linux	<ul style="list-style-type: none"> SSH server is configured to use a non-default TCP port. SSH server is configured to use public key authentication. OS Login is enabled on the instance in the Google Compute Engine service. For more information on OS Login as the instance access method, see the OS Login Compute Engine Documentation Google Cloud webpage. You plan to use pre-snapshot or post-snapshot commands and invoke them with a user account that you specify instead of the Google Cloud Platform identity which is used by default.
Microsoft Windows	<ul style="list-style-type: none"> WinRM server is configured to use the HTTP transport protocol or a non-default TCP port. You plan to use pre-snapshot or post-snapshot commands.

Configuring and assigning credential groups manually

A credential group is a configuration entity in HYCU for GCP that you can assign to an instance to enable access to and proper handling of its file system objects.


 **Note** Each credential group may include the following settings: access protocol, underlying transport protocol (for WinRM), network service client listening port, user name, password, private key associated with the user account.

Prerequisites


- A user account with sufficient privileges is configured within each instance. For instances running Microsoft Windows, you can configure the user account from the Google Cloud Platform Console or the gcloud command line. For more information, see the [Creating Passwords for Windows Instances | Compute Engine Documentation | Google Cloud](#) webpage.
- *For instances running GNU/Linux:*
 - *Only if the Authentication option in HYCU for GCP is set to either Password authentication or Private key authentication:* Ensure the following within the instance:
 - The specified user account is a member of the sudo user group.
 - The following line is included in the /etc/sudoers file:



```
<UserName> ALL=(ALL) NOPASSWD: /bin/lsblk, /bin/ls, /bin/mkdir, /bin/mv, /bin/umount, /bin/cp, /bin/rm, /bin/mount
```
 - *Only if you want HYCU for GCP to access the instance by using a specific user account with password authentication.* The SSH server is configured to allow password authentication for logging on to the instance.

Accessing the Instances panel

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

Procedure

1. In the Instances panel, select an instance to which you want to assign a credential group.
2. Click  **Credentials**. The Credential Groups dialog box appears.
3. Click **+ New**.
4. In the Credential group name text box, enter a name for the new credential group.
5. From the Protocol drop-down list, select one the following:
 - **Automatic:** Select this option to allow HYCU for GCP to automatically choose a protocol for accessing the objects in the instance file systems, and use the supplied user name and password to connect to the instance. This option is compatible with the default configuration of the network service server (SSH, WinRM) in the instance.
 - **SSH:** Select this option to use the SSH protocol for instance access and to enable the credential group adjustment for your actual SSH server configuration.
 - **WinRM:** Select this option to use the WinRM protocol for instance access and to enable the credential group adjustment for the actual WinRM server configuration.
6. Depending on your selection in the previous step, proceed as follows:




Automatic	<p>a. In the Username text box, enter the user name of the user account that has required permissions to access the instance, in the following format:</p> <ul style="list-style-type: none"> • GNU/Linux: <code><LocalOrDomainUserName></code> • Microsoft Windows: <code><LocalUserName></code>, <code><Domain>\<DomainUserName></code>, <code><DomainUserName>@<Domain></code> <p>b. In the Password text box, enter the corresponding password.</p>
SSH	<p>a. In the Port text box, enter the port number used by the SSH server in the instance.</p> <p>b. From the Authentication drop-down list, select one of the following and proceed accordingly:</p> <ul style="list-style-type: none"> • Automatic: This option enforces the same behavior as if no credential group is assigned to the instance, but adds the possibility to adjust the port number used in connections to the instance.  Important Do not select this option if OS Login is enabled on your instance. • Password authentication: Select this option to allow HYCU for GCP to use the specified user name and password for connections to the instance. Follow these steps: <ol style="list-style-type: none"> i. In the Username text box, enter the user name of the user account that has required permissions to access the instance, in the <code><LocalOrDomainUserName></code> format. ii. In the Password text box, enter the corresponding password. • Private key authentication: Select this option to allow HYCU for GCP to use the specified user name and the corresponding private key for connections to the instance. This selection is mandatory for utilization of the OS Login instance access method in the Google Compute Engine service or connection to an instance in which the SSH server is configured to use public key authentication. For more information, see the Choosing an access method Compute Engine Documentation Google Cloud and Setting up OS Login Compute Engine Documentation Google Cloud webpages. Follow these steps: <ol style="list-style-type: none"> i. In the Username text box, enter the user name of the user


	<p>account that has required permissions to access the instance, in the <code><LocalOrDomainUserName></code> format.</p> <p>ii. Click Browse. In the dialog box, browse for and then select the file with the private key of the user account whose user name you specified in the previous step. Click Open.</p> <p>For information on how to obtain the private key, see the Managing SSH keys in metadata Compute Engine Documentation Google Cloud webpage.</p> <p>iii. <i>Only when the private key is encrypted with a passphrase.</i> In the Private Key Passphrase text box, enter the effective passphrase.</p>
WinRM	<p>a. From the transport drop-down list, select the transport protocol of the WinRM server in the instance.</p> <p>b. In the Port text box, enter the port number used by the WinRM server in the instance.</p> <p>c. In the Username text box, enter the user name of the user account that has required permissions to access the instance, in the <code><LocalOrDomainUserName></code> format.</p> <p>d. In the Password text box, enter the corresponding password.</p>

7. Click **Save**.

8. Click **Assign**. The name of the assigned credential group appears in the Credential group column of the Instances panel, and the instance discovery starts.

Check the discovery task status to verify that the connection to the instance succeeded. You can always unassign the credential group from an instance or replace it with a different one.

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

 **Tip** You can use multiple selection to assign the same credential group to multiple instances or to unassign credential groups from multiple instances.

Chapter 5

Protecting instances

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

Limitations

- Local SSDs are not protected. Local SSD is a solid-state drive physically attached to the server that hosts the instance. The Google Compute Engine service does not support creating snapshots of local SSDs.
- Instance memory is not protected. The Google Compute Engine service only provides functionality for creating snapshots of instance disks.
- Based on how snapshots are taken in the Google Compute Engine service, instance disks are backed up as individual entities. Crash consistency of backup data is therefore guaranteed only for each disk individually.

Considerations

- To optimize the use of storage space in the Google Cloud Storage buckets, HYCU for GCP employs the HYCU changed block tracking technique on the backup data before storing it. Snapshots are divided into chunks which are checked for uniqueness (single instancing). The technique is applied to primary backup data as well as to backup image copies and data archives.
- Data in instances' backup images, copies of backup images, and data archives that HYCU for GCP creates is crash-consistent, but it may not always be application-consistent.

If you use appropriate pre-snapshot commands, you can achieve universal application consistency of backup data. If pre-snapshot commands are not provided, the application consistency of backup data is limited to (all conditions apply):

- Instances that are running Microsoft Windows.
- Instances that are based on an operating system image version v20160810 or later.
- Disk volumes that use one of the following file systems: NTFS, exFAT, ReFS (Resilient File System).
- Applications that are implemented as VSS writers for the Microsoft Volume Shadow

Copy Service (VSS).

- Applications that store their data on a single disk.

For more information, see the [Creating a Windows Persistent Disk Snapshot | Compute Engine Documentation | Google Cloud](#) webpage. For recommendations on how to improve application consistency with other operating systems and applications, see the [Best Practices for Persistent Disk Snapshots | Compute Engine Documentation | Google Cloud](#) webpage.

- Not all types of tasks can be running simultaneously for the same original instance. Consider the following:
 - No more than one backup task can run at the same time.
 - No more than one restore task (instance restore, instance cloning, or restore of individual files or folders) can run at the same time. For example, while HYCU for GCP is cloning an instance with a restore task, you cannot restore individual files or folders to it.
 - A backup task and a restore task can run in parallel.

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

- [“Configuring instance backup options” below](#)
- [“Backing up instances” on page 54](#)
- [“Restoring instances” on page 58](#)
- [“Restoring individual files or folders” on page 62](#)

Configuring instance backup options

Before you start protecting instances, you can adjust instance protection to the needs of your data protection environment by using instance backup options. These options allow you to:

- Exclude individual disks from a backup.
- Enable a restore of individual files or folders.
- Specify commands to be run just before or immediately after HYCU for GCP takes a snapshot of the instance. You can use such pre-snapshot and post-snapshot commands to prepare the instance or an application running on it for a backup or to perform any additional tasks afterward.


The backup options are usually configured for each individual instance. By using multiple selection, you can also configure them for a set of instances within the same protection set.

Considerations

- If you enable a restore of individual files or folders, keep in mind that this scenario has further prerequisites which must be fulfilled:


- *Only for original instances running Microsoft Windows.* At the time of the backup. For details, see the prerequisite list in the section [“Assigning backup policies manually” on page 55.](#)
- At the time of the restore of individual files or folders. For details, see the prerequisite list in the section [“Restoring individual files or folders” on page 62.](#)
- If you specify pre-snapshot or post-snapshot commands, keep in mind that further prerequisites must be fulfilled at the backup time. For details, see the prerequisite list in the [“Assigning backup policies manually” on page 55](#) section.

Accessing the Instances panel

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

Procedure

1. In the Instances panel, select an instance.

 **Tip** To configure the same backup options for multiple instances at once, select the desired 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 appears.

3. Depending on what you want to do, provide the required information:

- Exclude disks from a backup:

In the Disks drop-down list, click each desired disk to mark it as excluded from or included in the backup tasks.

 **Important** You cannot exclude the boot disk from a backup.

- Enable a restore of individual files or folders:


Use the **Enable restore of individual files** switch if you want to enable a restore of individual files or folders.

- Specify the commands to be run just before or immediately after HYCU for GCP takes a snapshot of the instance:

- In the Pre-Snapshot Command text box, enter the command that HYCU for GCP runs just before it creates a snapshot of the instance.
- In the Post-Snapshot Command text box, enter the command that HYCU for GCP runs immediately after it creates a snapshot of the instance.

The commands are run from the home directory of the user account that HYCU for GCP uses for running the commands. Depending on the operating system on the instance, the following user accounts are used:

- GNU/Linux:
 - The instance has no credential group assigned in HYCU for GCP: The user account that runs the backup task. For interactive backups, this is the sign-in user account. For scheduled backups, this is either the policy assignor (if no service account is assigned to the protection set) or the service account assigned to the protection set.
 - The instance has a credential group assigned: The user account specified in the credential group.
- Microsoft Windows: The user account that is assigned to the instance in HYCU for GCP through a credential group.

 **Important** A snapshot is created even if the pre-snapshot command fails. The post-snapshot command is run even if the pre-snapshot command, snapshot creation, or both actions fail. When a pre-snapshot or post-snapshot command returns an error, the status of the restore point is set to Done with errors.

Example

Examples of the pre-snapshot and post-snapshot commands for different instance operating systems.

GNU/Linux:

```
bash /home/<UserName>/freeze_db.sh
bash /home/<UserName>/thaw_db.sh
```

Microsoft Windows:

```
%USERPROFILE%\quiesce_db.bat
%USERPROFILE%\resume_db.bat
```

4. Click **Save**.

Backing up instances

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

Methods of invoking backups

Continuous protection of an instance is achieved by assigning it an appropriate backup policy. While the assigned policy automatically invokes periodic backups, you can also back up the instance manually any time. For instructions on how to back up instances manually, see [“Performing manual backups” on page 82](#).

Methods of assigning backup policies

Backup policies can be assigned to instances:

- Manually—for the instances selected in the HYCU for GCP web user interface.
See [“Assigning backup policies manually” below](#).
- Automatically—by configuring a default backup policy in HYCU for GCP for the corresponding protection set.
See [“Setting default backup policies” on page 42](#).
- Automatically—by utilizing existing instance custom metadata tags or labels (preferred) in the Google Compute Engine service and configuring proper mapping in the backup policy configuration.
See [“Assigning backup policies from the Google Compute Engine service” on page 43](#).
- Automatically—by adding a special custom metadata tag to instances in the Google Compute Engine service.
See [“Assigning backup policies from the Google Compute Engine service” on page 43](#).

Assigning backup policies manually

Prerequisites when planning to restore individual files or folders

- A restore of individual files or folders is supported on the operating system that is running on the original instance. See the *HYCU Data Protection as a Service for GCP Compatibility Matrix*.
- The source disk volume uses one of the supported file systems. See the *HYCU Data Protection as a Service for GCP Compatibility Matrix*.
- *Only for instances running Microsoft Windows:* The following prerequisites are fulfilled:
 - In the Google Cloud Platform Console, there is a network firewall rule applied to the instances—either to the entire network or to individual instances through the use of network (target) tags. For each target 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. To do so, add the following source tag to the network firewall rule:

Source tag

hycu-*<String>*

In the template, *<String>* denotes any string that is allowed to be specified in firewall rule tags.

For instructions on how to configure and apply the network firewall rule, see a corresponding Google Cloud Platform webpage:

- [Using Firewall Rules | VPC | Google Cloud](#)
- [Configuring Network Tags | VPC | Google Cloud](#)
- On the instances that you plan to protect, WinRM is configured to use a TCP port (by default, 5986) for HTTPS or HTTP communication. The firewall is configured to enable inbound network traffic through this port.
- Restore of individual files or folders is enabled for the instance. For instructions on how to enable restore of individual files or folders, see [“Configuring instance backup options” on page 52](#).
- *Only if you want to use custom credentials for the restore.* Correct credential group is assigned to the original instance, and the corresponding credentials belong to a user account with sufficient privileges. For instructions on how to assign access credentials, see [“Enabling access to objects inside instances” on page 46](#).

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

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

Optionally, you can make the network firewall rule more restrictive so that it allows network traffic only from legitimate sources. To do so, add the following source tag to the network firewall rule:

Source tag

hycu-*<String>*

In the template, *<String>* denotes any string that is allowed to be specified in firewall rule tags.


For instructions on how to configure and apply the network firewall rule, see a corresponding Google Cloud Platform webpage:

- [Using Firewall Rules | VPC | Google Cloud](#)
- [Configuring Network Tags | VPC | Google Cloud](#)
- *Only for instances running GNU/Linux:* On the instances that you plan to protect, an SSH server is installed and configured to use a TCP port (by default, 22) for SSH


communication. The firewall is configured to enable inbound network traffic through this port.

- *Only for instances running Microsoft Windows:* On the instances that you plan to protect, WinRM is configured to use a TCP port (by default, 5986) for HTTPS or HTTP communication. The firewall is configured to enable inbound network traffic through this port.
- *Only for instances running Microsoft Windows, and for instances running GNU/Linux with non-default configuration of SSH server or if you want to use custom user accounts for the command invocation.* A 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 objects inside instances” on page 46](#).


Accessing the Instances panel


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

Procedure

1. Select a protection set. For instructions, see [“Selecting HYCU for GCP protection sets” on page 31](#).
2. Select the instances that you want to back up. You can update the instance list by clicking  **Synchronize**. In a protection set with a large number of projects, the update may take a while.


To narrow down the list of displayed instances, use the filtering options as described in [“Filtering and sorting data in panels” on page 70](#).

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


 **Note** The exclude policy serves as a "container" for unprotected instances.

5. Click **Assign** to assign the backup policy to the selected instances.

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


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

You can also perform a manual backup of individual instances at any time. For details, see [“Performing manual backups” on page 82](#).


 **Tip** Assign the exclude policy to instances not needing protection, so that the corresponding statistic in the Dashboard panel remains accurate.

Restoring instances

HYCU for GCP enables you to restore either an entire instance or only some of its disks to a specific point in time. For both use cases, you can select between the following restore options:

Restore option	Description
Restore instance	<p>Enables you to restore an instance to its original location with original settings. Select this option if you want to replace the original instance or its individual disks with the restored instance or individual restored disks. For instructions, see “Restoring instances or their disks” below</p> <p> Important HYCU for GCP does not delete user-created objects (instances, instance disks) from the Google Compute Engine service. A manual delete action is required before you initiate the process for restoring data to its original location.</p>
Clone instance	<p>Enables you to create a clone of the original instance by restoring it or its individual disks to a new location with custom settings. Select this option if you want to keep the original instance. For instructions, see “Cloning instances” on the next page.</p>

Accessing the Instances panel


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

Restoring instances or their disks

You can restore an instance or its individual disks to their original location with original settings.

Prerequisite

The objects that you plan to restore are manually deleted from the Google Compute Engine service through the Google Cloud Platform Console or the gcloud command line.

 **Caution** Unprotected data (new data that has been stored on the instance after its last successful backup) is lost during the deletion.

The following table lists objects that you must delete in different restore scenarios.

Object to be restored	Objects that should be deleted
An entire instance (all instance disks)	The instance and all instance disks ¹
A boot disk	The instance and its boot disk ²


A non-boot disk	That particular disk
-----------------	----------------------


¹ In the Google Compute Engine service, disks may be configured for automatic deletion when the instance they are attached to is deleted.


² In the Google Compute Engine service, boot disks may be configured for automatic deletion when the instance they are attached to is deleted. Deleting the instance and its boot disk does not delete other disks that are attached to the instance and have been configured to be kept in case of instance deletion. You must attach such disks to the restored instance.

Procedure

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

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

2. In the Details section that appears at the bottom of the screen, select the desired restore point.
3. Click  **Restore Instance**. The Instance Restore Options dialog box opens.
4. Select **Restore**, and then click **Next**.
5. From the Disks drop-down list, select the instance disks that you want to restore.

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

6. Click **Restore**.

Cloning instances

To create a clone of the original instance, restore it with a new name to the original or a different location. When cloning, you can change the following properties: selection of the backed up disks, destination Google Cloud Platform project, Google Cloud Platform region and zone, instance network configuration.

Prerequisite

Only for cloning into a non-original project. In the Google Compute Engine service, your sign-in user account has the following roles granted on the project where you plan to clone your instances to:



- Compute Admin (`roles/compute.admin`)
- Service Account User (`roles/iam.serviceAccountUser`)


Limitation


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

Procedure

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

 **Note** The Details section appears only if you click an instance. Selecting the check box before the name of the instance will not open the Details section.
2. In the Details section that appears at the bottom of the screen, select the desired restore point.
3. Click  **Restore Instance**. The Instance Restore Options dialog box opens.
4. Select **Restore As**, and then click **Next**.
5. In the New Instance Name text box, specify a new name for the instance.
6. From the Disks drop-down list, select the instance disks that you want to restore.

 **Note** All disks of the instance are preselected for cloning. With such selection, the entire instance is cloned. The boot disk is cloned even if you do not select it.
7. From the Target Project drop-down list, select the project that you want to clone the instance to. The original project of the instance is preselected. You can choose from projects that are linked to the billing account of your HYCU for GCP subscription and that your sign-in user account can access.
8. From the Target Region and Target Zone drop-down lists, select the Google Cloud Platform region and zone to clone the instance to. The original region and zone of the instance are preselected.
9. Under Network Interfaces, review a 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



 **Note** You can add an additional network interface, edit settings of an existing network interface, or remove a network interface. For each configured network interface, you can separately adjust its external and internal IP address types. By default, cloning keeps external IP address configuration of the original instance.
10. *Only if you need to reconfigure network settings for the instance clone.* Do the following:

To add a network interface, follow these steps:

 - a. Click **Add Network Interface**.
 - b. From the Target Networks drop-down list, select a network that you want to add the cloned instance to. You can choose among the networks configured in the project of the instance and other networks that your sign-in user account has

access to.

- c. Select the external address type for the network interface. For more information, see [“External network address types”](#) below.
- d. *Only if required by the selected external address type.* From the External Address drop-down list, select the name of the desired external IP address resource.
- e. Select the internal address type for the network interface. For more information, see [“Internal network address types”](#) below.
- f. *Only if required by the selected internal address type.* Depending on the address type, do one of the following:
 - In the Internal Address text box, enter the desired IP address. For information on the allowed address range, point to the question mark icon.
 - From the Internal Address drop-down list, select the name of the desired internal IP address resource.
- g. Click **Save**.

You can also edit network interface settings (click  **Edit**) and make the required modifications) or remove the ones that you do not need anymore (click  **Delete**).

11. Click **Restore**.



External network address types

For the external address of a network interface, the following options are available:

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

Internal network address types

For the internal address of a network interface, the following options are available:

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

Restoring individual files or folders

You can restore one or more individual files or folders if they have been deleted for some reason and are now missing on the instance.

You can restore individual files or folders to:

- The original location or a new location on the original instance
- A bucket that is set up in the protection set which includes the project of the original instance

Prerequisites

- The original instance had the Enable restore of individual files option selected at backup time.
- *Only for a restore to the original instance.* The following prerequisites are fulfilled:
 - The instance still exists and is running.
 - The target disk volume uses one of the supported file systems. See the *HYCU Data Protection as a Service for GCP Compatibility Matrix*.
 - In the Google Cloud Platform Console, there is a network firewall rule applied to the instances—either to the entire network or to individual instances through the use

of network (target) tags. For each target instance, the rule must allow ingress network traffic through a specific TCP port from the entire subnetwork that the instance belongs to. The port number depends on the guest operating system of the instance and connection server configuration:

- GNU/Linux: TCP port 22 (or a different port if configured for SSH communication)
- Microsoft Windows: TCP port 5986 (or a different port if configured for WinRM communication)

Optionally, you can make the network firewall rule more restrictive so that it allows network traffic only from legitimate sources. To do so, add the following source tag to the network firewall rule:

Source tag

```
hycu-<String>
```

In the template, `<String>` denotes any string that is allowed to be specified in firewall rule tags.

For instructions on how to configure and apply the network firewall rule, see a corresponding Google Cloud Platform webpage:

- [Using Firewall Rules | VPC | Google Cloud](#)
- [Configuring Network Tags | VPC | Google Cloud](#)


- *Only for a restore to the original instance running GNU/Linux:* The following prerequisites are fulfilled:
 - 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 the Google Compute Engine service.* An appropriate credential group is assigned to the original instance.
- *Only for a restore to the original instance running Microsoft Windows:* The following prerequisites are fulfilled:
 - On the original instance, WinRM is configured to use a TCP port (by default, 5986) for HTTPS or HTTP communication. The firewall is configured to enable inbound network traffic through this port.
 - An appropriate credential group is assigned to the original instance, and the supplied credentials belong to a user account with sufficient privileges. Credential group assignment is performed automatically by HYCU for GCP. For instructions on how to manually assign credential groups, see [“Enabling access to objects inside instances” on page 46](#).
- *Only for a restore to a bucket.* At least one bucket is set up in the protection set that

includes the project of the original instance. For information on how to add user-created buckets to this list, see [“Setting up user-created buckets” on page 45](#).

Considerations


- For a restore of individual files or folders, HYCU for GCP 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.
- *Only for a restore to the original instance running GNU/Linux.* Depending on whether the instance has a credential group assigned in HYCU for GCP, the following user accounts are used for the restore task:
 - No credential group is assigned: The sign-in user account that invokes the task.
 - A credential group is assigned: The user account that is specified in the credential group.


Accessing the Instances panel



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

Procedure

1. In the Instances panel, click the instance that contains the file or folder that you want to restore. By doing so, you open the Details section.



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


2. In the Details section that appears at the bottom of the screen, select the desired restore point, and then click  **Restore Files**.

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

3. In the Choose Disks dialog box, in the Disks drop-down list, deselect the disks that do not contain the file or folder that you want to restore. This shortens the time that HYCU for GCP need to load the disk catalogs.

4. In the Choose File and Folders dialog box, from the list of available files and folders, select the one 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 text box and then pressing **Enter**.

The File Restore Options dialog box appears.

5. In the File Restore Options dialog box, select whether you want to restore the file or

folder to the original instance or a bucket.

Original instance	<p>To restore the file or folder to the original instance, follow these steps:</p> <ol style="list-style-type: none"> a. Select Restore to Instance, and then click Next. b. In the Restore to Instance dialog box, select the location on the instance where you want to restore the file or folder, and provide the required information. <p>Make a selection:</p> <ul style="list-style-type: none"> • Original location <p>Select how HYCU for GCP should act when there is already a file with the same name at the original location. You can choose among the following options:</p> <ul style="list-style-type: none"> ◦ Overwrite original The file on the disk is overwritten with restored data. ◦ Rename original The original file is renamed before the data is restored. ◦ Rename restored The restored file gets a new name. <p>For the naming conventions for renamed original and restored files, see “Objects created by the service” on page 115.</p> • Alternate location <p>In the Path on the Original Disk dialog box, specify the path to an alternate location on the same instance in the following format:</p> <ul style="list-style-type: none"> ◦ GNU/Linux: <pre style="background-color: #f0f0f0; padding: 2px; margin: 5px 0;">/<Path>/<FolderName></pre> ◦ Microsoft Windows: <pre style="background-color: #f0f0f0; padding: 2px; margin: 5px 0;"><DriveLetter>:\<Path>\<FolderName></pre> <p>The restored file overwrites the file with the same name that might exist at the alternate location.</p> <ol style="list-style-type: none"> c. Depending on your preference for the restore of the file access control lists (ACLs), do one of the following: <ul style="list-style-type: none"> • Select the Restore ACL option. This will make HYCU for GCP restore original ACLs and apply them on the restored file. • Keep the Restore ACL option unselected. This will make HYCU for GCP apply inherited ACLs on the restored file—according to the file system ACL inheritance rules.
-------------------	---

Bucket	<p>To restore files or folders to a bucket, follow these steps:</p> <ol style="list-style-type: none">a. Select Restore to Bucket, and then click Next.b. In the Restore to Bucket dialog box, from the Target bucket drop-down list, select a bucket to which you want to restore data. <p>Restored objects are copied into the selected bucket. For information on their exact locations in the bucket, see “Objects created by the service” on page 115.</p>
--------	--

6. Click **Restore**.

Chapter 6

Performing common tasks

To ensure secure and reliable performance of HYCU for GCP, the solution provides various mechanisms to support your daily activities.

Goal	Link to instructions
Achieve any of the following: <ul style="list-style-type: none">• Get an at-a-glance overview of the data protection status in your environment.• Identify possible bottlenecks.• Inspect different areas of the environment protected by HYCU for GCP.	"Using the HYCU for GCP dashboard" on the next page.
Achieve any of the following: <ul style="list-style-type: none">• View bucket information.• Deactivate or activate a bucket.• Edit or remove a bucket.	"Managing buckets" on page 84.
Achieve either of the following: <ul style="list-style-type: none">• View backup policy information.• Edit or delete a backup policy.	"Managing backup policies" on page 87.
Obtain HYCU for GCP reports on different aspects of the data protection environment.	"Using HYCU for GCP reports" on page 80
View the backup status of instances.	"Viewing detailed instance information" on page 75.
Back up data manually.	"Performing manual backups" on page 82.
Mark a restore point as expired.	"Manually marking restore points as expired" on page 83.
Achieve either of the following: <ul style="list-style-type: none">• Track tasks that are running in your environment.	"Checking task statuses" on page 74.

Goal	Link to instructions
<ul style="list-style-type: none"> Get insight into the status of a specific task. 	
View all events that occurred in your data protection environment.	"Viewing events" on page 74.
Achieve either of the following: <ul style="list-style-type: none"> Narrow down the list of displayed elements in panels. Sort the list of displayed elements in panels. 	"Filtering and sorting data in panels" on page 70.


Using the HYCU for GCP dashboard

The HYCU for GCP web user interface includes an intuitive dashboard. This dashboard:

- Provides an at-a-glance overview of the data protection status in your environment, substantiated with relevant statistics.
- Enables you to monitor all data protection activities and quickly identify areas that require your attention.

You can use the HYCU for GCP dashboard as a starting point for your data protection activities. It enables you to easily access the area of interest by simply clicking the corresponding widget title.

Accessing the Dashboard panel

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

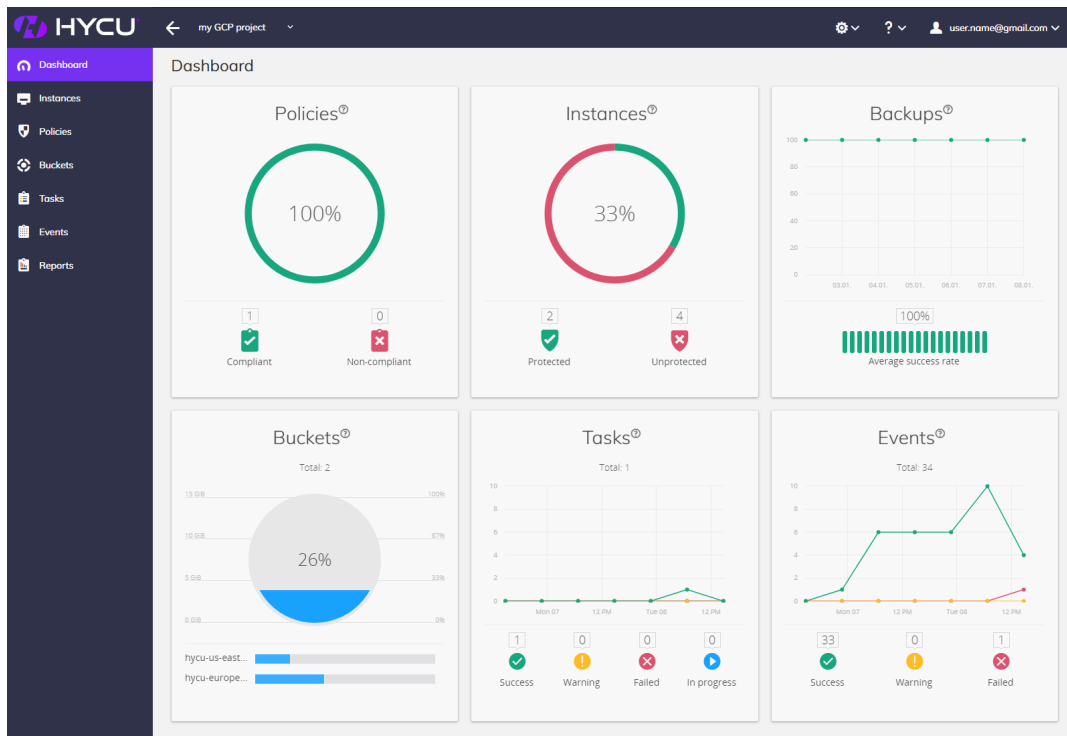



Figure 6–1: Dashboard in the web user interface of HYCU for GCP

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

Widget	Description
Policies	Shows the percentage of policies that are compliant, and the exact number of compliant and non-compliant policies. A backup policy is considered compliant if all instances that have this policy assigned are compliant. An instance is considered compliant when it is backed up so that the recovery point objective (RPO) in the assigned policy is met. For more information on policies, see “Defining your backup strategy” on page 33 .
Instances	Shows the percentage of protected instances in your environment, and the exact number of protected and unprotected instances. An instance is considered protected if it has a policy assigned and at least one backup image of the instance exists. Instances that have the exclude backup policy assigned are excluded from compliance and protection status determination. They are omitted from the figures depicted in this widget. For information about protecting instances, see “Backing up instances” on page 54 .
Backups	Shows the backup task success rate for the last seven days.
Buckets	Shows the number of existing buckets, their individual storage capacity utilization, and joint storage capacity utilization. For information about setting up buckets, see “Setting up buckets” on page 45 .
Tasks	Shows the number of tasks in the protected environment for the last 48 hours. The widget also shows how many tasks succeeded, failed, or are in progress or queued. For instructions on how to check the task status, see “Checking task statuses” on page 74 .
Events	Shows the number of events in the protected environment for the last 48 hours. The widget also shows the number of events for each event severity. For instructions on how to view event details, see “Viewing events” on page 74 .


 **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.

Filtering and sorting data in panels

HYCU for GCP 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 backup policies.

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 text box, 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 Instances panel” below](#)
- [“Filtering options in the Policies panel” on the next page](#)
- [“Filtering options in the Buckets panel” on page 72](#)
- [“Filtering options in the Tasks panel” on page 72](#)
- [“Filtering options in the Events panel” on page 73](#)

Filtering options in the Instances panel

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

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

Filtering option	Action
Project	From the drop-down list, select the Google Cloud Platform projects of interest.
Policy	From the drop-down list, select the backup policies that are

Filtering option	Action
	assigned to the instances.
Credential group	From the drop-down list, select the credential groups that are assigned to the instances.
Zone	From the drop-down list, select the Google Compute Engine instance zones.
Compliance	<p>Select one or more options to filter by the compliance status:</p> <ul style="list-style-type: none"> • Success: Instance is compliant. • Failure: Instance is not compliant. • Undefined: The exclude backup policy is assigned to the instance, or the instance does not have a backup policy assigned.
Protection	<p>Select one or more options to filter by the protection status:</p> <ul style="list-style-type: none"> • Yes: Instance is protected. • No: Instance is not protected. • Deleted: Instance no longer exists, but at least one of its backup images does. • Undefined: The exclude backup policy is assigned to the instance.
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 commands). • Failure: The instance could not be connected to. • Undefined: Connectivity to the instance has not been checked.

Filtering options in the Policies panel

You can enter a backup 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 instances that the backup policy is assigned to are compliant.

Filtering option	Action
	<ul style="list-style-type: none"> • Failure: Not all instances that the backup policy is assigned to are compliant. • Undefined: The backup policy is not assigned to any instance or this is the exclude backup policy.

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:

Filtering option	Action
Storage class	Select one or more options to filter by the Google Cloud Platform storage class: <ul style="list-style-type: none"> • Multi-regional • Regional • Standard • Coldline • Nearline • Archive
Health	Select one or more options to filter by the status of the bucket: <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
Project	From the drop-down list, select the Google Cloud Platform projects of interest.
Username	From the drop-down list, select items to filter the list to include only the tasks started by any of the selected user accounts. You can select Google Accounts and Google Cloud Platform service accounts.

Filtering option	Action
Status	Select one or more options to filter by the status of the task: <ul style="list-style-type: none"> • Done • Ready • Running • Failed • Done with errors • Aborted

Filtering options in the Events panel

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


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


Filtering option	Action
Project	From the drop-down list, select the Google Cloud Platform projects of interest.
Category	From the drop-down list, select items to filter the list to include only the selected event categories.
Username	From the drop-down list, select items to filter the list to include only the events resulting from the selected user accounts' actions. You can select Google Accounts and Google Cloud Platform service accounts.
Severity	Select one or more options to filter by the event severity: <ul style="list-style-type: none"> • Success • Warning • Failed

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.

Checking task statuses

You can use the Tasks panel to check the overall status of tasks.



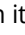
Accessing the Tasks panel

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


In the Tasks panel, you can do the following:

- Check the status of tasks that are currently running.
- Check the status of completed and stopped tasks.
- Check more details about a specific task.

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

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

- Generate a report about a specific task.

To generate the report, select a task, and then click  **View Report**. To copy the report to the clipboard, in the Task Report dialog box that opens, click **Copy to clipboard**.

- Abort a currently running task.

To abort a task, select it, and then click  **Abort Task**.


The following table shows the task information:

Task information	Description
Description	Summary of the task (for example, running a backup, performing a restore, restoring individual files or folders).
Status	Current status of a task (for example, Ready, a progress bar indicating the Running status, Done, Done with errors, Failed, or Aborted).
Started	The task's start date and time.
Finished	The task's finish date and time.

Viewing events

The Events panel enables you to view all events that occurred in your environment, to check details about the selected event, and to list events that match the specified filter.


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>Solution functional area to which the event belongs:</p> <ul style="list-style-type: none"> Targets: Events related to bucket management. Credentials: Events related to assignment of credentials to instances. Configuration: Events related to setting backup options of the instances. Archive: Events related to creating archives. Policies: Events related to backup policy management. Backup: Events that take place during backup, notifications about skipped backup tasks. Restore: Events that take place during restore. System: Events not related to any other category. Events of this type usually take place independently of your interaction with HYCU for GCP.
Timestamp	Event creation date and time.

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

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

Viewing detailed instance information

You can view detailed information about each instance in the Detail view section of the Instances panel. It includes the instance configuration and data protection summary and details of each of the restore points.





Viewing instance summary and restore point information


To view instance summary and restore point information, in the Instances panel, browse to and click the desired instance. 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.


To view details of a restore point, browse to the desired restore point in the list next to the Summary pane. 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.

The following information is available for the selected instance and the selected restore point:

Property name	Description
Summary	<p>Shows the following information about the selected instance:</p> <ul style="list-style-type: none"> • Virtual hardware properties and guest operating system properties • Current status of the instance • Instance location • Instance size and disk count • Time since last successful backup, backup success rate • Policy compliance success rate of the restore points • Policy assignor—the Google Cloud Platform identity that assigned the backup policy: a Google account (interactive assignment) or a service account (automatic assignment)
Restore point	<p>Shows the following information for the restore point:</p> <ul style="list-style-type: none"> • Creation date and time. • Available entities of the restore point, which determine your possibilities for a restore to that point in time: <ul style="list-style-type: none"> ◦ SNAP or S: Snapshot. Displayed if a snapshot of the instance exists. Snapshots allow faster completion of restore tasks. ◦ BCKP or B: Backup data in a bucket. Displayed if backup data of the instance is stored in a bucket. ◦ COPY or C: Copy of a backup image. Displayed if a copy of a backup image (snapshot or backup data in a bucket) exists in another bucket.





	<ul style="list-style-type: none"> ○ ARCH-D or D : Data archive—daily. Displayed if a daily data archive for the instance exists in a bucket. ○ ARCH-W or W : Data archive—weekly. Displayed if a weekly data archive for the instance exists in a bucket. ○ ARCH-M or M : Data archive—monthly. Displayed if a monthly data archive for the instance exists in a bucket. ○ ARCH-Y or Y : Data archive—yearly. Displayed if a yearly data archive for the instance exists in a bucket. ○ CTLG or C : Catalog. Displayed if a restore of individual files or folders is available. <p> Note A restore point may or may not include backup data of the entire instance. This depends on the disks included in the corresponding backup.</p> <p>Visual labels of the restore point entities may be specially marked to denote different restore point entity statuses. For more information, see “Statuses of restore point entities” on page 79.</p>
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 backup policy assigned to the instance was met). • The  icon: The backup is not compliant (the RPO setting in the backup policy assigned to the instance was not met). • 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 shown:</p> <ul style="list-style-type: none"> • Backup frequency in the backup policy that triggered the backup. • Elapsed time since the latest previous successful backup. • Expiration time for each available restore point entity (snapshot, backup data in a bucket, copy of the backup image, data archive).
Backup status	Shows the backup status of the restore point. For more information, see “Restore point backup statuses” on the next page .
Restore status	Shows a progress bar indicating the progress of an instance restore based on that restore point.




 **Note** By double-clicking a progress bar, you are directed to the Tasks panel where you can check details about the related task.


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

Restore point backup statuses

The backup status of a restore point indicates whether you can restore the instance or its individual disks to a state at the corresponding point in time. The following table lists possible status icons and the corresponding backup statuses of restore points as wholes.








Backup status	Restorability	Notes
 (Done)	✓	<p>May indicate one of the following:</p> <ul style="list-style-type: none"> All restore point entities were created successfully. Disk catalog creation failed. <p>You cannot restore individual files or folders.</p> <ul style="list-style-type: none"> Creation of either of the following failed or was aborted by the user: <ul style="list-style-type: none"> Copy of the backup image Data archive <p>The instance can still be fully restored from other restore point entities, for example, the backup image.</p>
 (Done with errors)	✓*	<p>Not all instance disks were backed up successfully.</p> <p>* The instance can be only partially restored. If backup of the boot disk failed, you may be unable to start the instance after the restore.</p>
 (Failed)	×	<p>Creation of a backup image (primary backup data—snapshot or backup data in a bucket) failed.</p>
 (Aborted)	×	None

Backup status	Restorability	Notes
 (Expired)	×	All entities of the restore point are marked as expired.
 (Inaccessible on GCP)	×	<i>Only if no instance snapshot is kept.</i> Access permissions on all Google Cloud Storage buckets where the backup data is stored are not sufficient.
 (Deleted from GCP)	×	<i>Only if no instance snapshot is kept.</i> Backup data is not available because buckets have been deleted from the Google Cloud Platform project.

 **Note** By pausing on a backup status icon, additional information about the restore point is shown: consistency indicator, backup duration, size of backup data in a bucket, names of the containing Google Cloud Storage buckets, and backup ID.

Statuses of restore point entities

Restore point entity labels may be visually marked to represent different statuses of individual restore point entities. The following table lists possible marks and the corresponding statuses of individual restore point entities (the visual label for backup data in a bucket is used as an example).

Status	Restorability
 (Done)	✓
 (Done with errors)	✓*
 (Failed)	×
 (Aborted)	×
 (Expired)	×
 (Inaccessible on GCP)	×
 (Deleted from GCP)	×


* The instance can be only partially restored. If backup of the boot disk failed, you may be unable to start the instance after the restore.

A particular status of restore point entities has the same meaning as the corresponding backup status of entire restore points. For details, see [“Restore point backup statuses” on the previous page](#).

Using HYCU for GCP reports

HYCU for GCP reports provide you with visual presentations of different metrics from your data protection environment. These visualizations allow you to better understand your data protection environment, identify potential bottlenecks, and efficiently eliminate them. You can easier analyze the resource consumption trends and therefore make optimum decisions when it comes to protecting your data.


Report data can be presented as a table or as a chart. The following report chart types are used to visualize the reports: a bar chart or an area chart.


 **Important** Data shown in a report reflects the state of your environment as it was not at the moment of report preview or generation but up to one hour earlier (at the moment of data collection).

After you get familiar with the reports as described in [“Getting started with reporting”](#) below, you can continue as follows:

- View reports. For instructions, see [“Viewing reports” on the next page](#).
- Generate reports. For instructions, see [“Generating reports” on the next page](#).
- Schedule reports. For instructions, see [“Scheduling reports” on page 82](#).

Accessing the Reports panel

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

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

Getting started with reporting

HYCU for GCP comes with a series of predefined reports.

Predefined reports



The predefined reports enable you to obtain information on the key aspects of your HYCU for GCP protection set such as the size of instance disks and the total size of instance backup data. Each report either comprises a specific time period or shows the state at a single point in time. The reports cannot be edited or deleted.


Name	Description
Backup data in buckets – per instance	Amount of backup data (including copies and archives) in buckets for each protected instance.
Backup data in buckets – per policy	Amount of backup data (including copies and archives) in buckets for each backup policy.

Name	Description
Backup data in buckets – per storage class	Amount of backup data (including copies and archives) in buckets for each storage class.
Instance compliance status	List of instances, their compliance statuses, assigned backup policies, and the corresponding pricing tiers.
Protected instance disk capacity – per policy	Amount of protected instance disk capacity for each backup policy.
Total backup data in buckets (trend)	Total amount of backup data (including archives) in buckets through time.
Total instance disk capacity (trend)	Total amount of instance disk capacity through time.

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.

Goal	Procedure
View a report on the current state of my data protection environment.	From the list of reports, select the desired report, and then click  Preview .
View a saved report version.	<ol style="list-style-type: none"> From the list of reports, select the desired report. In the Details section that appears at the bottom of the screen, select the desired report version, and then double-click it or click  View Report. <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 CSV format. To do so, click  **Download**, and then select **Download as CSV**.

Generating reports

When you generate a report, you actually 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.
2. In the Details section that appears at the bottom of the screen, click **+ Generate**. The Generate Report Version dialog box opens.
3. *Optional.* Enter a description for the report version.
4. Click **Generate**.

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

You can later do the following:

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

Scheduling reports

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

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.
2. In the Schedule date box, specify the date and the time of day when you want the report generation to begin.
3. From the Interval drop-down list, select how often you want the report versions to be generated (daily, weekly, or monthly).
4. Click **Save**.

💡 Tip Reports with automated report version generation 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**.

Performing manual backups

HYCU for GCP backs up your data automatically after you assign a backup policy to the selected instances. 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 backup policy other than the exclude backup policy is assigned to the instance.

Consideration

When the assigned backup policy uses a backup window, manual backups may prevent the scheduled backup for the same instance from starting within the defined time frame. If this happens, the instance becomes non-compliant with the policy settings until the next backup window or the next manual backup.

Procedure


1. In the Instances panel, select which instances you want to back up.
2. Click  **Backup** to invoke the backup of the selected instances.
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.


Manually marking restore points as expired

If there is a restore point that you do not want to use for a restore anymore, you can mark it as expired. You can do this also for restore points whose backup status is Failed or Aborted if you want to free storage space. Each restore point represents data that was backed up at a specified point in time. Restore points contain one or more restore point entities—snapshot, backup data in a bucket, copy of a backup image, data archive—that you can individually mark as expired. This means that your action can be one of the following:


- Expiring an entire restore point
 - Make sure that all entities of the restore point are marked for expiration.
- Leaving a restore point, but expiring one or more of its entities
 - Make sure that the entity you want to keep is not marked for expiration.

 **Important** Marking a restore point or its entities as expired cannot be undone.

Accessing the Instances panel

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


Procedure

1. In the Instances panel, select the corresponding instance.
2. In the Details section that appears at the bottom of the screen, select the restore point that you want to mark as expired.
3. Click  **Expire**. The Expire Backup dialog box appears.
4. If the backup status of the restore point is Failed or Aborted, go to step 5 of this

procedure.

The dialog box shows only the available entities and that you can mark as expired. The complete list is as follows:

- **Backup (Snapshot)**
- **Backup (Bucket)**
- **Copy**
- **Archive - daily**
- **Archive - weekly**
- **Archive - monthly**
- **Archive - yearly**

 **Note** HYCU for GCP preselects all available entities of the restore point.

Leave the default selection if you want to mark the entire restore point for expiration, or deselect the entities that you do not want to mark as expired.


5. Click **Yes** to confirm your choice and apply the changes.

When the entities or the entire restore point are marked as expired, the first next retention maintenance task in HYCU for GCP removes the corresponding data from the Google Compute Engine service (snapshots) or the Google Cloud Storage service (other restore point entities).

Managing buckets


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





Accessing the Buckets panel

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

Viewing bucket information


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

Property name	Description
Name	<p>Bucket name (globally unique).</p> <p>For information on how the automatic buckets are named, see “Objects created by the service” on page 115.</p> <p> Tip You can click the bucket name to open the bucket</p>

Property name	Description
	<p>details page of the Google Cloud Platform Console in your web browser.</p>
Location	Name of the Google Cloud Storage region in which the bucket resides.
Storage Class	Default object storage class of the bucket in the Google Cloud Storage service: Multi-regional, Regional, Standard, Nearline, Coldline, or Archive.
Status	<p>Status of the bucket:</p> <ul style="list-style-type: none"> • Active: You can use the bucket for backing up data, creating data archives, and restoring data. • Inactive: The bucket has been deactivated within HYCU for GCP. As long as it is not activated you can use it only for restoring data. • Inaccessible on GCP: Insufficient permissions are set on the bucket in the Google Cloud Storage service. HYCU for GCP cannot access the bucket. • Deleted from GCP: The bucket no longer exists in the Google Cloud Storage service. <p>For instructions on how to change the status of active or inactive buckets, see "Deactivating and activating buckets" on the next page.</p>
Size Limit	Maximum amount of the bucket storage space (expressed in MiB, GiB, or TiB) that is allowed to be used by backup data created by HYCU for GCP. The amount represents a soft limit, therefore actual usage may exceed it.
Health	<p>Health status of the bucket:</p> <ul style="list-style-type: none"> • The  icon: Indicates one of the following: <ul style="list-style-type: none"> ◦ The bucket health has not been determined yet. ◦ The bucket is inactive. • The  icon: The bucket is in a healthy state. Utilization of storage space for backup data in the bucket is less than 90 percent of the configured size limit. • The  icon: Utilization of storage space for backup data in the bucket is over 90 percent and under 100 percent of the configured size limit. • The  icon: Indicates one of the following: <ul style="list-style-type: none"> ◦ Bucket storage space occupied by backup data exceeds the configured size limit.

Property name	Description
	<ul style="list-style-type: none"> The bucket is not accessible due to an I/O error, insufficient permissions, or some other reason.
Utilization	Ratio (expressed in percentage) between the bucket storage space occupied by backup data and the configured size limit.
Automatic	Indicator of whether the bucket was created automatically by HYCU for GCP (✓) or not (✗).


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

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

Editing buckets


Editing a bucket means changing the value of the bucket's Size Limit property. This property represents a soft limit.

Procedure

1. In the Buckets panel, select the bucket that you want to edit, and then click  **Edit**. The Edit Bucket dialog box appears.
2. Edit the value in the Size text box or select another unit of measurement from the drop-down list as required.
3. Click **Save**.

Deactivating and activating buckets

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

 **Note** You cannot deactivate buckets that were created automatically by HYCU for GCP.



Prerequisites

- *Only for bucket deactivation.* The bucket is not specified in a Bucket option of any backup policy.
- *Only for bucket deactivation.* The bucket is not specified in the Bucket option of any data archive.

Consideration


After deactivating a bucket, the bucket cannot be selected for the Bucket option of a backup policy until it is activated again.

Procedure

1. In the Buckets panel, select the bucket that you want to deactivate or activate.
2. Change the status of the selected bucket: click  **Deactivate** or  **Activate**.
3. *Only for deactivation.* Click **Yes** to confirm that you want to deactivate the selected bucket.

Removing buckets

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

 **Note** You cannot remove buckets that were created automatically by HYCU for GCP.


Prerequisites

- The bucket contains no backup data.
- The bucket is not specified in a Bucket option of any backup policy.
- The bucket is not specified in the Bucket option of any data archive.

Consideration

After removing a bucket, no backup operations that include this bucket are possible anymore.


Procedure

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

Managing backup policies




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

Accessing the Policies panel



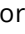
To access the Policies panel, in the navigation pane, click  **Policies**. Alternatively, in the Dashboard panel, click the **Policies** widget title.

Viewing backup policy information

You can view information about each backup policy in the list of backup policies in the Policies panel. This allows you to have an overview of the general status of the backup policies. The following information is available for each backup policy:

Property name	Description
Name	Backup policy name.
Compliance	Compliance status of the backup policy: <ul style="list-style-type: none"> • The  icon: The backup policy is compliant. • The  icon: The backup policy is non-compliant. • The  icon: Backup policy compliance is undefined. The policy is not assigned to any instance, or this is the exclude policy.
Instance Count	Number of the instances that have the backup policy assigned to them.
Description	Description of the backup policy.

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


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

Creating backup policies

See [“Creating custom backup policies” on page 35](#).


Editing backup policies

Procedure

1. In the Policies panel, select the backup policy that you want to edit, and then click  **Edit**. The Edit Policy dialog box appears.
2. Edit the selected backup policy as required. For detailed information about backup policy properties, see [“Creating custom backup policies” on page 35](#).
3. Click **Save**.

Deleting backup policies


Procedure

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


Viewing subscription information

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

Accessing the Subscription Information dialog box

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

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

Subscription — Billing Account	Name of the billing account that is used for the subscription.  Note With multiple HYCU for GCP subscriptions that you can use, choose which one you want to check the information about by selecting its corresponding billing account from the drop-down list.
Subscriber	
Name	Information about the person who subscribed to HYCU for GCP.
Surname	
Company	
Subscription details	
Subscription ID	An identification that is automatically assigned to the subscription by Google.
Subscription plan	The plan that your HYCU for GCP 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 recovery pricing” on page 17 .
Subscribed on	The date of subscribing to HYCU for GCP.

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

Chapter 7

Administering

While using HYCU for GCP, you can perform different tasks to administer and customize the solution for your data protection needs.

Task	Link to instructions
Get familiar with user permissions in HYCU for GCP.	"Managing users" below
Set a different billing account viewer for a HYCU for GCP subscription.	"Setting the billing account viewer" on page 31
Manage HYCU for GCP protection sets.	"Managing protection sets" on the next page
Configure protection set service accounts.	"Configuring protection set service accounts" on page 95
Be notified about different kinds of events in HYCU for GCP.	"Configuring event notifications" on page 98
Hide instances from HYCU for GCP.	"Excluding instances from synchronization" on page 101
Stop protecting individual projects.	"Stopping protection for individual projects" on page 102

If for whatever reason you decide that you no longer want to use HYCU for GCP or you do not need it anymore, you can easily stop using it. For information on properly doing so, see ["Ceasing to use HYCU for GCP" on page 110](#).

Managing users

This section provides information on managing HYCU for GCP users. HYCU for GCP does not provide means of direct user management. The service indirectly determines which user accounts are allowed to take specific actions based on:

- The billing account selected when you subscribe to the solution.
- Sets of permissions that are granted the corresponding Google Accounts in the Google Cloud Platform service suite.

The following table lists different goals that you want to achieve by using a specific sign-in user account and their corresponding preconditions.


Goal	Precondition
Sign in to the web user interface	The sign-in user account can access at least one of the scoped Google Cloud Platform projects.
See your Google Cloud Platform projects	The projects are linked to the billing account selected for the HYCU for GCP subscription.
Protect your Google Cloud Platform projects	The sign-in user account has the required roles granted in the following services of the Google Cloud Platform service suite: <ul style="list-style-type: none"> • Google Compute Engine • Google Cloud Storage

For more information on the required preconditions, see [“Extent of data protection” on page 13](#) and [“Subscribing to and signing in to HYCU for GCP” on page 23](#).

Managing protection sets

This section provides information on managing HYCU for GCP protection sets. By default, a single predefined protection set (default-protection-set) includes all projects that are linked to a billing account of a HYCU for GCP subscription.


You can adjust the default setup to better suit your needs by configuring additional protection sets and distributing your projects among them. As protection sets can consist of a single project or multiple projects, you can also address requirements of individual projects.

 **Important** For all protection set management operations, the web user interface displays all projects that are linked to the billing account of the currently selected HYCU for GCP subscription, regardless of the access permissions of the sign-in user account.

See the following sections for the corresponding instructions:

- [“Configuring protection sets” on the next page](#)
- [“Editing protection sets” on the next page](#)
- [“Deleting protection sets” on page 94](#)
- [“Excluding projects from any protection set” on page 95](#)

Accessing the Protection Sets dialog box

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

select **Protection Sets**.

Configuring protection sets

You may want to have different data protection setup for different groups of Google Cloud Platform projects, or use different service accounts to run data protection tasks for them. In both cases, you must configure additional protection sets.

Prerequisite

Your sign-in user account has access to at least one of the Google Cloud Platform projects that you want to join in a new protection set.

Consideration

When a project is moved to a different protection set, HYCU for GCP automatically unassigns backup policies from the instances in the project.

Procedure

1. In the Protection Sets dialog box, from the **Subscription — Billing Account** drop-down list, select the billing account of the desired HYCU for GCP subscription.
2. Click **+ New**.
3. In the Name text box, enter a name for the new protection set.
4. *Optional.* In the Description text box, enter a protection set description.
5. In the Project column, locate projects that you want to include in the protection set, and select them.
6. Click **Save**.

Editing protection sets

You can at any time rename protection sets or change the set of included Google Cloud Platform projects.

Limitation

You can edit a protection set only if it includes at least one project to which your sign-in user account has access.


Considerations

- When a project is moved to a different protection set, HYCU for GCP automatically unassigns backup policies from the instances in the project.
- HYCU for GCP prevents you from configuring a protection set that consists solely of projects to which your sign-in user account has no access. At least one accessible project must be included.

Recommendation

With multiple configured protection sets, the default predefined protection set should always have an easily distinguishable name.

Procedure

1. In the Protection Sets dialog box, from the **Subscription — Billing Account** drop-down list, select the billing account of the desired HYCU for GCP subscription.
2. In the Name column, select the protection set, and then click  **Edit**.
3. *Only if you want to rename the protection set or change its description.* In the Name text box, enter a new name for the protection set. In the Description text box, enter a new description.
4. *Only if you want to change the protection set configuration.* Do the following:
 - a. To add a project into the protection set, in the Project column, locate the project name, and then click the icon next to it.
 - b. To remove a project from the protection set, in the Project column, locate the project name, and then click the icon next to it.
 - c. Repeat steps a and b as appropriate.
5. Click **Save**.

Deleting protection sets

You can at any time delete protection sets that you no longer need. Keep in mind that the default predefined protection set cannot be deleted.


Prerequisites

- The protection set that you want to delete is empty with no included projects.
- The current web user interface scope is set to a protection set other than the protection set that you want to delete.

Consideration

When a protection set is deleted, HYCU for GCP data protection configuration of the protection set is permanently lost. Data protection configuration comprises backup policies, credential groups, data archive configuration, backup window specifications, service account assignments, event notification configuration, configured reports, and generated report versions.

Procedure

1. In the Protection Sets dialog box, from the **Subscription — Billing Account** drop-down list, select the billing account of the desired HYCU for GCP subscription.
2. In the Name column, select the name of the protection set, and then click  **Delete**.
3. In the Delete Protection Set dialog box, click **Yes** to confirm the deletion.

Excluding projects from any protection set

You may want to prevent specific projects from being protected with HYCU for GCP. To do so, you must remove them from any configured protection set.

Limitation

You can edit a protection set only if it includes at least one project to which your sign-in user account has access.

Consideration

HYCU for GCP prevents you from configuring a protection set that consists solely of projects to which your sign-in user account has no access. At least one accessible project must be included.

Procedure

1. In the Protection Sets dialog box, from the **Subscription — Billing Account** drop-down list, select the billing account of the desired HYCU for GCP subscription.
2. *Only if the desired project belongs to a protection set other than the default predefined protection set.* Do the following:
 - a. In the Name column, click the **>** icon next to the protection set with the desired project. The list of included projects appears.
 - b. Select the desired project, and then click **🗑 Remove**.
 - c. In the Remove Project dialog box, click **Yes** to confirm the removal. The removed project is added to the default predefined protection set.
3. Do the following:
 - a. In the Protection Sets dialog box, in the Name column, click the **>** icon next to the default predefined protection set. The list of the included projects appears.
 - b. Select the desired project, and then click **🗑 Remove**.
 - c. In the Remove Project dialog box, click **Yes** to confirm the removal. The removed project is no longer included in any protection set.

When a project is excluded from any protection set, HYCU for GCP no longer retrieves the project information from the Google Cloud Platform service suite.


Configuring protection set service accounts

This section provides information on how to configure protection set service accounts in HYCU for GCP. Protection set service account is a Google Cloud Platform service account that HYCU for GCP uses to invoke specific operations within a protection set. When configured, a protection set service account becomes a property of the protection set, and starts being used instead of the sign-in user account to:

- Restore an instance
- Clone an instance within the original project or to a different project of the same protection set

With configured protection set service accounts, also the following HYCU for GCP features become available:

- Policy assignment from the Google Cloud Platform service suite
- Use of the HYCU for GCP application programming interface (API)

 **Note** A protection set service account is used for identification of automated requests to HYCU for GCP within Google Cloud Platform projects of a protection set. Such requests must be authenticated similarly to those that you invoke interactively through the HYCU for GCP web user interface.

For more information about Google Cloud Platform service accounts, see the [Understanding service accounts | Cloud Identity and Access Management Documentation | Google Cloud](#) webpage.

Prerequisites

- The following application programming interfaces are enabled on the Google Cloud Platform project on which the service account was created:
 - Cloud Resource Manager API
 - Compute Engine API
 - Cloud Storage

For instructions on how to enable them, see the [Getting Started | Cloud APIs | Google Cloud](#) webpage.

- The service account is granted the following independent roles in the Google Cloud Platform service suite:
 - Compute Admin (`roles/compute.admin`), Storage Admin (`roles/storage.admin`), and Service Account User (`roles/iam.serviceAccountUser`) on all projects of the protection set that contain your protected instances.
 - Service Account Token Creator (`roles/iam.serviceAccountTokenCreator`) on the service account itself.
 - *Only for cloning into non-original projects by using the HYCU for GCP API.* Compute Admin (`roles/compute.admin`) and Service Account User (`roles/iam.serviceAccountUser`) on the projects where you plan to clone your instances to.

A service account becomes effective only after it is imported into HYCU for GCP and assigned to a protection set. See the following sections for the corresponding instructions:

- [“Importing service accounts” on the next page](#)
- [“Assigning imported service accounts to protection sets” on page 98](#)

Importing service accounts

Importing a service account makes it available to HYCU for GCP for assignment to a protection set.


Prerequisites

- The service account is configured in the Google Cloud Platform service suite.
- You have access to a valid JSON file that stores the service account information, including its private key.


Consideration

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

Accessing the Service Accounts dialog box


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

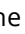
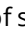

Procedure


1. In the Service Accounts dialog box, click  **Import**.
2. Click **Browse**. In the Choose File to Upload dialog box, browse for and then select the JSON file with the service account information. Click **Open**.
3. Review the attributes of the service account, and then click **Upload**.

As an indication of a successful import, the service account's email address appears in the list of service accounts.

4. Click **Close**.

 **Note** The email address of a successfully imported service account may be missing from the service account list. This happens when the service account belongs to a Google Cloud Platform project which is not linked to the billing account of the currently selected subscription. In this case, the service account is stored for use at a later time.

To delete an imported service account, in the Service Accounts dialog box, select its entry from the list of service accounts. If needed, click  or  to move between the pages. Click  to delete the selected service account from HYCU for GCP. Deletion of a service account that is used as the billing account viewer of a HYCU for GCP subscription is prevented.

 **Caution** Before deleting a service account, make sure that it is not assigned to any protection set in any HYCU for GCP subscription. Deletion unassigns the service account.


Assigning imported service accounts to protection sets

Assigning a service account to a protection set allows using it for restoring an instance, cloning an instance within the original Google Cloud Platform project or to a different project of the same protection set, and identification of automated requests to HYCU for GCP within projects of the protection set. After such assignment, also additional HYCU for GCP features are enabled for these projects.




Prerequisites


- The service account is imported into HYCU for GCP for the currently selected HYCU for GCP subscription.
- On all projects of the protection set to which you want to assign the service account, the service account is granted any Google Cloud Platform role that includes the `resourceManager.projects.get` permission. Most of the predefined roles, for example, the Compute Admin (`roles/compute.admin`) and Storage Admin (`roles/storage.admin`) roles, include this permission.
- The protection set to which you want to assign the service account is selected in the HYCU for GCP web user interface.


Accessing the Service Accounts dialog box

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

Procedure

1. In the service account list, select the service account that you want to assign. If needed, click  or  to move between the pages.
2. Click  **Assign** to assign the selected service account to the currently selected protection set.

As an indication of a successful assignment, the service account is marked with the  icon in the list of service accounts.

To unassign an imported service account from the selected protection set, in the Service Accounts dialog box, select its entry from the list of service accounts, and then click  **Unassign**.

Configuring event notifications

Notifications in HYCU for GCP are a convenient way of informing about events that occur in your data protection environment. They help you automate monitoring and make your supervision of the data protection activities more efficient. Properly configured notifications enable you to always react in time in situations that require doing so. HYCU for GCP supports two notification transmission methods which you can use independently:

- Email messages
- Webhooks


You configure notifications through a series of notification rules. Each rule defines who should be informed about what kind of events. For this purpose, a rule always includes at least one recipient, one or more categories (functional areas of HYCU for GCP), and one or more statuses (event severity levels) for which the notification should be sent.

For each email-based notification rule, you must specify:

- An email subject line
- Event categories of interest
- Event statuses of interest
- One or more recipients

For each webhook-based notification rule, you must specify:

- A name for the rule
- Event categories of interest
- Event statuses of interest
- Receiving endpoint URL
- *Only if the receiving endpoint requires sender's identification.* Secret token (webhook verification signature)

 **Important** Configured notification rules are effective within the selected HYCU for GCP protection set. Each protection set therefore requires its own configuration of event notifications.

See the following sections for the corresponding instructions:

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

Configuring email-based notifications

This section lists the steps that you must follow to configure email-based notifications.

Accessing the Notifications dialog box


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

Procedure

1. In the Notifications dialog box, the Email messages tab is preselected. Click **+ New** to create a new notification rule.
2. In the Subject text box, enter text for the email subject line.
3. From the Category drop-down list, select one or more categories. To include all

categories, click **Select All**. For description of categories, see [“Viewing events” on page 74](#).

4. From the Status drop-down list, select one or more statuses. To include all statuses, click **Select All**. For description of statuses, see [“Viewing events” on page 74](#).
5. In the Email address text box, enter one or more recipients (valid email addresses), separated by space characters.
6. Review the supplied data and click **Save**.

 **Note** An email notification rule becomes effective immediately after its configuration is saved.



7. Click **Close**.

To edit or delete an already configured notification rule, select its entry from the list of email notifications, and then click  **Edit** or  **Delete** as appropriate.

Configuring webhook-based notifications

This section lists the steps that you must follow to configure webhook-based notifications.

Accessing the Notifications dialog box


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

Procedure

1. In the Notifications dialog box, click the **Webhooks** tab.
2. Click **+ New** to create a new notification rule.
3. In the Name text box, enter a name for the rule.
4. From the Category drop-down list, select one or more categories. To include all categories, click **Select All**. For description of categories, see [“Viewing events” on page 74](#).
5. From the Status drop-down list, select one or more statuses. To include all statuses, click **Select All**. For description of statuses, see [“Viewing events” on page 74](#).
6. In the Post URL text box, enter a valid URL of the endpoint the callbacks should be sent to. The URL should match one of the following formats, depending on the actual endpoint address:

```
https://<Host>
https://<Host>/<Path>
```

7. *Only if the receiving endpoint requires sender's identification.* In the Secret text box, enter a valid secret token for authentication. If the receiving endpoint does not perform authentication, the token is ignored.
8. Review the supplied data and then click **Save**.

 **Note** A webhook notification rule becomes effective immediately after its configuration is saved.

9. Click **Close**.

To edit or delete an already configured notification rule, select its entry from the list of webhook notifications, and then click  **Edit** or  **Delete** as appropriate.

Excluding instances from synchronization



This section provides information on how to make selected instances invisible to HYCU for GCP. The needs of your environment may require that some instances are not protected by HYCU for GCP. For example, your Google Cloud Platform projects may include managed instance groups and employ an autoscaler. To leave some instance unprotected, you can exclude them from synchronization so that they are not visible to HYCU for GCP. The invisible instances cannot be assigned backup policies in any way.

Procedure

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

Key	Value
<code>hycu-instance-sync</code>	<code>false</code>


Custom metadata tags can be added from the Google Cloud Platform Console, the `gcloud` command line, or by using the Google Cloud Platform API. For instructions, see the [Storing and Retrieving Instance Metadata | Compute Engine Documentation | Google Cloud](#) webpage.

3. Repeat step 2 for each additional instance that you want to make invisible to HYCU for GCP.
4. Sign in to the HYCU for GCP web user interface.
5. Select the protection set that includes the same Google Cloud Platform project as you selected in step 1 of the procedure. For instructions on selecting protection sets in HYCU for GCP, see [“Selecting HYCU for GCP protection sets” on page 31](#).
6. In the navigation pane, click  **Instances**.
7. Click  **Synchronize** or wait until the next instance synchronization cycle.

In the Instances panel, the names of the instances that you excluded from synchronization are not present.

Stopping protection for individual projects

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

 **Note** If you want to stop using HYCU for GCP completely, see [“Ceasing to use HYCU for GCP” on page 110](#).

Procedure

1. In HYCU for GCP, unassign backup policies from all protected instances in the project. For instructions, see [“Unassigning backup policies” on page 111](#).
2. In HYCU for GCP, manually mark all restore points of all instances in the project as expired. For instructions, see [“Marking all restore points in a protection set as expired” on page 111](#).
3. Exclude the project from any protection set. For instructions, see [“Excluding projects from any protection set” on page 95](#).

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

Chapter 8

Troubleshooting

If you encounter problems while using HYCU for GCP, you can often solve them yourself. This chapter contains information that may help you in such cases. To get assistance from HYCU Customer Support straight away, see [“Getting assistance” on page 107](#).

We recommend that you use a troubleshooting flow depicted in the following figure.

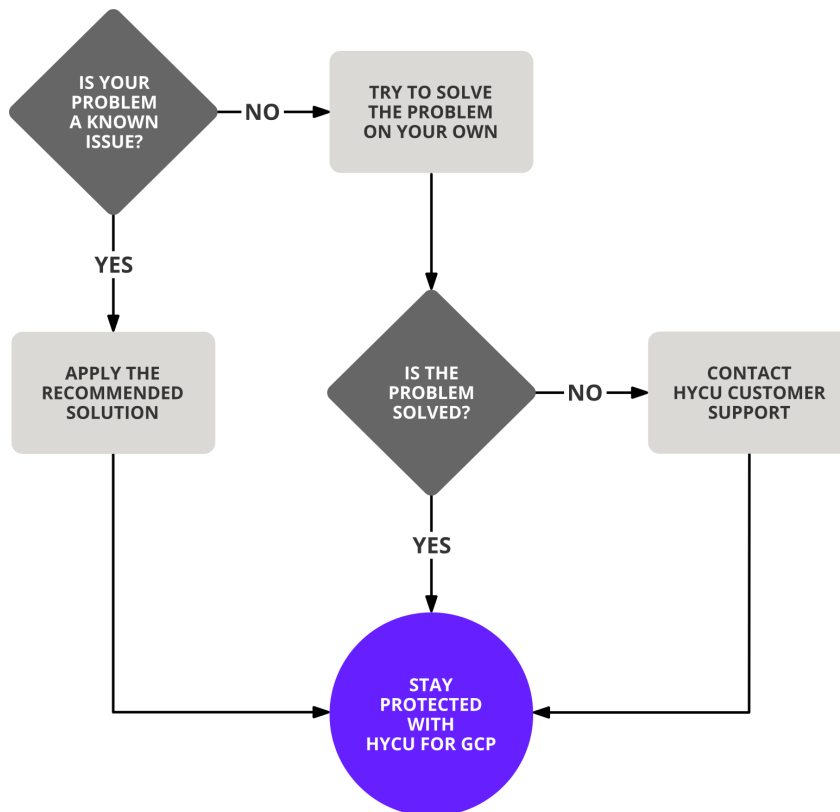


Figure 8-1: Overview of the troubleshooting process

General troubleshooting guidelines

When investigating an issue, first verify that:

- All subscription and usage prerequisites are fulfilled, and you performed configuration according to the provided instructions.

- You are not running into a known service limitation.
- Your issue is not related to third-party services (Google Cloud Platform). Otherwise, contact the respective service provider for assistance.
- The affected Google Cloud Platform instances are not running out of memory or storage space.

Problems and solutions

This section lists symptoms of common problems that you may encounter while using HYCU for GCP, together with proposed actions – resolution steps.

Missing Google Cloud Platform projects

Problem

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

Cause

The missing projects are not linked to any Google Cloud Platform billing account that was selected in the process of subscribing to the service.

Solution

Consult your organization's data protection administrator to choose between the following resolutions:

- In the Google Cloud Platform service suite, link the missing projects to a billing account that was selected for the HYCU for GCP subscription. For instructions, see the [Modify a Project's Billing Settings | Cloud Billing Documentation | Google Cloud](#) webpage.
- Subscribe to HYCU for GCP again. In the process, choose the billing account which your missing projects are linked to.

Inability to set up user-created buckets

Problem

In the web user interface, when you try to add a user-created Google Cloud Storage bucket, HYCU for GCP reports that the bucket is inaccessible.

Solution

In the Google Cloud Storage service, grant your Google Account the Storage Admin (`roles/storage.admin`) role on the Google Cloud Platform project of the problematic bucket.

For information on the required roles for general use of the service, see the prerequisite list in the section [“Signing in to HYCU for GCP” on page 25](#).

Backup policy assignment failures

Problem

After adding the `hycu-policy` custom metadata tag to an instance in the Google Compute Engine service, no backup policy is assigned to the instance in HYCU for GCP.

Cause

The symptom may indicate one of the following:

- The instance belongs to a project that is not included in any protection set.
- No service account is assigned to the protection set of the affected project in HYCU for GCP.
- The backup policy that is specified for the metadata tag value does not exist.

Solution

Find the corresponding entry in the event log to identify the root cause of the problem:

1. In the HYCU for GCP 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 failures

Problem

Whenever a backup task for any instance in a specific Google Cloud Platform project is started, the snapshot creation task fails and reports an error.

Solution

In the Google Compute Engine service, grant your Google Account the Compute Admin (`roles/compute.admin`) role on the problematic Google Cloud Platform project.

For information on the required roles for general use of the service, see the prerequisite list in the section [“Signing in to HYCU for GCP” on page 25](#).

Task progress indicator stuck at 0% forever

Problem

You experience one of the following symptoms:

- When you invoke a backup task, its child task for creating disk catalog never makes any progress.
- After you invoke a backup task or a restore task, the task gets started but it never makes any progress.

Solution

Check if the Google Cloud Platform 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 Platform Console.

Restore of individual files or folders ending with errors or failing

Problem

When a restore of individual files or folders completes, the status of the corresponding task is set to Done with errors or Failed. Closer inspection reveals that some or all of your selected objects have not been restored.

Cause

The original volume no longer exists, or the credential group that is assigned to the original instance in HYCU for GCP includes a user account with insufficient privileges.

Solution

Restore your files or folders to an alternate location on the original instance or to an available bucket, or update configuration of the credential group that is assigned to the original instance in HYCU for GCP.

Restore of individual files or folders failing

Problem

A restore of individual files or folders to the original instance fails because of unsuccessful mounting of the original disk.

Cause

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

Solution

Assign a credential group to the instance or make the necessary adjustments to the credential group configuration. For instructions, see [“Configuring and assigning credential groups manually” on page 47](#).

Inability to change the protection set or to sign in

Problem

Although you have access to Google Cloud Platform projects that are included in multiple protection sets in HYCU for GCP, only the currently selected protection set is available in the Protection Set Picker dialog box.

After your web user interface session ends, you are unable to sign in again.

Cause

Most probably the Google Cloud Platform identity (a Google Account or a service account) that is used as the billing account viewer lost the required role on the billing account of the HYCU for GCP subscription.

Solution

Contact HYCU Customer Support. For information on how to prevent the problem from reoccurring, see [“Setting the billing account viewer” on page 31](#).

Instance backup option reconfiguration failure

Problem

With an instance running Microsoft Windows, after you enable restore of individual files or folders in the Instance Configuration dialog box, automatic assignment of a credential group to the instance fails. HYCU for GCP is therefore unable to update 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 objects inside instances” on page 46](#).

Getting assistance

Depending on the required type of assistance, do the following:

- If you need assistance with service evaluation, contact HYCU Customer Support. See [“Customer support” on the next page](#).
- If you are already past your free trial period, you have a valid HYCU for GCP subscription, and you:
 - Require information about HYCU for GCP pricing, see [“Service pricing” on page 15](#).
 - Have an operational issue with HYCU for GCP, see [“Customer support” on the next page](#).

Customer support

If you have an issue with the service, collect the following information before contacting HYCU Customer Support:

- Sequence of actions leading to the problem
- Symptoms that you noticed and the expected behavior
- Messages that you received: description, date, and time
- Information about the first occurrence (and possible recurrence) of the problem

The listed pieces of information are required by HYCU Customer Support so that a support engineer can efficiently investigate the issue from the very beginning. When you have the information ready, do one of the following:

- *Preferred.* On the [HYCU Customer Support](#) webpage, submit your request (support case) with the information included.
- Send an email with the included information to support@hycu.com.

HYCU Customer Support will contact you shortly.

Getting additional information and latest updates

For additional information about HYCU for GCP, visit the [Backup & Recovery for Google Cloud Platform | HYCU](#) webpage.


For the most up-to-date documentation, go to the [HYCU Data Protection as a Service for GCP – HYCU Customer Support](#) webpage.

Before contacting HYCU Customer Support

If you cannot solve your issue, report it. Before contacting HYCU Customer Support, make sure that you:

- Perform the general checks. For details, see [“General troubleshooting guidelines” on page 103](#).
- Verify that your problem is not documented. For more information, see [“Problems and solutions” on page 104](#).
- Collect relevant information that might be required to send to HYCU Customer Support. For details, see [“Customer Support” on page 121](#) and [“Getting assistance” on the previous page](#).

HYCU Customer Support will provide you with further instructions.


 **Note** HYCU Customer Support is not qualified to solve issues with third-party services.

For information on how to reach HYCU Customer Support, see [“HYCU Customer Support and information” on page 121](#).

Chapter 9

Ceasing to use HYCU for GCP

At any time, you can decide to cease using HYCU for GCP to protect instances in your projects in the Google Cloud Platform service suite.

 **Note** If you wish to only unprotect one or more Google Cloud Platform projects and continue with service use, see [“Stopping protection for individual projects” on page 102](#).

When you decide to stop using HYCU for GCP completely, you must complete a proper process to eliminate unnecessary costs, keep tight control over access to your Google Account, and cancel HYCU for GCP subscription. The process helps you achieve three goals:

1. Stop being charged for using HYCU for GCP when you no longer require it. For a guidance on achieving it, see [“Stopping service charges” below](#).
2. Prevent further HYCU for GCP access to your Google Account. For a guidance on achieving it, see [“Preventing account access” on page 113](#).
3. Unsubscribe from HYCU for GCP. For a guidance on achieving it, see [“Canceling HYCU for GCP subscription” on page 114](#).

Stopping service charges

Protection of your Google Cloud Platform data with HYCU for GCP is billed by Google through one or more billing accounts that your projects are linked to. To avoid unnecessary charges for the backup and recovery service, perform the following actions in the suggested order. Each procedure affects a different share of the total data protection cost.



Step	Data protection cost share	Action
1.	Backup and recovery service cost	Unassign backup policies from all protected instances in all projects. For instructions, see “Unassigning backup policies” on the next page .
2.	Backup data storage cost ¹	Manually mark all restore points of all instances in all projects as expired. For

		instructions, see “Marking all restore points in a protection set as expired” below.
3.	Cost of storing backup data in buckets	Remove all backup data created by HYCU for GCP from the Google Cloud Storage buckets. For instructions, see “Removing backup data from buckets” on the next page.
4.	Snapshot storage cost	Remove all snapshots created by HYCU for GCP from the Google Compute Engine service. For instructions, see “Removing snapshots” on page 113.

¹ Includes the cost of storing backup data in buckets and the snapshot storage cost.


Unassigning backup policies

Procedure

1. Sign in to the HYCU for GCP web user interface.
2. Select the desired protection set. For instructions, see [“Selecting a different protection set”](#) on page 32.
3. In the web user interface, in the navigation pane, click  **Instances**.
4. Select every instance with an assigned backup policy, and then click  **Policies**.
5. Click **Unassign**.
6. Click **Yes** to confirm that you want to unassign the policies from the selected instances.

Marking all restore points in a protection set as expired

Procedure

1. Sign in to the HYCU for GCP web user interface.
2. Select the desired protection set. For instructions, see [“Selecting HYCU for GCP protection sets”](#) on page 31.
3. In the web user interface, in the navigation pane, click  **Instances**.
4. Mark all (and entire) restore points of an instance as expired. For instructions, see [“Manually marking restore points as expired”](#) on page 83.
5. Repeat step 4 for each additional instance and for each project of the protection set.

Removing backup data from buckets

Removal of backup data of all instances in your protected Google Cloud Platform projects includes:

- *Only if automatic buckets contain only backup data created by HYCU for GCP.* Deletion of automatic buckets
- *Only if automatic buckets contain also data other than the backup data created by HYCU for GCP.* Removal of backup data from automatic buckets
- Removal of backup data from user-created buckets


Before you can delete automatic buckets or remove backup data from them, identify the buckets in a Google Cloud Platform project. For the bucket naming convention, see [“Objects created by the service” on page 115](#).

Accessing the Google Cloud Platform Console

To access the Google Cloud Platform Console, open a web browser, go to the [Google Cloud including GCP & G Suite — Try Free | Google Cloud](#) webpage, and click **Sign in**. Then sign in with your Google Account.


Deleting automatic buckets

Procedure

1. In the toolbar of the Google Cloud Platform Console, click . The Select a project dialog box appears.
2. In the Select a project dialog box, click **ALL** and then click a project name in the Name column.
3. Follow instructions on the [Deleting Buckets | Cloud Storage Documentation | Google Cloud](#) webpage.

Removing backup data from automatic or user-created buckets

Procedure

1. In the Google Cloud Platform Console, open the Google Cloud Storage browser.
2. In the toolbar, click . The Select a project dialog box appears.
3. In the Select a project dialog box, click **ALL**, and then click a project name in the Name column.
4. In the Name column, click the name of your bucket, and then click **hycu/**.
5. Select the checkbox next to the **backups/** folder name, and then click **Delete**.
6. In the overlay window, confirm you want to delete the folder and its contents by clicking **Delete**.

Removing snapshots


You can remove snapshots created by HYCU for GCP from either the Google Cloud Platform Console or the `gcloud` command line. To do so, first identify the snapshots that are created by HYCU for GCP in a Google Cloud Platform project. For the snapshot naming convention, see [“Objects created by the service” on page 115](#).

Removing snapshots from the Google Cloud Platform Console

Accessing the Google Cloud Platform Console

To access the Google Cloud Platform Console, open a web browser, go to the [Google Cloud including GCP & G Suite — Try Free | Google Cloud](#) webpage, and click **Sign in**. Then sign in with your Google Account.

Procedure

1. Open the Google Compute Engine browser, and then click **Snapshots**.
2. In the toolbar of the Google Cloud Platform Console, click . The Select a project dialog box appears.
3. In the Select a project dialog box, click **ALL**, and then click a project name in the Name column.
4. Select the checkbox of the snapshot you want to delete.
5. Click **Delete**.
6. In the overlay window, click **Delete** to confirm your choice.

Removing snapshots from the `gcloud` command line

For instructions, see the [Restoring and Deleting Persistent Disk Snapshots | Compute Engine Documentation | Google Cloud](#) webpage.

Preventing account access

When you subscribed to HYCU for GCP, you granted the solution (a third-party app from the perspective of Google) 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 click **Sign in**.
2. Sign in with your Google Account.
3. Click **Security**.
4. Locate the Third party apps with account access section and click **Manage third party access**.

5. Under Third-party apps with account access, click **HYCU for GCP**, and then click **REMOVE ACCESS**.
6. Click **OK** to confirm revocation of the access permission.

For general information on permissions to access your Google Account, see the [Third-party sites & apps with access to your account - Google Account Help](#) webpage.

Canceling HYCU for GCP subscription

Once you unsubscribe from HYCU for GCP, you lose data protection for your Google Cloud Platform projects that are covered by your subscription.

Prerequisites

- You are signed in to Google with a Google Account that is granted the Billing Account Administrator (`roles/billing.admin`) role on the billing account of the HYCU for GCP subscription.
- Your currently selected project in the Google Cloud Platform Console is linked to the billing account of the HYCU for GCP subscription.

Procedure

1. Open a web browser and go to the [HYCU | Marketplace - Google Cloud Platform](#) webpage.
2. Click **Cancel service**.
3. In the Cancel HYCU subscription dialog box, click **CANCEL SUBSCRIPTION** to confirm your choice.

After you cancel your HYCU for GCP subscription, your data protection configuration is kept for 14 days before it is permanently deleted. Data protection configuration comprises backup policies, credential groups, data archive configuration, backup window specifications, service account assignments, event notification configuration, configured reports, and generated report versions. During this period you can still change your mind and decide that you want to resume using HYCU for GCP. In this case, resubscribe to HYCU for GCP and specify the billing account of the canceled subscription in the process.

Appendix A

Objects created by the service

To provide data protection, HYCU for GCP creates specific objects (referred to as HYCU objects) in your Google Cloud Platform projects. Some of these objects actually protect original data (snapshots, backup images, automatic buckets), some are the result of a restore (restored files on an instance disk or in a bucket), and others are auxiliary objects that exist only for the duration of a task. All these objects are visible in the Google Cloud Platform user interfaces.

Caution Apart from the restored files and unless specifically instructed to do so, never rename or delete any HYCU objects through the Google Cloud Platform Console or the `gcloud` command line.

The following table lists names or locations of objects that are created during a backup or restore task, and are preserved afterward.

Table A-1: Permanent HYCU objects

Object type
Name or location path template
Snapshot
<code>hycu-snap-<TaskID>-<DiskName></code>
Automatic bucket
<code>hycu-<CloudStorageRegionName>-<UUID></code>
Bucket folder with any of the following: backup image, copy of a backup image, data archive.
<code>hycu/backups/ <ProjectName>/<ZoneName>/<InstanceName>/disks/<DiskName>/<StorageClass></code>
Bucket folder with disk catalog, instance metadata, and instance disks' metadata.
<code>hycu/backups/<ProjectName>/<ZoneName><InstanceName>/tasks/<TaskID></code>

Renamed original file (at the original location on an instance)	<code><OriginalFileName>.hycu.orig[.<OriginalFileExtension>]</code>
Renamed restored file (at the original location on an instance)	<code><OriginalFileName>.hycu.restored[.<OriginalFileExtension>]</code>
Bucket folder with restored individual files or folders	<code>hycu/restores/<ProjectName>/<ZoneName>/<InstanceName>/<TaskID>/<DiskName>/<VolumeName>/<PathName></code>
External IP address resource automatically allocated by HYCU for GCP during cloning	<code>hycu-static-external-<UUID></code>
Internal IP address resource automatically allocated by HYCU for GCP during cloning	<code>hycu-static-internal-<UUID></code>

The following table lists names of auxiliary objects that exist only for the duration of a task.

Table A-2: Temporary HYCU objects

Object type	Name template
Temporary disk	<code>hycu-disk-tmp-<TaskID>-<OriginalDiskName></code>
Temporary instance ¹	<code>hycu-instance-tmp-<TaskType>-<TaskID>-<UUID></code>

¹ Applicable scenarios: instance rediscovery after assigning access credentials or enabling a restore of individual files or folders, backup tasks when a restore of individual files or folders is enabled for the instances, restore tasks.

Glossary

A

automatic bucket

A Google Cloud Storage bucket that is created by HYCU for GCP.

B

backup image

Primary backup data of an instance that can be used to restore the instance. It can be kept in one of the following forms: snapshot, or backup data in a bucket.

backup policy, policy

A set of rules that basically defines backup frequency for backing up your data and retention period for the backup images. It may also include a rule for creating copies of backup images and a rule for archiving protected data.

backup window

A user-defined time frame during which scheduled backup tasks are allowed to start.

backup window specification

A named set of backup windows that can be used in a backup policy.

billing account

An entity in the Google Cloud Platform service suite that defines who pays for a given set of resources. Each billing account defines a set of the Google Cloud Platform projects that are visible to a Google Account within the corresponding HYCU for GCP subscription.

billing account viewer

A Google Cloud Platform identity (a Google Account or a service account) that HYCU for GCP uses for determining which Google Cloud Platform projects are linked to the billing account.

bucket

A storage container in the Google Cloud Storage service that can store backup images, copies of backup images, and data archives created by HYCU for GCP.

C

copy of backup image

A copy of the primary backup data that is stored in a separate bucket and can use a different retention setting.

D

data archive

A daily, weekly, monthly, or yearly copy of a backup image that is retained in a bucket for future reference, usually for a longer period of time.

default backup policy

A backup policy that you define to be automatically assigned to newly identified instances in the Google Cloud Platform projects of a protection set.

disk

A persistent disk in the Google Compute Engine service that belongs to some instance. With HYCU for GCP you can restore it separately from other persistent disks.

disk catalog

A record of an instance disk's inventory that enables a restore of individual files or folders.

E

entity

An object that you can protect with a backup policy independently of others. In HYCU for GCP, such objects are instances in the Google Compute Engine service.

H

HYCU for GCP

A managed service implemented on the basis of the Google Cloud Storage and Google Compute Engine services. It provides data protection for instances in the Google Compute Engine service.

I

instance compliance

An indicator showing whether an instance is backed up in accordance with the recovery point objective (RPO) in the assigned policy.

instance protection

An indicator showing whether the instance has a policy assigned and at least one backup image of the instance exists.

instance, VM Instance

A virtual machine (VM) that is hosted in the Google Compute Engine service and for which HYCU for GCP provides data protection.

P

policy compliance

An indicator of joint compliance of all instances a policy is assigned to.

project

An entity in the Google Cloud Platform resource hierarchy which includes VM instances that you can protect with HYCU for GCP.

protected data

Data stored on a disk of a protected instance that is included in a backup.

protection set

A group of Google Cloud Platform projects which share the same data protection setup in HYCU for GCP: credential groups, backup policies, backup window specifications, data archives. Such projects may also share an assigned service account.

protection set service account

A Google Cloud Platform service account that is chosen, instead of the sign-in user account, to run certain types of tasks within a protection set. Protection set service accounts also enable specific features of HYCU for GCP.

R

recovery point objective (RPO)

A backup policy rule that defines the period between two consecutive backups for the same instance. Also referred to as backup frequency.

restore point

A state of data at backup time. The state can be recreated from a corresponding backup image, copy of backup image, or data archive. If a restore of individual files or folders is enabled at the backup time, a restore point also allows you to restore individual files or folders.

restore point entity

A snapshot, backup data in a bucket, copy of backup image, data archive, or disk catalog that belongs to a restore point.

S

sign-in user account

A Google Account that is used to sign in to the HYCU for GCP web user interface.

source

A resource which includes entities that you can protect. HYCU for GCP treats Google Cloud Platform projects as sources.

T

target

A storage location used for storing backup data. HYCU for GCP uses Google Cloud Storage buckets as targets.

W

web user interface (WUI)

A web-based graphical console that provides access to HYCU for GCP. It enables activities such as configuration and maintenance of your data protection environment, scheduling and starting backup tasks, monitoring task progress, and browsing the event log.

HYCU Customer Support and information

Use the communication channels listed in this section if you need:

- Help with the HYCU for GCP subscription process
- Assistance while using HYCU for GCP
- Additional information about HYCU for GCP
- Information about other HYCU products and services

Customer Support

Should you require additional information or assistance while using the service, contact the vendor that arranged its subscription for you.

If you have subscribed to the service yourself, and are experiencing a problem, search for a solution on the [HYCU Customer Support](#) webpage. In the absence of an article addressing your problem, ask HYCU Customer Support for assistance: on the webpage, sign in with a valid user account, click **Submit a request**, and then fill in the request form. You should have received user account information by email after subscribing to the service.

Important: Before submitting a request to HYCU Customer Support, collect troubleshooting information. For a list of the relevant pieces of information, check troubleshooting sections in the service documentation.

Company resources on the web

For more information about our company and other products and services in our offering, visit the [HYCU | Simplifying Multi-cloud Data Protection](#) website. For additional product- or service-related information, watch videos on the [HYCU, Inc. - YouTube](#) channel. HYCU is also present on social networks. Follow us on Twitter [🐦](#) and LinkedIn [in](#).

General information

For questions related to product or service business, purchase of HYCU products, or subscription to other HYCU services, send an email to info@hycu.com.

Feedback

For comments or suggestions about this service, including its documentation, send an email to info@hycu.com. We will be glad to hear from you!

